# DAGSTUHL REPORTS

**Volume 13, Issue 7, July 2023**



Dagstuhl Reports, Vol. 13, Issue 7 ISSN 2192-5283

**ISSN 2192-5283**

*Published online and open access by*
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at
https://www.dagstuhl.de/dagpub/2192-5283

*Publication date*
March, 2024

*Bibliographic information published by the Deutsche Nationalbibliothek*
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at https://dnb.d-nb.de.

*License*
This work is licensed under a Creative Commons Attribution 4.0 International license (CC BY 4.0).
In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:
- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Digital Object Identifier: 10.4230/DagRep.13.7.i

*Aims and Scope*
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:
- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

*Editorial Board*
- Elisabeth André
- Franz Baader
- Daniel Cremers
- Goetz Graefe
- Reiner Hähnle
- Barbara Hammer
- Lynda Hardman
- Oliver Kohlbacher
- Steve Kremer
- Rupak Majumdar
- Heiko Mantel
- Albrecht Schmidt
- Wolfgang Schröder-Preikschat
- Raimund Seidel (*Editor-in-Chief*)
- Heike Wehrheim
- Verena Wolf
- Martina Zitterbart

*Editorial Office*
Michael Wagner *(Managing Editor)*
Michael Didas *(Managing Editor)*
Jutka Gasiorowski *(Editorial Assistance)*
Dagmar Glaser *(Editorial Assistance)*
Thomas Schillo *(Technical Assistance)*

*Contact*
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
https://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 23271

# Human in the (Process) Mines

**Claudio Di Ciccio*[1], Silvia Miksch*[2], Pnina Soffer*[3], Barbara Weber*[4], and Giovanni Meroni†[5]**

1 **Sapienza University of Rome, IT.** claudio.diciccio@uniroma1.it
2 **TU Wien, AT.** silvia.miksch@tuwien.ac.at
3 **University of Haifa, IL.** spnina@is.haifa.ac.il
4 **Universität St. Gallen, CH.** barbara.weber@unisg.ch
5 **Technical University of Denmark – Lyngby, DK.** giom@dtu.dk

## Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 23271, "Human in the (process) mines". The seminar dealt with topics that are at the intersection of process mining and visual analytics, and can potentially contribute to both areas. Process mining is a discipline blending data science concepts with business process management. It utilizes event data recorded by IT systems for a variety of tasks, including the automated discovery of graphical process models, conformance checking between data and models, enhancement of process models with additional analytic information, run-time monitoring of processes and operational support. Ultimately, the purpose of process mining is to make sense of event data and answer business and domain-related questions to support domain-specific goals. Visual Analytics, defined as "the science of analytical reasoning facilitated by interactive visual interfaces," is a multidisciplinary approach, integrating aspects of data mining and knowledge discovery, information visualization, human-computer interaction, and cognitive science to support humans in making sense of various kinds of data. While these two research disciplines face similar challenges in different contexts, there have been few interactions and cross-fertilization efforts between the respective communities so far. This Dagstuhl Seminar is intended to bring together researchers from both communities and foster joint research efforts and collaborations to advance both fields and enrich future approaches to be developed.

**Seminar** July 2–7, 2023 – https://www.dagstuhl.de/23271
**2012 ACM Subject Classification** Applied computing → Business process management; Human-centered computing → Visual analytics
**Keywords and phrases** human in the loop, process mining, visual analytics
**Digital Object Identifier** 10.4230/DagRep.13.7.1

## 1 Executive Summary

*Claudio Di Ciccio (Sapienza University of Rome, IT)*
*Silvia Miksch (TU Wien, AT)*
*Pnina Soffer (University of Haifa, IL)*
*Barbara Weber (Universität St. Gallen, CH)*

**License** Creative Commons BY 4.0 International license
© Claudio Di Ciccio, Silvia Miksch, Pnina Soffer, and Barbara Weber

This summary provides an overview of the outcomes of our Dagstuhl Seminar *Human in the (Process) Mines*. It began with a general introduction to the aim, scope and context of the Dagstuhl Seminar. The preliminary presentation was followed by a sequence of two-minute

* Editor / Organizer
† Editorial Assistant / Collector

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

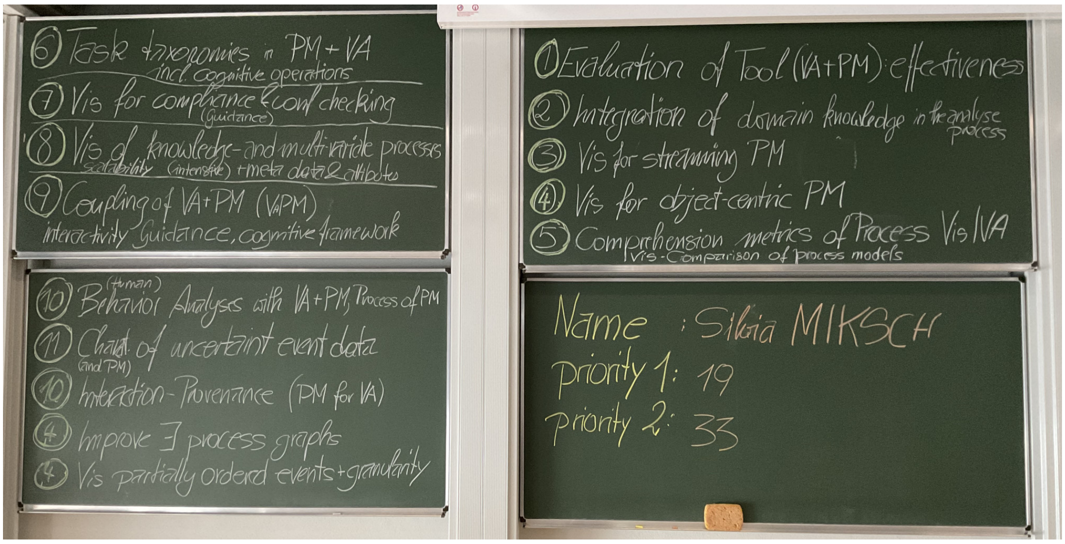


**Figure 1** Our brainstorming session.

speeches, with which every participant could give a brief overview of their own background, expertise and personal expectations for the seminar. The seminar participants included experts from both the process mining and visual analytics communities, with industry representatives and researchers in academia at different levels of seniority. The following four talks then set the scene to establish a shared understanding of the two research areas' core concepts, employed methods and pursued aims: "Process Mining and Visual Analytics: A Story of Twins separated at Birth," by Jan Mendling; "Visual Analytics for Time-Oriented Data" by Wolfgang Aigner; "The Interaction Side of Visual Analytics" by Christian Tominski; "The Process of Process Mining: Looking into Process Mining through the Lens of Process Analysts," by Francesca Zerbato. The reader can find the abstracts of those talks below. Thereupon, a plenary brainstorming session began, moderated by the seminar organizers. Every participant in the assembly could suggest a topic that they would like to investigate together with the other attendees. The whole assembly decided how to distill and join the topics together. Figure 1 shows the set of fourteen arguments gathered from the discussion, then merged into eleven points (see the right-top, left-top, and bottom-left quadrants in the figure). To form the working groups, every attendee was asked to write on a card their name and the top two topics they would like to discuss further, in order of priority (see Fig. 1 in the bottom-right corner). The organizers collected the cards and divided the participants into groups of five to six people based on the preferences declared on the cards while keeping the teams heterogeneous in terms of background, research interests and affiliation. To this end, the suggested topics underwent another merging round in order to draw out six key research questions at the intersection of process mining and visual analytics. The questions and the main results that were achieved follow.

**How to incorporate knowledge into an interactive analysis process of event data?** Addressing this question resulted in a conceptual model of knowledge-assisted interactive visual process mining, which extends an established model in the area of visual analytics. To validate this model, it has been instantiated for an example process mining task.

**How to make sense of multi-faceted process data?** This question has guided the development of an interactive framework called "Tiramisù". Based upon the interconnection of multiple visual layers, it displays process information under different perspectives and projects them onto a domain-friendly representation of the context in which the process unfolds. The feasibility of the framework was demonstrated through its application in use-case scenarios. A paper summarizing these results has been submitted to a workshop co-located with the 5th International Conference on Process Mining (ICPM 2023).

**What comprehension metrics can be used to evaluate visualizations for process mining tasks?** As a starting point for addressing this question, the ICE-T hierarchical evaluation framework was selected to adapt its metrics for the process mining domain. Adapting the methodological approach included the selection of process mining tasks and the use of the ICE-T framework for evaluation. An initial evaluation using the Disco process mining software and the road traffic fine management process data set was performed.

**How to select appropriate visualizations for object-centric and multi-dimensional process mining tasks?** Addressing this question has focused on investigating and developing a conceptual categorization that facilitates a structured design process of multidimensional visual process mining artefacts grounded in multi-faceted visual analytics design. Existing process mining visualizations have been mapped onto well-established visual analytics facets.

**How to design visualizations for the task of conformance checking?** This question deals with providing visual support to a major process mining task, of assessing conformance between a process model and the behaviour captured in event data, measuring it, spotting deviations, and analyzing them. Addressing this question, gaps in existing conformance-checking visualizations were analyzed. Then, using a visual analytics workflow, specific application scenarios of conformance checking were analyzed.

**How to support human-centred process mining by constructing a unified process mining and visual analytics task taxonomy?** Aiming to realize opportunities at the intersection of human-centred process mining activities, the various cognitive aspects of these activities, and visual ways of supporting them, a taxonomical approach was taken. The objective of the proposed taxonomy was to establish a common vocabulary and shared understanding across the three bodies of knowledge, serve as a point of reference, and inform measurements for evaluating process visualizations.

From the second day on, time was mostly devoted to working groups, each discussing one of the research questions that were raised in breakout sessions. This activity was accompanied by "lightning talks", i.e., brief presentations on issues that emerged during the discussions: "Process mining: A hands-on session" by Anti Alman, "Extended Cognition with Big Data and AI" by Irit Hadar, "How Visual Analytics People Work: Design Study Methodology in Practice" by Tatiana von Landesberger, and "Coping with Volume and Variety in Temporal Event Sequences" by Jan Mendling. Periodic intermediate group presentations reported on the advancement of the teams' work. The final presentations of the results achieved by the working groups and a discussion of the continuation of the freshly established collaborations took place on the final day. The seminar was concluded with the plan for a follow-up Dagstuhl Seminar, having Pnina Soffer, Katerina Vrotsou, Claudio Di Ciccio, and Christian Tominski as the proponents and prospective organizers. In the remainder of this report, the talks held on the first day of the *Human in the (Process) Mines* Dagstuhl Seminar and the final working group reports follow.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Visual Analytics for Time-Oriented Data

*Wolfgang Aigner (FH – St. Pölten, AT)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Wolfgang Aigner
**Main reference** Wolfgang Aigner, Silvia Miksch, Heidrun Schumann, Christian Tominski: "Visualization of Time-Oriented Data", Springer, 2023
**URL** https://browser.timeviz.net

Visual Analytics (VA) has been defined as the science of analytical reasoning supported by interactive visual interfaces. This hints towards the goal of supporting human users in making sense of large and complex data. In order to achieve that, VA intertwines computers and humans by combining automated analysis with interactive visualization methods. Processes and sequences are concepts inherently tied to the time dimension. Therefore, the focus of this talk is on Visual Analytics for time-oriented data. In particular, this includes aspects of modeling of time and time-oriented data, their visualization, as well as interactivity.

## 3.2 Process Mining and Visual Analytics: A Story of Twins separated at Birth

*Jan Mendling (HU Berlin, DE)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Jan Mendling
**Joint work of** Jan Mendling, Anton Yeshchenko
**Main reference** Anton Yeshchenko, Jan Mendling: "A Survey of Approaches for Event Sequence Analysis and Visualization using the ESeVis Framework", CoRR, Vol. abs/2202.07941, 2022.
**URL** https://arxiv.org/abs/2202.07941

In this talk I elaborate on the relationship and intersection of process mining and visual analytics research. Partially, I refer to a recent survey by Anton Yeshchenko and Jan Mendling on Approaches for Event Sequence Analysis and Visualization using the ESeVis Framework (https://arxiv.org/abs/2202.07941).

## 3.3 The Interaction Side of Visual Analytics

*Christian Tominski (Universität Rostock, DE)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Christian Tominski
**Main reference** Christian Tominski, Heidrun Schumann:: "Interactive Visual Data Analysis," AK Peters Visualization Series, CRC Press 2020, ISBN 9781498753982.
**URL** https://doi.org/10.1201/9781315152707

Visual Analytics combines computational, visual, and interactive methods to help users gain insight into large and complex data. In this talk, I highlight the importance of interactivity for the knowledge generation process. In the first part of the talk, I explain basic models and notions of interaction, including Shneiderman's information seeking mantra, Keim's visual analytics mantra, Norman's action cycle, Yi's interaction intents as well as direct manipulation and fluid interaction as defined by Shneiderman and Elmqvist. The second part of the talk showcases several techniques for interacting with graph structures, mostly from my own work, including responsive matrix cells for analyzing and editing multivariate graphs, interactive lenses for graph exploration, naturally-inspired visual comparison, and fluid in-situ unfolding of edges. I also briefly touch upon guidance in visual analytics, progressive visual analytics, and multi-modal interaction. I close with advocating human-data interaction as a key ingredient of visual analytics.

## 3.4 The Process of Process Mining: Looking into Process Mining through the Lens of Process Analysts

*Francesca Zerbato (Universität St. Gallen, CH)*

**License** Creative Commons BY 4.0 International license
© Francesca Zerbato
**Joint work of** Francesca Zerbato, Pnina Soffer, Barbara Weber, Lisa Zimmermann, Elizaveta Sorokina, Irit Hadar
**Main reference** Francesca Zerbato, Pnina Soffer, Barbara Weber: "Process Mining Practices: Evidence from Interviews", in Proc. of the Business Process Management – 20th International Conference, BPM 2022, Münster, Germany, September 11-16, 2022, Proceedings, Lecture Notes in Computer Science, Vol. 13420, pp. 268–285, Springer, 2022.
**URL** https://doi.org//10.1007/978-3-031-16103-2_19

Over the years, process mining research has mainly emphasized the development of automated techniques, giving less attention to the crucial role of (human) analysts involved in the analysis process. Research on the Process of Process Mining (PPM) aims to uncover how process analysts do their work in practice, i.e., how they act in the process mining tools and how they reason as the analysis unfolds. This talk provides an overview of recent research on the PPM, describing how the combination of different modalities of data, such as interaction traces and verbal data, can be combined to gain insights into the behavior and work practices of process analysts. The talk concludes by presenting two illustrative examples of analyzing such data. The first example demonstrates a bottom-up analysis approach utilizing interaction trace data to reveal what visualizations are used during the analysis, how they are used and why. The second example introduces a novel cognitive model derived from the theory of Prediction Error Minimization, which offers a top-down analytical perspective on the behavior of process analysts and their cognitive processes.

**Acknowledgments.** B. Weber, F. Zerbato and L. Zimmermann are supported by the ProMiSE project funded by the SNSF under Grant No.: 200021_197032.

# 4 Working groups

## 4.1 Visual Process Mining: from Interaction to Knowledge

*Wolfgang Aigner (FH – St. Pölten , AT), Chiara Di Francescomarino (University of Trento, IT), Daniel Schuster (Fraunhofer FIT – Sankt Augustin, DE), Cagatay Turkay (University of Warwick – Coventry, GB), and Francesca Zerbato (Universität St. Gallen, CH)*

**License** Creative Commons BY 4.0 International license
© Wolfgang Aigner, Chiara Di Francescomarino, Daniel Schuster, Cagatay Turkay, and Francesca Zerbato

Process mining (PM) techniques enable organizations to extract valuable insights and generate knowledge from the event data captured by their information systems. Visual analytics (VA) aims to intertwine analytical reasoning by humans based on visual perception with the processing power of computers for automated data analysis [5]. Both fields share the aim of generating knowledge based on data, but have developed rather independently from each other. Process mining research has emphasized the development of automated techniques, giving less attention to the role of (human) process analysts involved in the analysis [6]. Visual analytics literature focuses mostly on event data [3] and have paid little attention to the underlying process models and lack approaches to address some common tasks in process mining. There is therefore significant potential for the two communities to learn from each other and develop synergies across concepts, theories, methods, and artifacts.

This working group focused on studying the role of the human together with the integration of knowledge in the overall process. Presently, only a few process mining techniques incorporate knowledge alongside event data, which is the primary input for most techniques [4, 2]. For instance, the work by Schuster et al. [4] reviews domain knowledge utilization in process discovery. However, few knowledge-utilizing approaches exist compared to automated discovery techniques, indicating the need for further research.

Our group, comprising process mining and visualization researchers, started with exploring existing process mining and visual analytics models and approaches to develop a shared language across the team. We set out to bring together established models and techniques from the visual analytics and process mining literature with an objective to explore an interactive and visual approach to doing process mining that puts humans and domain knowledge at its core. We narrowed down our emphasis to conceptual frameworks that emphasize knowledge and how it is elicited and externalized through the use of visual analytics and process mining techniques.

Specifically, we based our work on a conceptual model of knowledge-assisted visual analytics (KAVA) [1] that brings together visualization, automated analysis, analytical reasoning of human users, and the knowledge created based on these. Central to the KAVA model is the utilization of visualizations and interaction to extract knowledge from data.

Building upon the KAVA model, we focused on adapting it to the process mining domain. To this end, we evaluated various common process mining tasks such as process discovery, conformance checking, temporal performance analysis, predictive process mining, comparative process mining, and action-oriented process mining in the context of the model. As a result, we derived a conceptual model that allows us to understand how different process mining approaches involve and consider humans and how process analysts generate and externalize knowledge through their interaction with process mining tools and techniques.

Figure 2 reports the derived conceptual model.

Going beyond the classical PM approaches that automatically learn from data, the model focuses on the user-machine system, and especially on how the human user visualizes (V), interacts (E), as well as generates (through perceptions and interactions) and externalizes (X) tacit knowledge in PM tasks.

We derived the model in several iterations. First, in dialogue with the process mining literature, we started untangling and interrogating the different components of the KAVA model. Our goal with this exercise was twofold: 1) to understand the model's fit and coverage for process mining, and 2) to revise the model as a descriptive framework for a knowledge-centric interactive and visual process mining approach.

Then, we explored the literature in both visual analytics and process mining to associate techniques, concepts and terms to each of the components in the model, e.g., we listed automated data processing techniques from both areas specifically designed for event-based and process data. We then focused our attention on the transitions between these elements and added new transitions and links to strengthen how well the model underpins our envisioned knowledge-centric interactive and visual process mining approach.

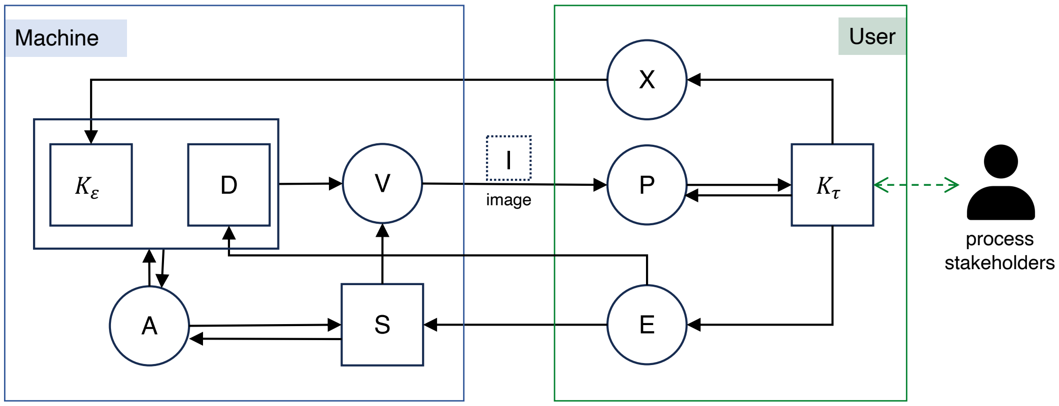


**Figure 2** Our knowledge-assisted interactive visual process mining model (D...data; V...visualization; I...image; P...perception; $K_\tau$...tacit knowledge; E...exploration/manipulation; S...specification; A...automated analysis; X...knowledge externalization; $K_\epsilon$...explicit knowledge). Boxes depict artifacts, while circles represent transformations. (Figure partly adapted from [1]).

Once we established a first version of the model, our next exercise was to go over a range of process mining sessions as a case study and explore how comprehensively our model captures the process mining process. This enabled us to put our model to test. This exercise resulted in a number of revisions which led to the model we introduced in Figure 2. We concluded our working group activities by identifying a range of process mining tasks that will enable us to more systematically evaluate our model in the next steps.

Summing up, the discussions of the working group resulted in two main outcomes: 1) a conceptual model for interactive visual process mining that takes a knowledge-centric perspective; 2) an instantiation of the model on a real PM task so as to validate the suitability of the model to the field, as well as to identify challenges and opportunities for both communities.

Going forward, the working group will further refine the proposed model with a goal to better understand the role of tacit and explicit knowledge in the context of different process mining tasks. We will use the model as a vehicle to identify challenges, open problems and opportunities in supporting users in generating and externalizing knowledge during visual process mining. Through these we envision to bring the process mining and visual analytic communities together along a shared research agenda on advancing interactive and visual process mining.

**Acknowledgments.** F. Zerbato is supported by the ProMiSE project funded by the SNSF under Grant No.: 200021_197032.

## References

1 Paolo Federico, Markus Wagner, Alexander Rind, Albert Amor-Amoros, Silvia Miksch, and Wolfgang Aigner. The role of explicit knowledge: A conceptual model of knowledge-assisted visual analytics. In Brian D. Fisher, Shixia Liu, and Tobias Schreck, editors, *12th IEEE Conference on Visual Analytics Science and Technology, IEEE VAST 2017, Phoenix, AZ, USA, October 3-6, 2017*, pages 92–103. IEEE Computer Society, 2017.

2 Chiara Di Francescomarino, Chiara Ghidini, Fabrizio Maria Maggi, Giulio Petrucci, and Anton Yeshchenko. An eye into the future: Leveraging a-priori knowledge in predictive business process monitoring. In Josep Carmona, Gregor Engels, and Akhil Kumar, editors, *Business Process Management – 15th International Conference, BPM 2017, Barcelona, Spain, September 10-15, 2017, Proceedings*, volume 10445 of *Lecture Notes in Computer Science*, pages 252–268. Springer, 2017.

**3** Phong H. Nguyen, Rafael Henkin, Siming Chen, Natalia V. Andrienko, Gennady L. Andrienko, Olivier Thonnard, and Cagatay Turkay. VASABI: hierarchical user profiles for interactive visual user behaviour analytics. *IEEE Trans. Vis. Comput. Graph.*, 26(1):77–86, 2020.
**4** Daniel Schuster, Sebastiaan J. van Zelst, and Wil M. P. van der Aalst. Utilizing domain knowledge in data-driven process discovery: A literature review. *Comput. Ind.*, 137:103612, 2022.
**5** James J. Thomas and Kristin A. Cook. A visual analytics agenda. *IEEE Computer Graphics and Applications*, 26(1):10–13, 2006.
**6** Francesca Zerbato, Pnina Soffer, and Barbara Weber. Process mining practices: Evidence from interviews. In Claudio Di Ciccio, Remco M. Dijkman, Adela del-Río-Ortega, and Stefanie Rinderle-Ma, editors, *Business Process Management – 20th International Conference, BPM 2022, Münster, Germany, September 11-16, 2022, Proceedings*, volume 13420 of *Lecture Notes in Computer Science*, pages 268–285. Springer, 2022.

## 4.2 Towards a visual analytics framework for sensemaking of multi-faceted process information

*Anti Alman (University of Tartu, EE), Alessio Arleo (TU Wien, AT), Iris Beerepoot (Utrecht University, NL), Andrea Burattin (Technical University of Denmark – Lyngby, DK), Claudio Di Ciccio (Sapienza University of Rome, IT), and Manuel Resinas (University of Sevilla, ES)*

**License** 🅭 Creative Commons BY 4.0 International license
© Anti Alman, Alessio Arleo, Iris Beerepoot, Andrea Burattin, Claudio Di Ciccio, and Manuel Resinas

### 4.2.1 Introduction

Process mining (PM) is the discipline aimed at extracting information from events recorded by information systems executing processes [22]. The operational context of process mining is multi-variate in that data potentially stems from multiple sources and pertains to diverse dimensions (control flow, time, resources, object lifecycle). As argued by the seminal work of Beerepoot et al. [2], a critical, yet largely unaddressed, issue of process mining is the fixed-granularity level of process analysis, namely the inability to navigate through distinct, and possibly domain-specific, dimensions. This problem makes the investigation of process improvement less effective, as studied by Kubrak et al. [13], as it entails a partial view over an inherently complex search space, which is typical of knowledge-intensive processes [5, 6]. Visual Analytics (VA) is the discipline of supporting users' analytical reasoning through the use of interactive visual interfaces [12]. VA is intended to keep the user *within* the analysis loop, so that human comprehension and perception actively contribute to generating new insights and increasing confidence with the analysis results. VA has extensively studied the problem of representing complex, multi-faceted phenomena (see, e.g., [11, 16]), which is also the type of phenomena often encountered within real-life business processes.

Thus far, there has been limited cooperation between PM and VA. In this paper, we argue that VA can play a pivotal role in addressing the above limitation of fixed-granularity in process mining. With a cross-pollinating effect, PM can equip VA with a collection of established algorithms and techniques for the automated generation of process-oriented representations of system dynamics.

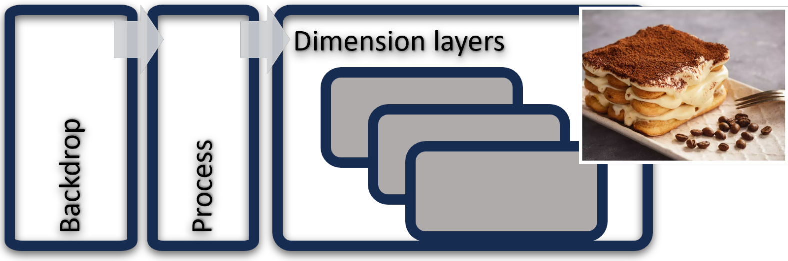


**Figure 3** An exemplification of the `Tiramisù` framework. On the right, a tiramisù.

During the seminar, we envisioned a novel framework that resorts to VA for the interactive investigation of factual evidence of process insights revolving around information mined from data sources that go under the name of *sequence event data*, to use VA's terminology, or *event logs*, as per the PM nomenclature. Our framework will follow a multi-layer approach, providing end users with a context-aware visualization integrating classical PM representation elements (such as workflow nets [19], directly-follows graphs [21], or declarative process maps [7]) with additional diagrams and visual cues tailored for context-variable and metadata representations (e.g., timelines and calendars for time, geographical or building maps for space). We introduce a couple of use cases that could benefit from such a framework. The interconnection of the different layers and their anchorage to a *backdrop* onto which the representations are projected provides the user with means to foster explainability for process analysis and, eventually, enhancement.

In the remainder of this section, Sec. 4.2.2 describes our framework, illustrating its key concepts and the interplay of its components. Section 4.2.3 showcases two use cases demonstrating the application of the `Tiramisù` approach in the area of knowledge-intensive processes. Section 4.2.4 acknowledges some known limitations of our work, paving the path for future work discussed in Sec. 4.2.5.

### 4.2.2 The `Tiramisù` framework

In `Tiramisù` we envision a visualization framework designed to support sensemaking when dealing with complex PM event sequences. Our goal is to augment existing process models with additional dimensions that can provide further context during the analysis, easing the generation of insights and generally providing a more comprehensive understanding of the process and phenomenon under investigation. With this goal in mind, we envision `Tiramisù` as a multi-level framework (see Fig. 3), reminiscent of the popular *tiramisù* cake. The framework will be structured with a "backdrop", a main layer with the process model, and one or more dimension layers. These are meant to be superimposed on the first two, in a details-on-demand fashion, thus reminding the different layers of a *tiramisù*.

### 4.2.3 Use-case scenarios

In this section, we illustrate some scenarios that motivate our work. We focus on two classes of knowledge-intensive processes, pertaining to healthcare [15] and personal information management [4].

#### 4.2.3.1 Behavioral deviation analysis in healthcare

Healthcare is one of the most difficult, but at the same time, one of the most promising domains to tackle both in the field of PM [15] as well as in the area of VA [3]. As such, healthcare is a natural domain for applying the `Tiramisù` approach. To demonstrate this, we have chosen the specific task of analyzing the sleeping routine of patients suffering from dementia or other similar diseases. Naively applying the already existing process mining approaches to this task can, for example, lead to relying on simple process maps where the nodes represent various activities performed by the patients, and the arcs represent dependencies or other temporal constraints among these activities [8]. While the above models may be sufficient for a process mining expert to extrapolate meaningful process insights, it can quickly become challenging for domain experts, who may lack training in interpreting these formalisms: as discussed in [8, 10], the most likely users of such a system would, in practice, not be a process expert, but the doctors and the nurses responsible for the care of the specific patient.

#### 4.2.3.2 Personal productivity analysis in work processes

Personal information management (PIM) pertains to the organization of one's own activities, contacts, etc., through the use of software on laptops and smart devices. Similarly, personal informatics systems resort to one individual's own information to pursue the objective of aiding people to collect and reflect on their personal information [14]. Several techniques can be used to collect personal information like non-participant observation, screen recording, and timesheet techniques, each with their own advantages and disadvantages [17]. Regardless of the technique used, collected personal information can be seen as an event log, which can be analyzed using process mining techniques to discover personal work processes.

Depending on the characteristics of the work, personal work processes can be knowledge-intensive and significantly unstructured, which means they present the aforementioned challenges. In this use case, we focus on the personal work processes performed by an academic during her daily work, which involves conducting research, preparing lessons, grading students or reviewing research papers, among many other activities.

Specifically, our focus is on the retrospective analysis of the influence of the personal work processes a person has followed during a certain period of time on the positive or negative outcome of a task. For instance, an academic who has missed the deadline for submitting a review to a conference might be interested in knowing what happened during the period of time in which she was working in the review that caused the delay.

### 4.2.4 Limitations

Evidence-based discovery approaches suffer from the potential bias given that only a part of processes (especially the knowledge-intensive ones) are recorded by computerized systems [9]. Our framework is in this sense no different. The injection of additional facts stemming from domain experts and users is complementary to this investigation and paves the path for future work (see Section 4.1 for further details in this regard).

Also, we assume the user knows what visualizations best suit their needs, although this may be a challenge on its own. Investigating this aspect goes beyond the scope of this report. However, the interested reader is referred to the work of Sirgmets et al. [18], which presents a framework aimed to guide design choices for the effective visualization of analytical data.

Data-based process analysis tasks often involve a significant amount of time for extracting, reformatting, and filtering event logs from information systems [20]. Our framework requires these preliminary operations too, with the addition that oftentimes knowledge-intensive processes record executions over a heterogeneous set of applications and devices in partially structured or unstructured formats [6]. Furthermore, the notion of a process instance (case) tends to be less defined in such contexts, thus requiring a prior customizable reconciliation of shared events [1]. The emergence of the novel and event data meta-models that are less centered around the concept of case could be beneficial to the information processing we envision in `Tiramisù` [23].

Finally, our framework covers the visualization aspects related to VA in PM. However, the *interaction* aspects are only partly discussed, leaving a gap which further iteration of this framework should address. This would be made easier and more effective if a *task taxonomy* about VA in PM were made available, opening a further research direction with the potential to bring together the two disciplines.

### 4.2.5 Conclusions and future remarks

In this report, we have envisioned `Tiramisù`, a framework based on the multi-layered representation of mined process information helping the user navigate the multi-faceted information at hand while keeping the data under the focus consistently linked and navigable across different dimensions.

We foresee the following research endeavors in the future. A further elaboration of the framework and its application to the use case scenarios identified. An evaluation based on an empirical study involving users is in our plans to assess the efficacy and effectiveness of our solution. The implementation of a working prototype is crucial to this extent and is part of our agenda. We envision the customization and refactoring of layers, their aggregation rules, and backdrops as key challenges to be addressed in the future. Finally, we observe that stepping from a single-backdrop multi-layer design to a tree-like hierarchical structure of backdrops could initiate a new path toward a significant extension of the expressive richness guaranteed by our framework.

**Acknowledgments.** The work of A. Alman was supported by the European Social Fund via "ICT programme" measure and by the Estonian Research Council grant PRG1226. The work of C. Di Ciccio was supported by the Italian Ministry of University and Research (MUR) under PRIN grant B87G22000450001 (PINPOINT) and by project SERICS (PE00000014) under the NRRP MUR program funded by the EU-NGEU. The work of M. Resinas was partially supported by projects PID2021-126227NB-C21/ AEI/10.13039/501100011033/ FEDER, UE and TED2021-131023B-C22/ AEI/10.13039/501100011033/ Unión Europea NextGenerationEU/PRTR.

**References**

1 Dina Bayomie, Claudio Di Ciccio, and Jan Mendling. Event-case correlation for process mining using probabilistic optimization. *Inf. Syst.*, 114:102167, 2023.
2 Iris Beerepoot et al. The biggest business process management problems to solve before we die. *Comput. Ind.*, 146:103837, 2023.
3 Jesus J. Caban and David Gotz. Visual analytics in healthcare – opportunities and research challenges. *J. Am. Medical Informatics Assoc.*, 22(2):260–262, 2015.
4 Tiziana Catarci, A.J. Dix, Akrivi Katifori, Giorgos Lepouras, and Antonella Poggi. Task-centred information management. In *DELOS Conference*, volume 4877 of *Lecture Notes in Computer Science*, pages 197–206. Springer, 2007.
5 Jochen De Weerdt, Annelies Schupp, An Vanderloock, and Bart Baesens. Process mining for the multi-faceted analysis of business processes - A case study in a financial services organization. *Comput. Ind.*, 64(1):57–67, 2013.

6 Claudio Di Ciccio, Andrea Marrella, and Alessandro Russo. Knowledge-intensive processes: Characteristics, requirements and analysis of contemporary approaches. *J. Data Semant.*, 4(1):29–57, 2015.

7 Claudio Di Ciccio and Marco Montali. Declarative process specifications: Reasoning, discovery, monitoring. In Wil M. P. van der Aalst and Josep Carmona, editors, *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*, pages 108–152. Springer, 2022.

8 Gemma Di Federico and Andrea Burattin. Do you behave always the same? In Marco Montali, Arik Senderovich, and Matthias Weidlich, editors, *Process Mining Workshops*, pages 5–17, Cham, 2023. Springer Nature Switzerland.

9 Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A. Reijers. *Fundamentals of Business Process Management, Second Edition*. Springer, 2018.

10 Gemma Di Federico, Andrea Burattin, and Marco Montali. Human behavior as a process model: Which language to use? In Andrea Marrella and Daniele Theseider Dupré, editors, *Proceedings of the 1st Italian Forum on Business Process Management co-located with the 19th International Conference of Business Process Management (BPM 2021), Rome, Italy, September 10th, 2021*, volume 2952 of *CEUR Workshop Proceedings*, pages 18–25. CEUR-WS.org, 2021.

11 Steffen Hadlak, Heidrun Schumann, and Hans-Jörg Schulz. A survey of multi-faceted graph visualization. In *EuroVis (STARs)*, pages 1–20, 2015.

12 Daniel Keim, Gennady Andrienko, Jean-Daniel Fekete, Carsten Görg, Jörn Kohlhammer, and Guy Melançon. *Visual analytics: Definition, process, and challenges*. Springer, 2008.

13 Kateryna Kubrak, Fredrik Milani, and Alexander Nolte. A visual approach to support process analysts in working with process improvement opportunities. *Business Process Management Journal*, 29(8):101–132, 2023.

14 Ian Li, Anind K. Dey, and Jodi Forlizzi. A stage-based model of personal informatics systems. In *CHI*, pages 557–566. ACM, 2010.

15 Jorge Munoz-Gama, Niels Martin, et al. Process mining for healthcare: Characteristics and challenges. *J. Biomed. Informatics*, 127:103994, 2022.

16 Renata Georgia Raidou. Visual analytics for the representation, exploration, and analysis of high-dimensional, multi-faceted medical data. *Biomedical Visualisation: Volume 2*, pages 137–162, 2019.

17 Tea Sinik, Iris Beerepoot, and Hajo A. Reijers. A peek into the working day: Comparing techniques for recording employee behaviour. In *Research Challenges in Information Science: Information Science and the Connected World – 17th International Conference, RCIS 2023*, volume 476 of *Lecture Notes in Business Information Processing*, pages 343–359. Springer, 2023.

18 Marit Sirgmets, Fredrik Milani, Alexander Nolte, and Taivo Pungas. Designing process diagrams – A framework for making design choices when visualizing process mining outputs. In Hervé Panetto, Christophe Debruyne, Henderik A. Proper, Claudio Agostino Ardagna, Dumitru Roman, and Robert Meersman, editors, *OTM*, volume 11229 of *Lecture Notes in Computer Science*, pages 463–480. Springer, 2018.

19 Wil M. P. van der Aalst. The application of petri nets to workflow management. *Journal of Circuits, Systems, and Computers*, 8(1):21–66, 1998.

20 Wil M. P. van der Aalst. *Process Mining – Data Science in Action, Second Edition*. Springer, 2016.

21 Wil M. P. van der Aalst. Foundations of process discovery. In Wil M. P. van der Aalst and Josep Carmona, editors, *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*, pages 37–75. Springer, 2022.

22 Wil M. P. van der Aalst and Josep Carmona, editors. *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*. Springer, 2022.
23 Moe Thandar Wynn, Julian Lebherz, et al. Rethinking the input for process mining: Insights from the XES survey and workshop. In Jorge Munoz-Gama and Xixi Lu, editors, *ICPM workshops*, volume 433 of *Lecture Notes in Business Information Processing*, pages 3–16. Springer, 2021.

## 4.3 Comprehension Metrics for Process Mining Visualizations

*Axel Buehler (Olympus – Hamburg, DE), Irit Hadar (University of Haifa, IL), Monika Malinova Mandelburger (Wirtschaftsuniversität Wien, AT), Andrea Marrella (Sapienza University of Rome, IT), Silvia Miksch (TU Wien, AT), and Shazia Sadiq (University of Queensland – Brisbane, AU)*

**License** Creative Commons BY 4.0 International license
© Axel Buehler, Irit Hadar, Monika Malinova Mandelburger, Andrea Marrella, Silvia Miksch, and Shazia Sadiq

The working group explores the evaluation of visualizations in process mining. Process mining is a valuable tool for analyzing and improving business processes based on event logs (event sequences). Our approach emphasizes the cognitive load involved in the process of process mining and the need for effective visualizations to support analysts in solving process mining tasks. The research question addressed is: What comprehension metrics can be used to evaluate visualizations for process mining tasks? The ICE-T hierarchical evaluation framework is selected as a starting point towards adapting its metrics for the process mining domain. We outline the methodological approach, including the selection of process mining tasks and the use of the ICE-T framework for evaluation. We perform the evaluation using the Disco process mining software and the road traffic fine management process data set. We conclude with the first prototypical evaluation results, highlighting the need for customization and validation of the ICE-T metrics for the process mining domain and outlining future steps for this research.

### 4.3.1 Motivation and Introduction

Process mining enables process analysts to provide insights into the business processes of organizations. Process mining tools use event logs as data inputs in order to make sense of the organization's processes. An event log includes data that has been generated by various systems within an organization. Process mining aims to discover, analyze, monitor and ultimately improve business processes, all based on the event log. The process of process mining is a highly interactive process. The general aim faced by analysts is to make sense of the process data, where the data serves as the input coming from the external world. This sense-making process entails an iterative cycle of data exploration toward revealing insights about the process investigated. This process entails a high cognitive load, and as such it requires high cognitive effort. Process analysts are supported by tools that enable the generalization of various visualizations of the business processes as they are executed in real time. Appropriate visualization can reduce cognitive load, and in fact serve as an extended cognition mechanism supporting and enhancing the cognitive process [11].

However, for a visualization to enable this cognitive benefit, we need to ensure that it is a good fit for both the task and the user. We therefore ask the following research question: What comprehension metrics can be used for effectively evaluating visualizations for solving process mining tasks?

In order to answer this question, we looked for relevant metrics from the field of visualization and visual analytics, and found the ICE-T hierarchical evaluation framework for evaluating visualizations as a good starting point. Our aim is to adapt this framework for the process mining domain and validate it as reliable metrics for evaluating the comprehension of a process mining visualizations.

### 4.3.2 Methodological Approach

#### 4.3.2.1 Selected Tasks in Process Mining

As one result of the co-operation between the process mining and visual analytics communities, we require a reference evaluation framework for typical tasks performed in the context of a process mining analysis to be available for general use. Our aim for the preliminarily selected evaluation framework was therefore to adapt it to the context of process mining, i.e., a process analyst solving specific process mining tasks. Accordingly, we needed to select tasks suitable for initial testing of the suitability of the selected evaluation framework.

As this reference framework is not yet available, we reviewed and brainstormed typical tasks of a process analyst by walking through a prototypical analysis session, enriched by typical visual analysis tasks performed while designing solutions for users applicable to the process mining use case.

Typical tasks we identified are:

- Check the quality of the event log during the generation and refinement thereof.
- Detect and communicate the core structure of the process, i.e. its most frequent paths, number of cases, events, and aggregated metrics for time passing between events.
- Detect potential issues in the process (e.g., loops, non-conformance, bottlenecks, long waiting times, outliers).
- Explore if the presence of a certain activity correlates with a change in throughput time.
- Identify differing process behavior within a given event log, e.g.,
  - by clustering traces or variants by various case attributes or variant features
  - by exploring the data set to detect attributes that potentially derive segmentation
- Detect, explore, understand causes and consequences of observed process behavior

There are various existing theories and comprehension metrics for evaluating different types of visualizations:

- Cognitive Load Theory: The construct representing the mental effort that, when a learner performs a particular task, imposes itself on the learner's cognitive system.
- Cognitive Fit Theory: The visualization should fit the task at hand. It provides an explanation for performance differences among users across different presentation formats of information.
- Physics of Notation: Provides principles for modeling notations to be easily understood by their users. It highlights the semantic understanding of the symbols forming the notation.
- Representation Theory: Diagrams are better than sentential representations in terms of information comprehension and inferencing.
- Split-attention Effect: Split-source information can generate a heavy cognitive load in the process of information assimilation.
- The ICE-T: A Heuristic Approach to Value-Driven Evaluation of Visualizations/Visual Analytics approaches.

As a starting point, we selected one of the tasks for prototypical evaluation using the ICE-T framework, as detailed in Section 3.

#### 4.3.2.2 ICE-T Framework

For our evaluation, we used a heuristic-based evaluation methodology, called ICE-T. ICE-T is a hierarchical value framework consisting of four components (i.e., metrics): insight, time, essence, and confidence. Each component includes a guideline, which describes important aspects of the high-level component. Each guideline is then comprised of a small set of low-level heuristics that are designed to be actionable, rate-able statements reflecting how a visualization achieves that guideline. Evaluators who have knowledge of visualization and visual analytics should assess the design with respect to these heuristics (using 7-point Likert scale: 1-strongly disagree to 7-strongly agree, or N/A-not applicable). The component scores are a simple average of the associated mid-level guideline scores, and the mid-level guideline scores are a simple average of the ratings for the low-level heuristics. According to the ICE-T guidelines, five experts are sufficient and a visualization design is successful when the mean score exceeds five.

#### 4.3.2.3 Data Set and Process Mining Tool: Analysis of the Road Traffic Fine Management Process in Disco

The data set used to perform the prototypical evaluation is derived from the management of road-traffic violations by the Italian local police and the information system used for this purpose. The process begins with the "Create Fine" transition, and offenders can pay the fine partially or in full at different stages. If the entire amount is paid or overpaid, the fine management is closed. Notifications and penalties are involved if the offender fails to pay within a certain period. Appeals can be made, and if successful, the process ends. Otherwise, the case proceeds with the "Receive Result" step. If the offender still does not pay, the case may be handed over for credit collection. See [8] for a detailed discussion of the data set. The data set is publicly available (`https://data.4tu.nl/articles/_/12683249/1`)

### 4.3.3 First Prototypical Evaluation of ICE-T

To evaluate the effectiveness of process mining visualizations for solving process mining tasks, we first had to think of tasks that are typically solved by process analysts when using process mining, as well as the various visualizations generated by the different process mining tools. As one of the most typical uses of process mining is process discovery, we decided to take one of the most common process discovery tasks "Discover the most frequent process path" as a starting point. The process mining visualization that is mostly used by a process analyst to perform this task is the directly-follows graph (DFG). The DFG generated from the above-mentioned data set is shown in Fig. 4 and has been generated using the process mining tool Disco. We used the customized ICE-T comprehension metrics to evaluate the effectiveness of the DFG with regard to the task of discovering the most frequent process path. Five experts evaluated the DFG visualization using the four components of the ICE-T framework. From the five experts, one is a process mining practitioner, and four are researchers. The four researchers come from the process mining, conceptual modeling, and visual analytics disciplines.

As a first step, each expert evaluated the DFG from Figure 4 individually. The ICE-T framework uses a 7-point Likert scale. After the individual evaluations were done, the five experts discussed their perceived understanding of the heuristics listed for each ICE-T component, as well as the difficulty to evaluate each of the heuristics against the DFG. As a result, we were able to identify two aspects that should be modified in order to fit the purpose of evaluating process mining visualizations for solving process mining tasks. The

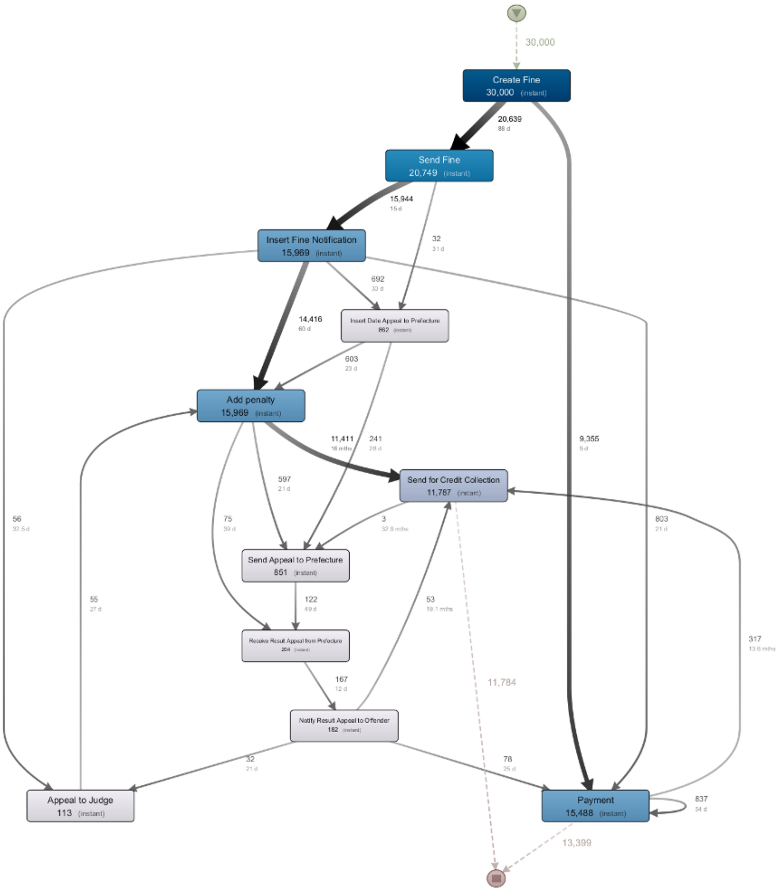


**Figure 4** The Road Traffic Fine Management Process visualization using DFG generated by Disco.

first aspect is with regard to the terminology used to describe the heuristics. Certain terms from the visual analytics community share the same semantics but are referred to differently in the process mining community. For example, "data case" in visual analytics is referred to a "process trace" in the process mining community. Therefore, we had to reformulate the terminology into the terminology used and understood by process mining analysts. Second, we found that some heuristics should be split into two heuristics, because the heuristic includes two different concepts that might lead to different evaluations of the same heuristic. We split such composed heuristics into two separate heuristics.

### 4.3.4 Next Steps

The next steps for continuing this research stream are:

- Fully customize the ICE-T metrics for the process mining domain
- Validate the customized ICE-T (prototypical)
  - Select existing process mining tasks

**Figure 5** Our lovely group.

  - Select existing and emerging visualizations from the other Working Groups of the Dagstuhl seminar to support the process mining tasks
  - Conduct the ICE-T survey for each task and visualization with 5 process mining experts
- Prepare a vision/research paper
- Generalize the customized ICE-T (experiment)

## References


1 Palash Bera. Does cognitive overload matter in understanding bpmn models? *J. Comput. Inf. Syst.*, 52(4):59–69, 2012.
2 Christian Bors. *Facilitating data quality assessment utilizing visual analytics: tackling time, metrics, uncertainty, and provenance.* PhD thesis, Wien, 2019.
3 Matthew Brehmer and Tamara Munzner. A multi-level typology of abstract visualization tasks. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2376–2385, 2013.
4 Davide Ceneda, Natalia V. Andrienko, Gennady L. Andrienko, Theresia Gschwandtner, Silvia Miksch, Nikolaus Piccolotto, Tobias Schreck, Marc Streit, Josef Suschnigg, and Christian Tominski. Guide me in analysis: A framework for guidance designers. *Comput. Graph. Forum*, 39(6):269–288, 2020.
5 Tianwa Chen, Shazia W. Sadiq, and Marta Indulska. Sensemaking in dual artefact tasks – the case of business process models and business rules. In Gillian Dobbie, Ulrich Frank, Gerti Kappel, Stephen W. Liddle, and Heinrich C. Mayr, editors, *Conceptual Modeling – 39th International Conference, ER 2020, Vienna, Austria, November 3-6, 2020, Proceedings*, volume 12400 of *Lecture Notes in Computer Science*, pages 105–118. Springer, 2020.
6 Lei Han, Tianwa Chen, Gianluca Demartini, Marta Indulska, and Shazia W. Sadiq. On understanding data worker interaction behaviors. In Jimmy X. Huang, Yi Chang, Xueqi Cheng, Jaap Kamps, Vanessa Murdock, Ji-Rong Wen, and Yiqun Liu, editors, *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 269–278. ACM, 2020.
7 Monika Malinova and Jan Mendling. Cognitive diagram understanding and task performance in systems analysis and design. *MIS Q.*, 45(4):2101–2158, 2021.

8 Felix Mannhardt, Massimiliano de Leoni, Hajo A. Reijers, and Wil M. P. van der Aalst. Balanced multi-perspective checking of process conformance. *Computing*, 98(4):407–437, 2016.
9 Jan Mendling, Mark Strembeck, and Jan Recker. Factors of process model comprehension – findings from a series of experiments. *Decis. Support Syst.*, 53(1):195–206, 2012.
10 Daniel L. Moody. The "physics" of notations: Toward a scientific basis for constructing visual notations in software engineering. *IEEE Trans. Software Eng.*, 35(6):756–779, 2009.
11 Albert Newen, Shaun Gallagher, and Leon De Bruin. 4e cognition: Historical roots, key concepts, and central issues. 2018.
12 Jan Recker and Alexander Dreiling. Does it matter which process modelling language we teach or use? an experimental study on understanding process modelling languages without formal education. *ACIS 2007 Proceedings*, page 45, 2007.
13 Hajo A. Reijers, Jan Mendling, and Remco M. Dijkman. Human and automatic modularizations of process models to enhance their comprehension. *Inf. Syst.*, 36(5):881–897, 2011.
14 Emily Wall, Meeshu Agnihotri, Laura E. Matzen, Kristin Divis, Michael Haass, Alex Endert, and John T. Stasko. A heuristic approach to value-driven evaluation of visualizations. *IEEE Trans. Vis. Comput. Graph.*, 25(1):491–500, 2019.
15 Wei Wang, Tianwa Chen, Marta Indulska, Shazia Wasim Sadiq, and Barbara Weber. Business process and rule integration approaches – an empirical analysis of model understanding. *Inf. Syst.*, 104:101901, 2022.
16 Shaochen Yu, Tianwa Chen, Lei Han, Gianluca Demartini, and Shazia Sadiq. Dataops-4g: On supporting generalists in data quality discovery. *IEEE Trans. Knowl. Data Eng.*, 35(5):4668–4681, 2023.

## 4.4 Multi-Faceted Process Visual Analytics

*Stef Van den Elzen (TU Eindhoven, NL), Mieke Jans (Hasselt University, BE), Christian Tominski (Universität Rostock, DE), Sebastiaan Johannes van Zelst (Celonis Labs GmbH – München, DE), and Mari-Cruz Villa-Uriol (University of Sheffield, GB)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Stef Van den Elzen, Mieke Jans, Christian Tominski, Sebastiaan Johannes van Zelst, and Mari-Cruz Villa-Uriol

As the application of process mining is becoming more relevant and is being used in a broader context, there is a need for new interactive visual representations that are better focused and use-case specific. For example, the recent object-focused refinement of event data, i.e., often referred to as *object-centric process mining*, defines a multi-dimensional process perspective, requiring corresponding interactive multi-dimensional visual representations. However, the design process of existing process mining visualizations is relatively unstructured and not well-studied. As such, the working group has focused on investigating and developing a conceptual categorization that facilitates a structured design process of multidimensional visual process mining artifacts grounded in *multi-faceted visual analytics design* [4, 3]. The working group has mapped existing process mining visualizations onto well-established facets from the domain of visual analytics. Correspondingly, blind-spots and opportunities (e.g., by means of combining novel facets) have been identified. Furthermore, the categorization enables a more structured approach to exploring and designing new visual representations.

### 4.4.1 Multi-Faceted Visual Analytics

In order to arrive at a well-balanced multi-faceted representation, it makes sense to follow a two-step design procedure [9]. First, a base representation has to be defined for the primary data facet whose depiction will govern the overall display. In alignment with the *effectiveness* principle [5], the most powerful visual variables should be used for the primary facet (usually, this means mapping the primary facet to the visual variable of position). Second, the additional data facet(s) will be incorporated into the base representation using further visual channels (e.g., color, size, orientation, or shape). In the next section, we describe the categorization of data facets applicable to process mining artifacts.

### 4.4.2 Data Facets

A wide variety of facets is identifiable in process mining, where the relevance of a facet as a driving force for a visual design largely depends on the perspective from which users need to look at their data. In general, there are key facets to be considered (in some cases to be refined hierarchically) according to the visual analytics literature. We mostly discussed the following facets inspired by [3] and [9] (see Figure 6):

Note that some *entities*, e.g., activities and objects, can also act as a *given attribute*. To distinguish between these roles in existing visual representations of process mining data and models, we studied the use of *marks* and *visual channels*, as the information-bearing building blocks in visualization idioms. Data that are explicitly represented as marks are interpreted as entities. If data are encoded using visual channels, we interpret them as attributes.

### 4.4.3 A Faceted View on Existing PM Visualizations

To develop a common understanding of existing process mining visualizations, we worked on a categorization that assigns to commonly used visual designs the data facets that they represent. In addition to the visualization-facet mapping, we also cover the marks and visual channels for selected examples:

- **Dotted chart (scatter plot)**
  - *Primary facet*: Time
  - *Additional facets*: Topology (entities – events) and Attributes
  - *Mark*: Event ↦ dot
  - *Visual channel*: Activity ↦ color
- **Performance spectrum (Marey's diagram)**
  - *Primary facet*: Time (temporal order)
  - *Additional facets*: Time (temporal distance) and Attributes
  - *Mark*: Derived Attributes (temporal distance of two events) ↦ line
  - *Visual channel*: end-point positioning ↦ color
- **Maps**
  - *Primary facet*: Space
  - *Additional facets*: Attribute
- **Frequency diagrams (basic charts)**
  - *Primary facet*: Attribute
  - *Additional facets*: *anything*
- **Variant diagrams**
  - *Primary facet*: Time (temporal order)
  - *Additional facets*: Topology (entities – activities) and Attributes
- **Process matrix**
  - *Primary facet*: Topology (relations)
  - *Additional facets*: Attribute

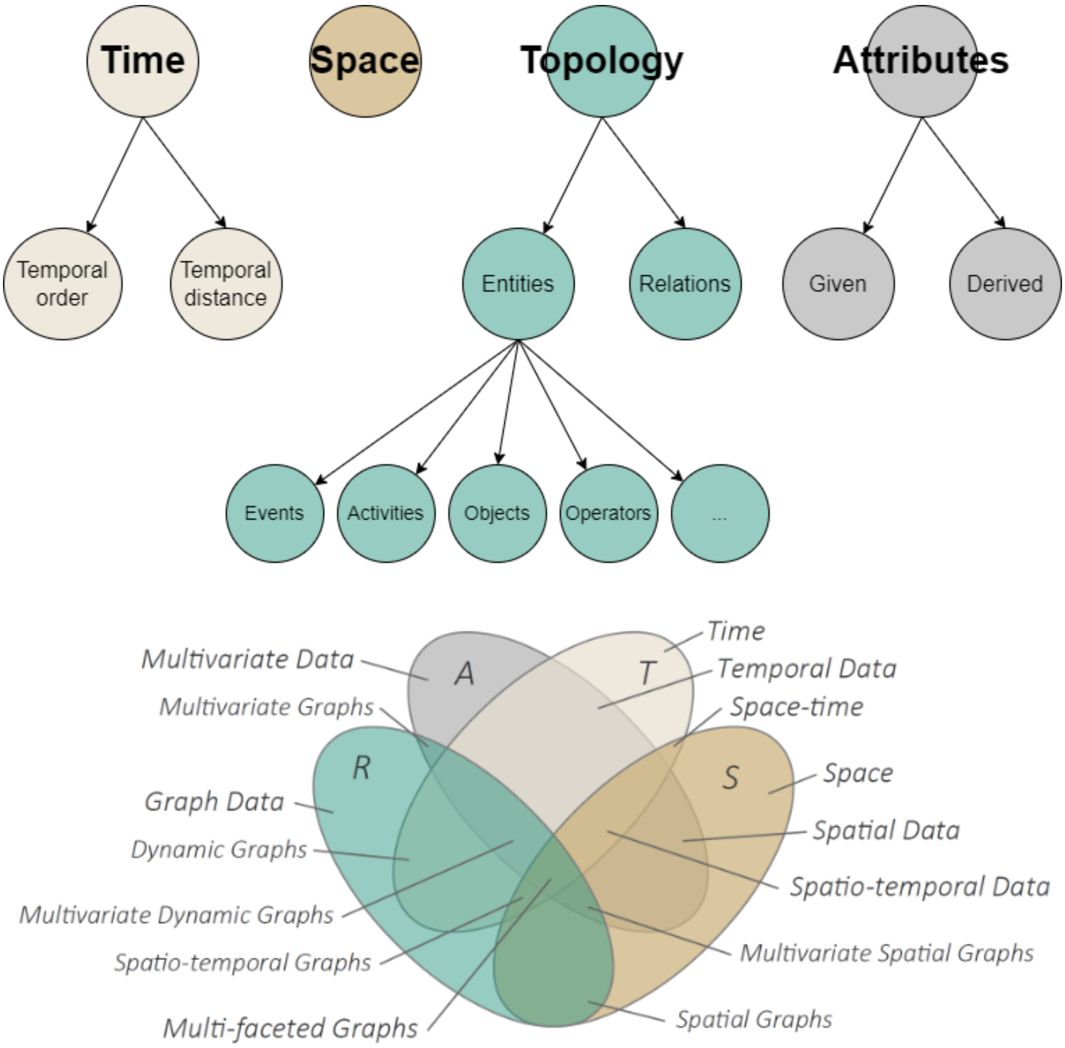


**Figure 6** Overview of facets and their combination relevant to the visual representation of process mining artifacts. Adapted from [9] under CC-BY licence.

- **Directly-Follow Graphs**
  - *Primary facet*: Time (temporal order)
  - *Additional facets*: Topology (relations – activities) and Attributes
- **BPMN**
  - *Primary facet*: Time (temporal order)
  - *Additional facets*: Topology (relations – activities and operators) and Attributes
- **Petri Nets**
  - *Primary facet*: Time (temporal order)
  - *Additional facets*: Topology (relations – activities and operators) and Attributes
- **Process Trees**
  - *Primary facet*: Time (temporal order)
  - *Additional facets*: Topology (relations – activities and operators) and Attributes
- **Object-Centric Petri Nets**
  - *Primary facet*: Time (temporal order)

- *Additional facets*: Topology (relations – activities and operators) and Attributes (objects)

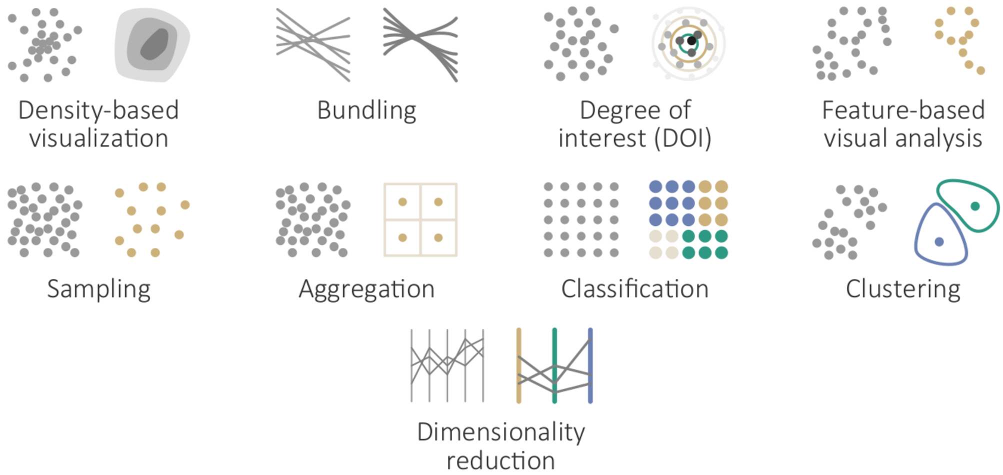


**Figure 7** Overview of automatic computational methods to support interactive visual data analysis by reducing the complexity of the data and their visual representations. Taken from [9] under CC-BY licence.

### 4.4.4 Discussion

From the existing process mining visual representations classified according to the earlier defined categorization, we observed the following:

- Time tends to be the primary facet, mostly focusing on temporal order.
- Space is rarely included in visual representations as a data facet.
- Attribute as a primary facet is only used in basic charts; more complex advanced multivariate visualizations are not used (*e.g.,* scatterplot-matrix, parallel-coordinate plot, glyph-based approaches).
- BPMN, Petri Nets, and Process trees all use the same facets from a data visualization point of view, although they use different semantics.
- In contrast to the other process modeling languages, DFGs do not have operators that can be used as a mark.
- Current object-centric process mining maps use the most important element, *object*, only as a visual channel: if color is removed, it can be considered a traditional Petri Net.

Based on the identified facets and observations, we plan to condense our work into a journal publication. The publication is intended to present the following contributions.

1. Propose a multi-faceted framework for the categorization of PM visualizations
2. Categorization of existing process mining visualizations
3. Identification of blind spots and challenges in PM visualizations
4. Outline opportunities how to enhance PM with interactivity and analytical abstractions

As defined in Munzner [5], visual idioms are described with *Data, Marks, Visual Channels,* and *Tasks*. With the Facets from the categorization, we cover the *Data* element. Furthermore, we identify for each of the existing visual process mining representations the *Marks* and *Visual Channels*. We did not cover the *Tasks* in our working group as this was the topic of working group *Group 6: Task taxonomies in process mining and visual analytics*. By combining the results of both working groups, we can ultimately construct novel visual representations for specific process mining tasks.

The working group also discussed the importance of analytical abstractions and interactivity for advancing existing process mining practices. The visual analytics community provides helpful categorizations of analytical abstractions (see Figure 7), which basically can be applied to each data facet. This would be something to be considered in future work. Similarly, interaction techniques such as view coordination [1], focus+context [2], interactive lenses [8], interactive comparison [7], or dynamic view adaptations [6] would be helpful to integrate into existing process mining tools.

**Figure 8** Members of the working group: Mieke, Bas, Stef, Mari-Cruz, Christian.

## References


1 Michelle Q. Wang Baldonado, Allison Woodruff, and Allan Kuchinsky. Guidelines for using multiple views in information visualization. In Vito Di Gesù, Stefano Levialdi, and Laura Tarantino, editors, *Proceedings of the working conference on Advanced visual interfaces, AVI 2000, Palermo, Italy, May 23-26, 2000*, pages 110–119. ACM Press, 2000.

2 Andy Cockburn, Amy K. Karlson, and Benjamin B. Bederson. A review of overview+detail, zooming, and focus+context interfaces. *ACM Comput. Surv.*, 41(1):2:1–2:31, 2008.

3 Steffen Hadlak, Heidrun Schumann, and Hans-Jörg Schulz. A Survey of Multi-faceted Graph Visualization. In *EuroVis State of the Art Reports*, pages 1–20. Eurographics Association, 2015.

4 Johannes Kehrer and Helwig Hauser. Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Trans. Vis. Comput. Graph.*, 19(3):495–513, 2013.

5 Tamara Munzner. *Visualization Analysis and Design*. A K Peters/CRC Press, 2014.

6 Christian Tominski, Gennady Andrienko, Natalia Andrienko, Susanne Bleisch, Sara Irina Fabrikant, Eva Mayr, Silvia Miksch, Margit Pohl, and André Skupin. Toward Flexible Visual Analytics Augmented through Smooth Display Transitions. *Visual Informatics*, 5(3):28–38, 2021.

7 Christian Tominski, Camilla Forsell, and Jimmy Johansson. Interaction Support for Visual Comparison Inspired by Natural Behavior. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2719–2728, 2012.
8 Christian Tominski, Stefan Gladisch, Ulrike Kister, Raimund Dachselt, and Heidrun Schumann. Interactive Lenses for Visualization: An Extended Survey. *Computer Graphics Forum*, 36(6):173–200, 2017.
9 Christian Tominski and Heidrun Schumann. *Interactive Visual Data Analysis.* AK Peters Visualization Series. CRC Press, 2020.


## 4.5 Visualization for Conformance Checking

*Tatiana von Landesberger (Universität Köln, DE), Jan Mendling (HU Berlin, DE), Giovanni Meroni (Technical University of Denmark – Lyngby, DK), Luise Pufahl (TU München – Heilbronn, DE), and Jana-Rebecca Rehse (Universität Mannheim, DE)*

**License** Creative Commons BY 4.0 International license
© Tatiana von Landesberger, Jan Mendling, Giovanni Meroni, Luise Pufahl, and Jana-Rebecca Rehse

### 4.5.1 Introduction

Process mining defines a set of techniques that take event sequence data in the form of event logs as an input in order to produce evidence-based insights into the execution of business processes. Different groups of techniques are distinguished, including automatic process discovery, conformance checking, enhancement, performance analysis, or deviance analysis [10, 3]. Conformance checking is of particular importance in this context. It allows analysts to investigate to which extent observed event sequences deviate from normative process models and which behavior constraints are violated. This can for example be used for compliance checking to check how far the execution of business processes complies with certain regulations, such as the General Data Protection Regulation (GDPR) or clinical guidelines.

The current practice of conformance checking builds on functionality implemented in professional process mining tools. These tools, in essence, support basic conformance checking tasks: They allow users to (i) quantify the conformance of a process, (ii) compare the conformance values of, e.g., different locations, (iii) localize the specific deviations in a process [6], and (iv) explaining those deviations by means of external factors [8]. Hence, we observe that support for conformance checking has been implemented in various ways. So far, research has not investigated which of the various approaches are most effective for analysts in order to fulfill their tasks. Overall, much of conformance checking research has been driven by technical challenges of efficiently calculating alignments between event sequences and process models [2], defining precision and fitness measures with sound formal properties [7], and providing feedback to business users who are not familiar with formal concepts [4].

In general, research on process mining can be framed at different levels [11]. Empirical research has been conducted at the level of organizational anchoring and impact [1] but mostly abstracted from the analyst as a user. What is by and large missing is a research strategy that considers the tasks and user characteristics as a focus of conformance checking research [5]. More specifically, several research questions miss specific answers from empirical and engineering research on process conformance checking, including the following.

- What kind of characteristics do the users have?
- Which are the conformance checking tasks?
- What data is required for each of the tasks?

- Which techniques can be used for specific tasks? What is missing?
- What are the functional/non-functional requirements of the systems?
- Which are effective conformance checking algorithms for the tasks?
- Which are effective visualizations for the conformance checking tasks?
- How to design an effective interaction between the conformance checking algorithms and visualizations?
- How to support annotations and editing to the original data?
- How to support decision-making and actions (i.e., implement interventions) in the real world?
- How to cope with updates of the real world in the conformance checking visualization system?

This report presents the discussion of the working group on conformance checking of the Dagstuhl Seminar 23271 on "Human in the (Process) Mines", held at Schloss Dagstuhl on 2–7 July 2023. We conclude by summarizing a set of action points that our working group identified.

### 4.5.2 Discussions of the Working Group

As a first step, the working group identified the limitations of current visualization techniques adopted by conformance checking tools. It emerged that deviations are often not represented in an effective way that stands out from the normal execution flow. This is typically caused by a gap between visualization producers, users who create the graphical notations, and visualization consumers, users who have to interpret the graphical notation.

The working group then discussed the typical workflow that visual analytics researchers adopt to address these issues [9]. This typically consists in involving domain experts to identify what kind of tasks (actions) that they need to perform to obtain insights. Then conformance checking techniques able to address these tasks are adopted or developed, which then produce output data. Models and visuals will then be built to effectively represent output data. Finally, the visual representation will be checked by domain experts to see if they are effective for accomplishing their tasks, what is missing, and what can be improved.

Following this workflow, the working group analyzed different application domains in which they had direct experience (e.g., by previous research projects) in order to identify the tasks (actions) that the user need to perform to obtain insights on the compliance of its processes. From this analysis, it emerged that, for auditing purposes, it is more relevant to produce a list of cases that caused a deviation and see their impact. If one has a large number of deviations, these need to be grouped into classes. The hypothesis in conformance checking is that if the trace is the same, then they have the same meaning. However, information on resources, data, etc, is currently missing.

This lead to the definition of which information should be the input for a system combining visual analytics with conformance checking, and which interactions with the domain expert this system should support (see Fig. 9).

To apply conformance checking one needs to create a model from regulations and norms, and an event log from information on a database or other data sources. The model can use different formalisms (DFG, BPMN, Declare, etc.), and the logs can be at different levels (task execution, task start/complete, changes in data, etc). Then, the outcome of conformance checking would be deviations. Deviations are grouped in harmful/not harmful, variants, or having additional or missing activities. They are then typically qualified based on risk by domain experts, and quantified by computing measures from the types. Then, deviations are explained. Qualifying, Quantifying and Grouping can be performed multiple times.

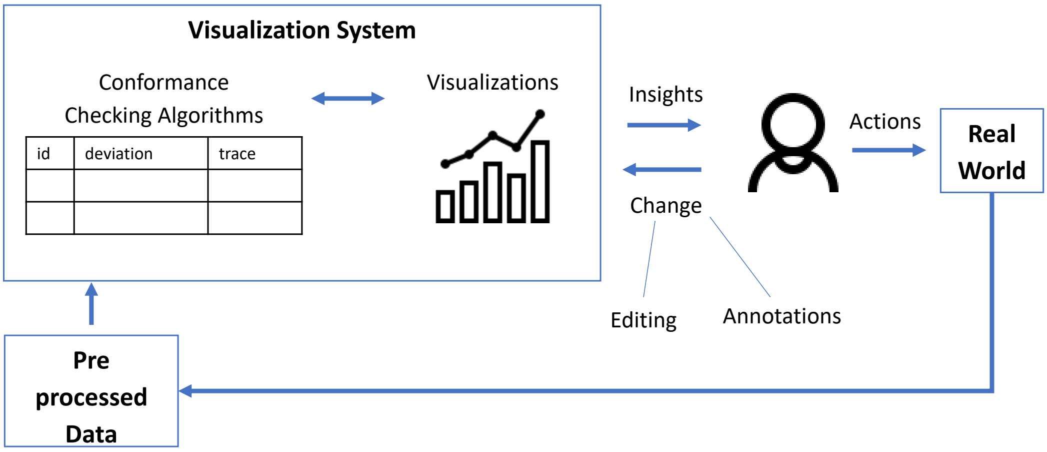


**Figure 9** Usage of a conformance checking visualization system.

Thank to conformance checking, one can perceive how the actual process is performed, then make sense of process-related information, which leads to data-driven decision making, that can eventually lead to implementing interventions on the process.

VA can 1) apply an algorithm and then visualize the results: you choose the algorithm and visualize/interact only with the output. 2) Interact in a cyclic behavior: after observing the result, parameters are changed and the algorithm is rerun. 3) computational steering: while the algorithm is running one changes the parameters and adjusts on-the-fly what is happening.

One can also edit VA results with annotations and corrections. Corrections directly change the data sources (e.g., fixing a wrong timestamp). Annotations highlight relevant behaviors in the data, without changing the data sources (e.g., noting that a regulation has changed, explaining why a trace is non-conformant).

Next, we discussed what input data is available and needed for proper and useful conformance checking. We identified, on the one hand, the event log or a stream of events that represents the process execution. On the other hand, we discussed that the business analyst usually does not provide a standard process model, but rather a set of rules to which the process needs to comply. This can be control flow (activity A must be followed by B), time (activity A must complete within 30 minutes), resource (activity A must be executed by the financial department), and data-related rules (activity A must produce an invoice with a fee not higher than 5000 EUR). These rules need to be formalized, such that conformance checking techniques can automatically identify deviations in the traces of an event log, the individual process executions. Sometimes, business analyst has no rules at hand, but they have hypotheses in mind about how their process should be executed and what should not happen. For such a case, the business analyst needs support in formulating hypotheses and confirming them.

### 4.5.3 Next steps

Our discussions this week have led us to believe that visual analytics can provide a completely new angle on process conformance, with a much stronger emphasis on the needs of the user and the tasks that conformance checking is supposed to fulfill. The working group has agreed on two immediate next steps. First, we would like to gather our newly established insights into the potential of visual analytics for the advancement of conformance checking into a workshop paper, outlining the ideas and concepts already presented in this report. Second, we would like to further develop, substantiate, and structure these ideas, with the goal of submitting a joint research proposal.

Concretely, the next steps that we want to take for the substantiation of our ideas are the following:

- Identify additional application scenarios, such as healthcare, purchasing, infection control, ambient assisted living, or logistics, and characterize them with regard to their similarities and differences to obtain a clear set of objectives
- Conduct an design study [9] with process analysts or process managers to verify and extend the preliminary list of tasks that conformance checking should fulfill
- Collect and analyze process-related regulations to develop a better understanding of the different types of rules that business processes are subject to
- Analyze existing visualizations of timed event sequence data, such as [12], from the VA domain to find out which functionalities they already provide

## References


1 Peyman Badakhshan, Bastian Wurm, Thomas Grisold, Jerome Geyer-Klingeberg, Jan Mendling, and Jan Vom Brocke. Creating business value with process mining. *The Journal of Strategic Information Systems*, 31(4):101745, 2022.

2 Josep Carmona, Boudewijn van Dongen, Andreas Solti, and Matthias Weidlich. Conformance checking. *Switzerland: Springer.[Google Scholar]*, 56, 2018.

3 Marlon Dumas, Marcello La Rosa, Jan Mendling, Hajo A Reijers, et al. *Fundamentals of business process management*, volume 2. Springer, 2018.

4 Luciano García-Bañuelos, Nick RTP Van Beest, Marlon Dumas, Marcello La Rosa, and Willem Mertens. Complete and interpretable conformance checking of business processes. *IEEE Transactions on Software Engineering*, 44(3):262–290, 2017.

5 Jan Mendling, Djordje Djurica, and Monika Malinova. Cognitive effectiveness of representations for process mining. In *International Conference on Business Process Management*, pages 17–22. Springer, 2021.

6 Giovanni Meroni. *Artifact-Driven Business Process Monitoring – A Novel Approach to Transparently Monitor Business Processes, Supported by Methods, Tools, and Real-World Applications*, volume 368 of *Lecture Notes in Business Information Processing*. Springer, 2019.

7 Artem Polyvyanyy, Andreas Solti, Matthias Weidlich, Claudio Di Ciccio, and Jan Mendling. Monotone precision and recall measures for comparing executions and specifications of dynamic systems. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 29(3):1–41, 2020.

8 Jana-Rebecca Rehse, Luise Pufahl, Michael Grohs, and Lisa-Marie Klein. Process mining meets visual analytics: the case of conformance checking. In *Hawaii International Conference on System Sciences*, 2023.

9 Michael Sedlmair, Miriah Meyer, and Tamara Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE transactions on visualization and computer graphics*, 18(12):2431–2440, 2012.

10 Wil van der Aalst. *Data science in action*. Springer, 2016.

11 Jan vom Brocke, Mieke Jans, Jan Mendling, and Hajo A Reijers. A five-level framework for research on process mining. *Business & Information Systems Engineering*, 63(5):483–490, 2021.

12 Marcel Wunderlich, Kathrin Ballweg, Georg Fuchs, and Tatiana von Landesberger. Visualization of delay uncertainty and its impact on train trip planning: A design study. In *Computer Graphics Forum*, volume 36, pages 317–328. Wiley Online Library, 2017.

## 4.6 Milana – A Task Taxonomy for Supporting Human-Centered Process Mining with Visual Analytics

*Lisa Zimmermann (Universität St. Gallen, CH), Philipp Koytek (Celonis Labs GmbH – München, DE), Shazia Sadiq (University of Queensland – Brisbane, AU), Pnina Soffer (University of Haifa, IL), Katerina Vrotsou (Linköping University, SE), and Barbara Weber (Universität St. Gallen, CH)*

**License** Creative Commons BY 4.0 International license
© Lisa Zimmermann, Philipp Koytek, Shazia Sadiq, Pnina Soffer, Katerina Vrotsou, and Barbara Weber

Process Mining is an activity that can significantly benefit from human-in-the-loop approaches. The current body of knowledge is lacking a full exploitation of opportunities that sit at the intersection of human-centered process mining activities, the various cognitive aspects of these activities, and visual ways of supporting them. To realize these opportunities we undertake a taxonomical approach to identify the key concepts and relationships across the three bodies of knowledge. The objective of the proposed taxonomy, namely Milana, is to establish a common vocabulary and shared understanding, serve as a point of reference, and inform measurements for evaluating process visualizations.

### 4.6.1 Introduction

Over the last two decades, process mining has developed into a collection of techniques and tools that have the potential to provide immense value to business users. Although the techniques have gained a significant level of sophistication in terms of their algorithmic approaches and outcomes, supporting the human within the process of process mining [7] has so far received little attention.

Therefore, the goal of this work is to **develop a task taxonomy for supporting human-centered process mining with visual analytics**.

We argue that the development of such a taxonomy, which we call *Milana*, needs to be guided by multiple perspectives that include process mining, visual analytics and cognitive underpinnings. The current body of knowledge is lacking a full exploitation of opportunities that sit at the intersection of human-centered process mining activities, the various cognitive aspects of these activities, and visual ways of supporting them.

We envision that *Milana* can provide a reference to apply a three-way mapping between analytical tasks in process mining, cognitive tasks, and visualization tasks. Such a mapping will establish a common vocabulary and contribute to a shared understanding between the different communities. Moreover, it can be leveraged to justify design choices of process visualizations, identify gaps to inform future work, and be used to inform the measurements for evaluating process visualizations.

### 4.6.2 Description of the theme

A taxonomy (or taxonomic classification) is a scheme of classification, especially a hierarchical classification, in which things are organized into groups or types [8]. Typically, the development of a taxonomy involves a number of steps including (i) identifying the concepts

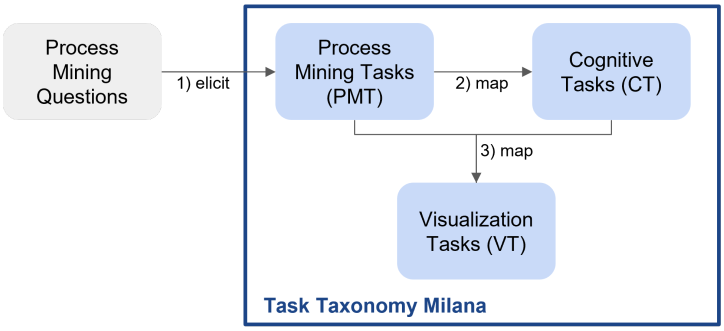


**Figure 10** Taxonomy Development Approach.

and relationships within the domain(s) of interest, (ii) establishing a nomenclature for the representation of the concepts, (iii) constructing the taxonomy using established notation (usually hierarchical), (iv) evaluating or validating it against its intended goal, and (v) setting up a mechanism for its continued maintenance.

During the seminar, we derived an approach to develop *Milana*. As shown in Fig 10, it consists of three steps to identify concepts and relationships of the taxonomy:

1. Elicit process mining tasks (PMT) (Sec. 4.6.2.1)
2. Mapping process mining tasks to cognitive tasks (CT) (Sec. 4.6.2.2)
3. Mapping combinations of process mining tasks and cognitive tasks to visualization tasks (VT) (Sec. 4.6.2.3)

Taking these steps as a starting point, we can build on the large body of knowledge on visualization and visual analytics that gives us a proven link on relevant and most suited visual representations which is the main goal of developing *Milana*.

Considering that the concepts will be derived from three different bodies of knowledge, namely process mining, visual analytics and cognitive theories, establishing the nomenclature is a critical and non-trivial step and will continue to be refined as the taxonomy is further developed. The work on the taxonomy construction, validation and plan for maintenance is yet to be undertaken.

#### 4.6.2.1 Elicit Process Mining Tasks (PMT)

The first step of our work is to elicit a comprehensive overview of process mining tasks. This requires choosing a valid source for identifying such tasks and to align on a suitable abstraction level for these tasks.

To identify tasks, we evaluated different approaches, of which the elicitation from questions seems to be the most promising one. In visual analytics literature, questions and tasks are often used as synonyms [1] and also in the process mining literature, for example Zerbato et al. [10], identified the analysis question as an important factor that determines the analysis approach (task).

As a source for process mining questions, we plan to start from the business process intelligence challenges (BPIC) published by the *IEEE Task Force on Process Mining* of which ten editions, dating back to 2011, exist [3]. These challenges contain real-world event logs and business questions posted by different organizations. Moreover, we can leverage out-of-the-box analyses for questions commonly asked by customers in existing process mining tools (e.g., Celonis Business Miner). These questions serve as an additional source and complement the BPIC questions.

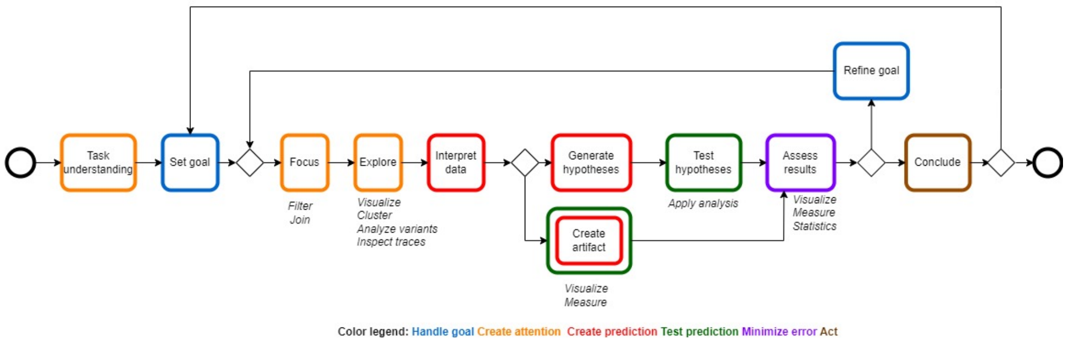


**Figure 11** The PEM4PPM model (Prediction Error Minimization for the Process of Process Mining [7]).

Exemplary questions are "Are there cases that spend a very short time at Service Line 1 before being escalated?"(Source: BPIC 2013), or "How many unwanted activities are in my process?" (Source: Celonis Business Miner).

Depending on the complexity of the question it might manifest in several process mining tasks (low level of abstraction). For the first question, for example, one would need to (i) identify cases processed by service line 1, (ii) calculate the throughput time for the respective sub-process, (iii) identify what "short" refers to and derive a threshold, and (iv) count the remaining cases.

As part of the taxonomy development we aim to harmonize the tasks over the available questions until we reach saturation. For this purpose, we plan to identify the (main) operations performed to answer the question, the process mining entity that is the subject of interest (e.g. case, activity, ...), its respective cardinality, and the attribute of the event log that is required to execute that operation.

#### 4.6.2.2 Mapping to Cognitive Tasks (CT)

At this step we consider the cognitive tasks that are involved in the process of process mining [7], depicted in Fig. 11. This model relates to the individual process of process mining from a cognitive perspective where different process mining tasks may be performed iteratively. For example, conducting one or many process mining tasks may entail focusing on a part of the log, creating a process model (an artifact), exploring the model to identify "interesting" patterns, generating hypotheses which explain these patterns, applying additional analysis algorithms for testing these hypotheses, and assessing whether the result is as expected and can be used for completing the task.

Our premise is that the cognitive tasks that underpin various process mining tasks provide an important perspective on the information needs of the process analyst and are therefore relevant for the identification of suited visualizations supporting these tasks. Hence, we propose to include this perspective in the taxonomy which requires a mapping of the process mining tasks identified in the previous step (cf. Sec 4.6.2.1) to the cognitive tasks.

#### 4.6.2.3 Mapping to Visualization Tasks (VT)

The final perspective we consider for *Milana*, is the mapping to the visualization task that links to the combination process mining task and cognitive task. As a starting point for this third mapping step, we surveyed the currently available taxonomies in the field of visualization and visual analytics and identified that they vary with respect to their focus and levels of abstraction. Examples include users' interactions intents [9] and visualization usage [2, 4].

Brehmer & Munzner [2] proposed a multi-level typology of visualization tasks that attempts to specify the *WHY* and *HOW* of a visualization task that is performed. The aim of their taxonomy is to enable the expression of complex sequences of activities in terms of simpler tasks. Moreover, its focus is on a high abstraction level that is not tied to specific application domains. This allows the freedom to make comparisons between domains but also provides the opportunity to customize it for the purpose of a specific domain; which is the ambition of this working group.

To derive our approach, we draw inspiration from Rind et al. [5] who introduced a three-dimensional conceptual space for user tasks in visualization, called *TaskCube*. The three considered dimensions are *perspective*, *composition*, and *abstraction*. Perspective is concerned with the objectives of the user, corresponding to the "why" of the taxonomy by Brehmer & Munzner [2], and actions which correspond to the "how". Composition relates to the break-down of the task from low-level, including small and concise steps, to high-level steps concerned with broader tasks. Finally, abstraction classifies tasks according to their level of abstraction from concrete to abstract.

In their work, Rind et al. [5] have classified a number of previously related taxonomies in terms of the dimensions of the *TaskCube*. This classification provides further motivation and guidance for our work towards a taxonomy. To this end, we intend to extend the generic visualization tasks defined by Brehmer & Munzner [2] to a more concrete abstraction level by applying it to the formulated process mining tasks (PMTs) (Sec. 4.6.2.1). In terms of the *TaskCube* classification this would place our proposed taxonomy on a concrete abstraction level with respect to data type (i.e. process data) and potentially also domain (e.g., "Business Processes").

We envision to follow an approach comparable to Rind et al. [6] which extends the proposed taxonomy of interaction intents by Yi et al. [9] by concretizing the intents with sub-intents specific to the healthcare domain.

### 4.6.3 Outline of the Next Steps

In the future we plan to continue our work as follows.

1. Concerning the identification of process mining tasks, we will continue to apply the approach to more questions that are collected from the BPICs until we reach saturation (cf. Sec. 4.6.2.1).
2. For mapping the identified process mining tasks to the cognitive tasks that would be required, a critical evaluation of the relevant cognitive tasks that could guide the identification of visual representations is outstanding (cf. Sec. 4.6.2.2).
3. After the process mining tasks are mapped to cognitive tasks, visualization tasks will be linked to these. Starting from the taxonomy by Brehmer & Munzner [2] appropriate options will be identified (cf. Sec. 4.6.2.3).
4. Finally, *Milana* will be validated (e.g., case studies, user studies, and experiments).

*Milana* will pave the way for both evaluating existing process analysis visualizations in terms of their appropriateness for the tasks as well as provide a blueprint for the design of novel visual representations.

**Acknowledgments.** B. Weber and L. Zimmermann are supported by the ProMiSE project funded by the SNSF under Grant No.: 200021_197032.

## References


1 Natalia V. Andrienko and Gennady L. Andrienko. *Exploratory analysis of spatial and temporal data – a systematic approach.* Springer, 2006.
2 Matthew Brehmer and Tamara Munzner. A multi-level typology of abstract visualization tasks. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2376–2385, 2013.
3 BPI Challenge. BPI Challenge. `https://www.tf-pm.org/competitions-awards/bpi-challenge`, 2023. [Online; accessed 25-July-2023].
4 Jeffrey Heer and Ben Shneiderman. Interactive dynamics for visual analysis: A taxonomy of tools that support the fluent and flexible use of visualizations. *Queue*, 10(2):30–55, 2012.
5 Alexander Rind, Wolfgang Aigner, Markus Wagner, Silvia Miksch, and Tim Lammarsch. Task cube: A three-dimensional conceptual space of user tasks in visualization design and evaluation. *Inf. Vis.*, 15(4):288–300, 2016.
6 Alexander Rind, Taowei David Wang, Wolfgang Aigner, Silvia Miksch, Krist Wongsuphasawat, Catherine Plaisant, and Ben Shneiderman. Interactive information visualization to explore and query electronic health records. *Found. Trends Hum. Comput. Interact.*, 5(3):207–298, 2013.
7 Elizaveta Sorokina, Pnina Soffer, Irit Hadar, Uri Leron, Francesca Zerbato, and Barbara Weber. Pem4ppm: A cognitive perspective on the process of process mining. In *Business Process Management – 21th International Conference, BPM 2023.* Accepted for publication.
8 Wikipedia, the free encyclopedia. Taxonomy. `https://en.wikipedia.org/wiki/Taxonomy`, 2023. [Online; accessed 25-July-2023].
9 Ji Soo Yi, Youn ah Kang, John T. Stasko, and Julie A. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1224–1231, 2007.
10 Francesca Zerbato, Pnina Soffer, and Barbara Weber. Process mining practices: Evidence from interviews. In Claudio Di Ciccio, Remco M. Dijkman, Adela del-Río-Ortega, and Stefanie Rinderle-Ma, editors, *Business Process Management – 20th International Conference, BPM 2022, Münster, Germany, September 11-16, 2022, Proceedings*, volume 13420 of *Lecture Notes in Computer Science*, pages 268–285. Springer, 2022.

## Participants

- Wolfgang Aigner
FH – St. Pölten , AT
- Anti Alman
University of Tartu, EE
- Alessio Arleo
TU Wien, AT
- Iris Beerepoot
Utrecht University, NL
- Axel Buehler
Olympus – Hamburg, DE
- Andrea Burattin
Technical University of Denmark – Lyngby, DK
- Claudio Di Ciccio
Sapienza University of Rome, IT
- Chiara Di Francescomarino
University of Trento, IT
- Irit Hadar
University of Haifa, IL
- Mieke Jans
Hasselt University, BE
- Philipp Koytek
Celonis Labs GmbH – München, DE
- Monika Malinova Mandelburger
Wirtschaftsuniversität Wien, AT
- Andrea Marrella
Sapienza University of Rome, IT
- Jan Mendling
HU Berlin, DE
- Giovanni Meroni
Technical University of Denmark – Lyngby, DK
- Silvia Miksch
TU Wien, AT
- Luise Pufahl
TU München – Heilbronn, DE
- Jana-Rebecca Rehse
Universität Mannheim, DE
- Manuel Resinas
University of Sevilla, ES
- Shazia Sadiq
University of Queensland – Brisbane, AU
- Daniel Schuster
Fraunhofer FIT – Sankt Augustin, DE
- Pnina Soffer
University of Haifa, IL
- Christian Tominski
Universität Rostock, DE
- Cagatay Turkay
University of Warwick – Coventry, GB
- Stef Van den Elzen
TU Eindhoven, NL
- Sebastiaan Johannes van Zelst
Celonis Labs GmbH – München, DE
- Mari-Cruz Villa-Uriol
University of Sheffield, GB
- Tatiana von Landesberger
Universität Köln, DE
- Katerina Vrotsou
Linköping University, SE
- Barbara Weber
Universität St. Gallen, CH
- Francesca Zerbato
Universität St. Gallen, CH
- Lisa Zimmermann
Universität St. Gallen, CH

Report from Dagstuhl Seminar 23272

# Epistemic and Topological Reasoning in Distributed Systems

**Armando Castañeda*[1], Hans van Ditmarsch*[2], Roman Kuznets†[3], Yoram Moses*[4], and Ulrich Schmid*[5]**

1 **National Autonomous University of Mexico, MX.** armando.castaneda@im.unam.mx
2 **CNRS – Toulouse, FR & Université de Toulouse, FR.** hans.van-ditmarsch@irit.fr
3 **TU Wien, AT.** roman@logic.at
4 **Technion – Haifa, IL.** yoram.moses@gmail.com
5 **TU Wien, AT.** s@ecs.tuwien.ac.at

**Abstract**

This report documents the program and the outcomes of Dagstuhl Seminar 23272 "Epistemic and Topological Reasoning in Distributed Systems." The seminar brought together experts in combinatorial topology and epistemic logic interested in distributed systems, with the aim of exploring the directions that the recent interaction between those approaches can take, identifying challenges and opportunities.

**Seminar** July 2–7, 2023 – https://www.dagstuhl.de/23272

**2012 ACM Subject Classification** Computing methodologies → Distributed algorithms; Mathematics of computing → Graph theory; Mathematics of computing → Topology; Software and its engineering → Distributed systems organizing principles; Software and its engineering → Formal methods; Theory of computation → Distributed computing models; Theory of computation → Data structures design and analysis; Theory of computation → Logic

**Keywords and phrases** combinatorial topology, distributed systems, epistemic logic, multi-agent systems, interpreted systems, dynamic epistemic logic, simplicial semantics, knowledge-based approach, distributed computing

**Digital Object Identifier** 10.4230/DagRep.13.7.34

## 1 Executive Summary

*Armando Castañeda (National Autonomous University of Mexico, MX)*
*Hans van Ditmarsch (CNRS – Toulouse, FR & Université de Toulouse, FR)*
*Yoram Moses (Technion – Haifa, IL)*
*Ulrich Schmid (TU Wien, AT)*

**License** Creative Commons BY 4.0 International license
© Armando Castañeda, Hans van Ditmarsch, Yoram Moses, and Ulrich Schmid

Distributed services cover a wide range of our everyday activities. Examples of such services include cloud storage, cryptocurrencies, and collaborative editing, as well as concurrent software that governs modern multicore computers. Reasoning about distributed systems that provide these services is notoriously difficult, however, due to the many sources of uncertainty that can occur: varying execution speeds, unpredictable transmission delays, and

* Editor / Organizer
† Editorial Assistant / Collector

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

partial failures. The design and analysis of protocols and algorithms for distributed systems is, hence, a difficult and error-prone task.

Two approaches have proved to be successful in raising the level of abstraction in the modeling, design, and analysis of distributed systems: combinatorial topology and the epistemic (or knowledge-based) approach for multi-agent systems. It is also worth noting that several different flavors of topology have been successfully applied in general modal logic and related fields as well.

In the context of distributed systems, the epistemic and the topological approach have evolved fairly independently for over more than three decades. In the last few years, however, researchers have started to combine the two approaches in productive ways. In particular, this is true for both of the two main variants of epistemic reasoning in such systems, the more traditional *interpreted systems* (IS) modeling and *dynamic epistemic logics* (DEL), where approaches and/or methods from combinatorial topology have been used successfully already. This is primarily due to a duality between the Kripke models, which underlie epistemic reasoning, and simplicial complexes, which are central to the analysis of distributed protocols using combinatorial topology. And indeed, meanwhile, there are encouraging relations and options for cross-fertilization with the communities of dynamic epistemic logic, knowledge-based analysis, and topological modal logics. This development makes a new level of abstraction possible, based on mutual incorporation of the extensive results established independently by all these approaches.

## Seminar Goals and Structure

The main purpose of this Dagstuhl Seminar was to further stimulate the process of making this integration happening: Whereas there are researchers with expertise in more than one of the above fields, a majority of them is firmly grounded in only one of those. Consequently, we brought together 28 experts from the respective scientific communities, in an attempt to:

**(A)** Introduce the core topics, methods and accomplishments obtained in some field to the others.

**(B)** Stimulate discussions among members from different fields, to identify possible cross-fertilization opportunities.

**(C)** Further explore particularly interesting/promising/challenging cross-fertilization opportunities and shape collaborations on these.

The cornerstones of the seminar program were:

**(i)** 5 tutorials (90 min) by experts in the following specific areas:

- *Alexandru Baltag*
  Epistemics and Topology
- *Hans van Ditmarsch*
  Dynamic Epistemic Logic
- *Guy Goren*
  Knowledge and Action
- *Jérémy Ledent*
  Simplicial Models: From Global States to Local States, and What Lies In-between
- *Sergio Rajsbaum*
  An Overview of Combinatorial Topology and Its Distributed Computing Perspective

**(ii)** 17 presentations (20 min) devoted to specific research topics, from all areas.

**(iii)** 3 short presentations (5 min) in a rump session, devoted to research topics that were identified in the course of the seminar.

**(iv)** Public sessions devoted to identifying particularly promising cross-fertilization opportunities.

**(v)** 3 focused discussion groups, devoted to the following cross-fertilization opportunities:
- Extending Topology to Handle Data Structures
- Interpreted Systems and Dynamic Epistemic Logic
- Representing Epistemic Attitudes via Simplicial Complexes

Multiple discussion sessions of these groups were run in parallel, with two public reporting sessions. Seminar participants could join/swap the discussion groups freely.

## 2 Table of Contents

**Working Groups**

# 3 Overview of Talks

## 3.1 Coalgebraic Modal Logic and Fibrations

*Henning Basold (Leiden University, NL)*

License Creative Commons BY 4.0 International license
© Henning Basold

In this talk, I provide an overview over coalgebraic modal logic and how it can be abstractly viewed through fibrations. The idea of coalgebras is that they provide an abstraction of state-based systems and enable a category-theoretical study of behaviour, just like (universal) algebra provides tools to reason about different algebraic structures. While in algebra we are interested in structures with operations on them, such as monoids, groups, etc., in coalgebra we consider observations as the fundamental building block. In the talk, we discuss labelled transitions systems as a particular example, in which we can observe for every state the propositional atoms that are true at that state and the outgoing labelled transitions [1]. These can represent epistemic models, linear dynamical systems, Kripke frames, neighbourhood frames, descriptive frames [2], and even proofs of coinductive predicates. To reason about the behaviour induced by coalgebras, we discuss predicates that arise via so-called liftings [3] and rely on the theory of fibrations [4]. We will see the theory illustrated on the labelled box modality for labelled transition systems and we will show how a modal logic with conjunction and box can be given semantics over a coalgebra by interpreting the conjunction via the algebraic structure of meet-semilattices and the modality via a so-called distributive law that is induced by the previously defined lifting. I will present the basic principles of coalgebraic modal logic and how the approach of fibrations can be used to give semantics to modal logic in quite general coalgebras, and to prove soundness and completeness results that extend those of Kupke and Rot [5].

### References

1. Lawrence S. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96(1–3):277–317, March 1999. doi:10.1016/S0168-0072(98)00042-6.
2. Guram Bezhanishvili, Luca Carai, and Patrick J. Morandi. The Vietoris functor and modal operators on rings of continuous functions. *Annals of Pure and Applied Logic*, 173(1):103029, January 2022. doi:10.1016/j.apal.2021.103029.
3. Ichiro Hasuo, Kenta Cho, Toshiki Kataoka, and Bart Jacobs. Coinductive predicates and final sequences in a fibration. In Dexter Kozen and Michael Mislove, editors, *Proceedings of the Twenty-ninth Conference on the Mathematical Foundations of Programming Semantics, MFPS XXIX*, volume 298 of *Electronic Notes in Theoretical Computer Science*, pages 197–214. Elsevier, 2013. doi:10.1016/j.entcs.2013.09.014.
4. Bart Jacobs. *Categorical Logic and Type Theory*. Number 141 in Studies in Logic and the Foundations of Mathematics. Elsevier, 1999.
5. Clemens Kupke and Jurriaan Rot. Expressive logics for coinductive predicates. *Logical Methods in Computer Science*, 17(4):19:1–19:30, 2021. doi:10.46298/lmcs-17(4:19)2021.

## 3.2 Asynchronous Speedup Theorems and FLP-Style Proofs

*Pierre Fraigniaud (CNRS – Paris, FR & Université Paris Cité, FR)*

**License** Creative Commons BY 4.0 International license
© Pierre Fraigniaud
**Joint work of** Hagit Attiya, Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum

The talk compares two generic techniques for deriving lower bounds or impossibility results in distributed computing. First, we formalize the notion of FLP-style proofs, aiming at capturing the essence of the seminal consensus-impossibility proof [1] and using forward induction. Second, we prove a speedup theorem à la [2], but for wait-free colorless algorithms, aiming at capturing the essence of the seminal round-reduction proof establishing a lower bound on the number of rounds for 3-coloring a cycle [3], and going by backward induction. We show that despite their very different natures, these two forms of proof are tightly connected. In particular, we show that for every colorless task $\Pi$, if there is a round-reduction proof establishing the impossibility of solving $\Pi$ using wait-free colorless algorithms, then there is an FLP-style proof establishing the same impossibility. For 1-dimensional colorless tasks (but for an arbitrarily large number $n \geq 2$ of processes), we prove that the two proof techniques have exactly the same power, and more importantly, both are complete: if a 1-dimensional colorless task is not wait-free solvable by $n \geq 2$ processes, then the impossibility can be proved by both proof techniques. Moreover, a round-reduction proof can be automatically derived, and an FLP-style proof can be automatically generated from it. Finally, we illustrate the use of these two techniques by establishing the impossibility of solving any colorless covering task of arbitrary dimension by wait-free algorithms.

### References

1 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985. `doi:10.1145/3149.214121`.
2 Sebastian Brandt. An automatic speedup theorem for distributed problems. In *PODC'19: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 379–388. Association for Computing Machinery, 2019. `doi:10.1145/3293611.3331611`.
3 Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, February 1992. `doi:10.1137/0221015`.

## 3.3 Epistemic Analysis of the FRR Problem

*Krisztina Fruzsa (TU Wien, AT)*

**License** Creative Commons BY 4.0 International license
© Krisztina Fruzsa
**Joint work of** Krisztina Fruzsa, Roman Kuznets, Ulrich Schmid
**Main reference** Krisztina Fruzsa, Roman Kuznets, Ulrich Schmid: "Fire!", in Proc. of the Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2021, Beijing, China, June 25-27, 2021, EPTCS, Vol. 335, pp. 139–153, Open Publishing Association, 2021.
**URL** https://doi.org//10.4204/EPTCS.335.13

We present an epistemic analysis of a fundamental distributed computing problem called the Firing Rebels with Relay (FRR) within the asynchronous byzantine fault-tolerant model of distributed systems. Essentially, the FRR problem requires all correct agents to perform an action (FIRE) in an all-or-nothing fashion, however, not necessarily at the same time. By using our framework for modeling asynchronous byzantine fault-tolerant distributed systems, we establish the necessary level of knowledge that needs to be acquired by any correct agent in order to perform FIRE in every protocol that meets the requirements of the FRR problem specification. The corresponding level of knowledge involves the so-called common eventual hope modality, which plays a crucial role in reaching (eventual) agreement among agents.

## 3.4 Semitopology and (Paraconsistent) Logic

*Murdoch James Gabbay (Heriot-Watt University – Edinburgh, GB)*

**License** Creative Commons BY 4.0 International license
© Murdoch James Gabbay
**Joint work of** Murdoch James Gabbay, Giuliano Losa
**Main reference** Murdoch J. Gabbay, Giuliano Losa: "Semitopology: a new topological model of heterogeneous consensus", CoRR abs/2303.09287, 2023. [1]
**URL** https://doi.org//10.48550/ARXIV.2303.09287

A semitopology is like a topology, but without the constraint that intersections of open sets be open. This models an idea in distributed consensus that participants can progress locally if they reach agreement amongst a local quorum (= open neighbourhood) of trusted peers.

Semitopologies are related to Horn clause theories: just fix some predicate atoms, write any Horn clause theory over those atoms that you like, and consider the set of assignments of truth values to the atoms that satisfy the clauses. These assignments generate the closed sets of a semitopology as their corresponding answer sets (the set of atoms to which the assignment assigns "true"). It is very easy to see that this is a semitopology.

But now we want to assert and reason about semitopological properties, by applying techniques from logic (such as derivability) to the semitopology as represented as a Horn clause theory. In essence, we have converted safety properties of distributed consensus into derivability (or non-derivability) questions about a simple logical theory.

For this, we require a paraconsistent modal logic (over three values: T, B, and F), with a modality $K$ asserting validity in all valuations.

Time permitting, I will sketch the maths above and show the logic in action on some simple examples.

Some parts of the material described above are in a draft journal paper [1] online.

### References

1 Murdoch J. Gabbay and Giuliano Losa. Semitopology: a new topological model of heterogeneous consensus. Eprint 2303.09287, arXiv, 2023. doi:10.48550/arXiv.2303.09287.

## 3.5 Optimal Eventual Byzantine Agreement Protocols with Omission Failures

*Joseph Y. Halpern (Cornell University – Ithaca, NY, US)*

**License** Creative Commons BY 4.0 International license
© Joseph Y. Halpern
**Joint work of** Kaya Alpturer, Joseph Y. Halpern, Ron van der Meyden
**Main reference** Kaya Alpturer, Joseph Y. Halpern, Ron van der Meyden: "Optimal Eventual Byzantine Agreement Protocols with Omission Failures", in Proc. of the 2023 ACM Symposium on Principles of Distributed Computing, PODC 2023, Orlando, FL, USA, June 19-23, 2023, pp. 244–252, ACM, 2023.
**URL** https://doi.org//10.1145/3583668.3594573

Work on optimal protocols for Eventual Byzantine Agreement (EBA) – protocols that, in a precise sense, decide as soon as possible in every run and guarantee that all nonfaulty agents decide on the same value – has focused on full-information protocols (FIPs), where agents repeatedly send messages that completely describe their past observations to every other agent. While it can be shown that, without loss of generality, we can take an optimal protocol to be an FIP, full information exchange is impractical to implement for many applications due to the required message size. We separate protocols into two parts, the information-exchange protocol and the action protocol, so as to be able to examine the effects of more limited information exchange. We then define a notion of optimality with respect to an information-exchange protocol. Roughly speaking, an action protocol $\mathbf{P}$ is optimal with respect to an information-exchange protocol $\mathcal{E}$, if with $\mathbf{P}$, agents decide as soon as possible among action protocols that exchange information according to $\mathcal{E}$. We present a knowledge-based EBA program for omission failures all of whose implementations are guaranteed to be correct and are optimal if the information exchange satisfies a certain safety condition. We then construct concrete programs that implement this knowledge-based program in two settings of interest that are shown to satisfy the safety condition. Finally, we show that a small modification of our program results in an FIP that is both optimal and efficiently implementable, settling an open problem posed by Halpern, Moses, and Waarts [1].

### References

1 Joseph Y. Halpern, Yoram Moses, and Orli Waarts. A characterization of eventual Byzantine agreement. *SIAM Journal on Computing*, 31(3):838–865, 2001. doi:10.1137/S0097539798340217.

## 3.6 Hypergraphs for Knowledge

*Roman Kniazev (École Polytechnique – Palaiseau, FR)*

**License** Creative Commons BY 4.0 International license
© Roman Kniazev
**Main reference** Éric Goubault, Roman Kniazev, Jérémy Ledent: "A Many-Sorted Epistemic Logic for Chromatic Hypergraphs", in Proc. of the 32nd EACSL Annual Conference on Computer Science Logic, CSL 2024, February 19-23, 2024, Naples, Italy, LIPIcs, Vol. 288, pp. 30:1–30:18, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
**URL** https://doi.org//10.4230/LIPICS.CSL.2024.30

After recalling semantics of epistemic logic in epistemic coverings, we discuss how chromatic hypergraphs generalize them and can resolve some inconveniences appearing in practice.

## 3.7 Domain-theoretical Constraint Systems for Knowledge and Belief in Multi-agent Systems

*Sophia Knight (University of Minnesota – Duluth, MN, US)*

**License** Creative Commons BY 4.0 International license
© Sophia Knight
**Joint work of** Michell Guzmán, Stefan Haar, Sophia Knight, Catuscia Palamidessi, Prakash Panangaden, Salim Perchy, Santiago Quintero, Sergio Ramírez, Camilo Rueda, Frank Valencia
**Main reference** Michell Guzmán, Sophia Knight, Santiago Quintero, Sergio Ramírez, Camilo Rueda, Frank Valencia: "Reasoning About Distributed Knowledge of Groups with Infinitely Many Agents", in Proc. of the 30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands, LIPIcs, Vol. 140, pp. 29:1–29:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.
**URL** https://doi.org//10.4230/LIPICS.CONCUR.2019.29

We begin by reviewing the connection between topology and domain theory. We discuss the importance and difficulties in representing knowledge, belief, and information in multi-agent distributed systems, especially systems with fallible agents. Then we explain how to generalize concurrent constraint programming (CCP) with an underlying epistemic constraint system, representing agents' knowledge and information. We discuss some of the details of epistemic constraint systems and the more general multi-agent spatial constraint systems. We present a new semantics for CCP extended with epistemic and spatial constraint systems and an agent space operator. We discuss how to extend these systems to infinite groups of agents. In a purely logical setting, representing group, common, and distributed knowledge of infinite groups of agents presents difficulties, but in the domain-theoretic constraint system setting, these knowledge operators can be represented in a natural way. Finally, we discuss epistemic and spatial information systems, and their correspondence to constraint systems.

## 3.8 Impure Simplicial Complexes: Local View

*Roman Kuznets (TU Wien, AT)*

**License** Creative Commons BY 4.0 International license
© Roman Kuznets
**Joint work of** Hans van Ditmarsch, Roman Kuznets, Rojo Randrianomentsoa
**Main reference** Rojo Randrianomentsoa, Hans van Ditmarsch, Roman Kuznets: "Impure Simplicial Complexes: Complete Axiomatization", Logical Methods in Computer Science, Vol. 19(4), October 2023.
**URL** https://doi.org//10.46298/LMCS-19(4:3)2023

We discuss the choices that the local view suggests for impure simplicial complexes, i.e., complexes where some of the agents may have crashed. We settle on a three-valued semantics with the third value being "undefined" and provide a sound and complete axiomatization by means of a novel canonical simplicial model construction.

## 3.9 The Impossibility of Approximate Agreement on a Larger Class of Graphs

*Shihao (Jason) Liu (University of Toronto, CA)*

**License** Creative Commons BY 4.0 International license
© Shihao (Jason) Liu
**Main reference** Shihao Liu: "The Impossibility of Approximate Agreement on a Larger Class of Graphs", in Proc. of the 26th International Conference on Principles of Distributed Systems, OPODIS 2022, December 13-15, 2022, Brussels, Belgium, LIPIcs, Vol. 253, pp. 22:1–22:20, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
**URL** https://doi.org//10.4230/LIPICS.OPODIS.2022.22

Approximate agreement is a variant of consensus in which processes receive input values from a domain and must output values in that domain that are sufficiently close to one another. We study the problem when the input domain is the vertex set of a connected graph. In asynchronous systems where processes communicate using shared registers, there are wait-free approximate agreement algorithms when the graph is a path or a tree, but not when the graph is a cycle of length at least 4. For many graphs, it is unknown whether a wait-free solution for approximate agreement exists.

We introduce a set of impossibility conditions and prove that approximate agreement on graphs satisfying these conditions cannot be solved in a wait-free manner. In particular, the graphs of all triangulated $d$-dimensional spheres that are not cliques satisfy these conditions. The vertices and edges of an octahedron are an example of such a graph. We also present a reduction from approximate agreement on one graph to another graph. This enables us to extend known impossibility results to even more graphs. Finally, we show that extension-based proofs cannot be used to prove the impossibility of wait-free approximate agreement on any connected graph, which demonstrates the necessity of using combinatorial arguments.

## 3.10 Product Updates for Partial Epistemic Models and Logical Obstruction to Task Solvability

*Susumu Nishimura (Kyoto University, JP)*

**License** Creative Commons BY 4.0 International license
© Susumu Nishimura
**Joint work of** Masaki Muramatsu, Daisuke Nakai, Susumu Nishimura

We introduce the notion of partial product update. This refines the original notion of product update to encompass distributed tasks and protocols modeled by impure simplicial complexes. With this refined notion, we obtain a logical method that is generalized to allow the application of logical obstruction to show unsolvability results in a distributed environment where the failure of agents is detectable. We demonstrate the use of the logical method by giving a concrete logical obstruction and showing that the consensus task is unsolvable by single-round synchronous message-passing protocol.

## 3.11 Mixing Combinatorial and Point-set Topology

*Thomas Nowak (ENS – Gif-sur-Yvette, FR)*

**License** Creative Commons BY 4.0 International license
© Thomas Nowak
**Joint work of** Hagit Attiya, Armando Castañeda, Thomas Nowak
**Main reference** Hagit Attiya, Armando Castañeda, Thomas Nowak: "Topological Characterization of Task Solvability in General Models of Computation", CoRR abs/2301.13837, 2023.
**URL** https://doi.org//10.48550/ARXIV.2301.13837

The famous asynchronous computability theorem (ACT) relates the existence of an asynchronous wait-free shared-memory protocol for solving a task with the existence of a simplicial map from a subdivision of the simplicial complex representing the inputs to the simplicial complex representing the allowable outputs. The original theorem relies on a correspondence between protocols and simplicial maps in finite models of computation that induce a compact topology. This correspondence, however, is far from obvious for computational models that induce a non-compact topology, and indeed previous attempts to extend the ACT have failed.

We show first that in every non-compact model, protocols solving tasks correspond to simplicial maps that need to be continuous. This correspondence is then used to prove that the approach used in ACT that equates protocols and simplicial complexes actually works for every compact model, and to show a generalized ACT that applies also to non-compact computation models.

Our study combines combinatorial and point-set topological aspects of the executions admitted by the computational model.

## 3.12 A Speedup Theorem for Asynchronous Computation

*Ami Paz (CNRS – Gif-sur-Yvette, FR)*

**License** Creative Commons BY 4.0 International license
© Ami Paz
**Joint work of** Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum
**Main reference** Pierre Fraigniaud, Ami Paz, Sergio Rajsbaum: "A Speedup Theorem for Asynchronous Computation with Applications to Consensus and Approximate Agreement", in Proc. of the 2022 ACM Symposium on Principles of Distributed Computing, PODC 2022, Salerno, Italy, July 25 – 29, 2022, pp. 460–470, ACM, 2022.
**URL** https://doi.org//10.1145/3519270.3538422

Speedup theorems have recently gained increasing attention in studying distributed graph algorithms in synchronous systems. Using topological tools, we present a variant of this technique that applies to asynchronous shared-memory systems. At the core of our technique is a round reduction theorem: given a distributed task $\Pi$, we define a new task called the closure of $\Pi$, such that if $\Pi$ is solvable in $t$ rounds, then its closure is solvable in $t - 1$ rounds.

We illustrate the power of our speedup theorem by providing new proof of the wait-free impossibility of consensus using read/write registers. This is done merely by showing that the closure of consensus is consensus itself. The simplicity of our technique allows us to study additional communication objects, namely test&set and binary consensus. We analyze the approximate agreement task in systems augmented with the two objects and show that while these objects are more powerful than read/write registers from the computability perspective, they do not help to reduce the time complexity of solving approximate agreement.

## 3.13 On Simplicial Semantics

*Rojo Randrianomentsoa (TU Wien, AT)*

**License** Creative Commons BY 4.0 International license
© Rojo Randrianomentsoa
**Joint work of** Hans van Ditmarsch, Roman Kuznets, Rojo Randrianomentsoa
**Main reference** Hans van Ditmarsch, Roman Kuznets, Rojo Randrianomentsoa: "On Two- and Three-valued Semantics for Impure Simplicial Complexes", in Proc. of the Fourteenth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2023, Udine, Italy, 18-20th September 2023, EPTCS, Vol. 390, pp. 50–66, Open Publishing Association, 2023.
**URL** https://doi.org//10.4204/EPTCS.390.4

Impure simplicial complexes are used to model distributed systems with crashing agents. We discuss language extension for three-valued semantics, and provide the translations from that semantics to the two-valued facet semantics. In my talk, I present some of our results as well as some of the challenges we face in our current work.

## 3.14 Do Genes Argue? Abstract Argumentation and Boolean Networks

*David A. Rosenblueth (National Autonomous University of Mexico, MX)*

**License** Creative Commons BY 4.0 International license
© David A. Rosenblueth
**Joint work of** Eugenio Azpeitia, Stalin Muñoz, David A. Rosenblueth, Octavio Zapata

Stemming from the work by Phan Minh Dung in the 90s, there is a plethora of formalisms of abstract argumentation, from Dung's original work to abstract dialectical frameworks. Stemming from the work by François Jacob and Jacques Monod in the 60s, there are numerous mathematical models of gene regulation, from differential equations to Boolean networks. These two areas resemble each other: synchronous Boolean networks are reminiscent of abstract dialectical frameworks. The question is whether or not there is more than this resemblance. We will explore possibilities of bridging these two areas.

## 3.15 Logics for Data Exchange

*Sonja Smets (University of Amsterdam, NL)*

**License** Creative Commons BY 4.0 International license
© Sonja Smets
**Joint work of** Alexandru Baltag, Sonja Smets

We present a new family of dynamic logics that can model complex acts of data exchange and communication between different systems or agents. In the context of multi-agent systems, these are acts by which individual agents, as well as groups of agents, can publicly or privately access and transmit all the information stored at specific locations. In addition to having the full power of standard dynamic epistemic logics (DEL) to model acts of propositional communication, our logics can handle the type of communication protocols in which the data that is being communicated cannot be explicitly captured by a proposition.

## 3.16 Synergistic Knowledge

*Thomas Studer (Universität Bern, CH)*

**License** Creative Commons BY 4.0 International license
© Thomas Studer
**Joint work of** Christian Cachin, David Lehnherr, Thomas Studer
**Main reference** Christian Cachin, David Lehnherr, Thomas Studer: "Synergistic Knowledge", in Proc. of the 25th International Symposium on Stabilization, Safety, and Security of Distributed Systems, SSS 2023, Jersey City, NJ, USA, October 2-4, 2023, LNCS, Vol. 14310, pp. 552–567, Springer, 2023.
**URL** https://doi.org//10.1007/978-3-031-44274-2_41

In formal epistemology, group knowledge is often modeled as the knowledge that the group would have if the agents shared all their individual knowledge. However, this interpretation does not account for relations between agents. In this work, we propose the notion of synergistic knowledge, which makes it possible to model those relationships. As an example, we investigate the use of consensus objects.

## 3.17 Pattern Models and Consensus

*Diego A. Velázquez (National Autonomous University of Mexico, MX)*

**License** Creative Commons BY 4.0 International license
© Diego A. Velázquez
**Joint work of** Armando Castañeda, Hans van Ditmarsch, David A. Rosenblueth, Diego A. Velázquez

Recently, dynamic epistemic logic has started to be exploited for analyzing distributed systems. Pattern model logic was designed to that end. We present a simple-to-test sufficient condition that ensures consensus unsolvability for a given oblivious dynamic-network model.

# 4 Working Groups

The presentations and discussions in the unique setting of Dagstuhl, with its relaxed and stimulating atmosphere, fully achieved their purpose: Long discussions took place during the official seminar, and many fruitful cross-community interactions spontaneously occurred during the free times, which even exceeded the amount of available time.

As a conclusion, we are convinced that our seminar will contribute to the further development of epistemic reasoning and topology in distributed computing, and the integration of the respective approaches.

## 4.1 Extending Topology to Handle Data Structures

*Henning Basold, Armando Castañeda, Faith Ellen, Guy Goren, Shihao (Jason) Liu, Susumu Nishimura, Thomas Nowak, Sergio Rajsbaum, Ulrich Schmid*

License Creative Commons BY 4.0 International license
© Henning Basold, Armando Castañeda, Faith Ellen, Guy Goren, Shihao (Jason) Liu, Susumu Nishimura, Thomas Nowak, Sergio Rajsbaum, Ulrich Schmid

### 4.1.1 Context

The combinatorial topology approach to distributed computing has been shown to be powerful for studying solvability of fundamental distributed problems such a consensus, set agreement and renaming [42]. However, the current approach has a clear limitation as it can be applied only to distributed problems that can be modeled as (distributed) tasks. A *task* models a one-shot distributed problem where each process starts with a private input, and the processes are required to communicate with each other in order to decide private outputs that satisfy the task's specification. Nonetheless, there are fundamental distributed problems that cannot be specified as tasks [17, 18]. An important class of such problems are *concurrent objects* (or data structures), e.g., queues, stacks, trees and dictionaries. One seeks for implementations of them that are *linearizable* [43]. Roughly speaking, in a concurrent object, a process can perform an unbounded number of the operations provided by the object (differently from tasks where there is a single operation) and, in a linearizable implementation, responses to operations must be valid for the object that is implemented, somehow respecting the order operations are performed in a given execution.

### 4.1.2 Known results

Formally, a task is a triple $T = \langle I, O, \Delta \rangle$, where $I$ and $O$ are input and output complexes specifying the input and output domains of the task, and $\Delta$ is a carrier map specifying the complex $\Delta(\sigma) \subseteq O$ with valid outputs for each input simplex $\sigma \in I$. An execution $E$ is *valid* for the task if $\tau \in \Delta(\sigma)$, where $\sigma$ and $\tau$ are the simplexes with inputs and outputs in $E$, respectively.

Our discussion centered around finding a topological specification for concurrent objects. It first centered on previous work [17, 18], which already showed that tasks are too weak to even represent a one-shot queue, where each process can only access once by performing a fixed operation. The main problem is that carrier maps of tasks have limited ability to specify the order in which operations are performed in an execution.

The aforementioned work defined *refined tasks* (called *long-lived task* in [17]), where a vertex in the input complex represents a specific invocation to the object by a specific process. When using a simplex to represent an input configuration, multiple invocations of the same operation by the same process are represented by multiple distinct vertices in the simplex. Starting from any simplex in the input complex, each execution $E$ specifies a simplex $\tau_E$ where each vertex of $\tau_E$ consists of a process identifier, an invocation to the object by this process, the response to this invocation in $E$, and the *set-view* when the response was returned. A set-view is the set of all invocations (by all processes) on the object up to a particular point in the execution. The set of set-views in $\tau_E$ fully captures the interleaving of operations in the execution $E$. The output complex consists of simplexes $\tau_E$ specified by all possible executions (concurrent and sequential) $E$ whose operations appear in simplexes of the input complex. A refined task is a triple $T = \langle I, O, \Delta \rangle$ where $\Delta(\sigma) \subseteq O$ contains simplexes $\tau_E$ specified by all possible sequential $E$ where processes perform the operations in a simplex $\sigma \in I$, in some order. A concurrent execution $E'$ with operations in a simplex $\sigma \in I$ is considered *valid* for the refined task if there exists a simplex $\tau_E \in \Delta(\sigma)$ such that (1) the outputs to operations in $E'$ and $E$ agree, and (2) the set-view of each vertex in $\tau_{E'}$ contains the set-view of its corresponding vertex in $\tau_E$.

Refined tasks have been shown to be expressive enough to model any linearizable concurrent object. However, this does not seem to resolve the limitation mentioned above because the notion of solvability for refined tasks is based on two components: outputs and set-views. The problem is that set-views are features of the execution, i.e., they are not produced or implied by processes' states. It is not clear how set-views can be captured by simplicial maps in a way that the combinatorial topology approach can be applied to study solvability of refined tasks. Let us recall that in the combinatorial topology approach protocols are modeled as simplicial maps that specify outputs as function on *only* process states, and the main question that needs to be answered is whether there exists a simplicial map from the protocol complex to the output complex, respecting the specification of the task in question. It may be possible that refined tasks can be used to study solvability of linearizable concurrent object $S$ in an indirect manner as suggested in [33], where a task $T$ is constructed such that if there is a linearizable implementation of $S$ then $T$ is solvable.

### 4.1.3 Proposed approaches

The group discussed two possible alternative approaches to specify concurrent objects that may be more suitable for the combinatorial topology approach. Applications of these approaches remain to be seen.

#### 4.1.3.1 Approach I: using a simplicial set to represent the possible outputs instead of an simplicial complex.

Here each vertex in the input complex $I$ is defined by a process and a sequence of operations to be performed by the process. Each vertex in the output simplicial set $O$ is defined by a process, a sequence of operations performed by the process, and the sequence of responses to these operations. Each facet of the simplicial set is labelled by a linearization, which is a total order of the operations in all of its vertices. The linearization labelling each facet needs to respect the order of operations within each of its vertices. A map $\Delta$ specifies the subset of allowed output simplexes for each input simplex. Starting from any simplex $\sigma \in I$, an execution $E$ with inputs described by $\sigma$ is *correct* if there exists a linearization of the operations in $E$ and a simplex $\tau \in O$ labelled by this linearization such that, for each process, $p$, there is a vertex in $\tau$ with process identifier $p$, the sequence of operations performed by $p$ during $E$, and the sequence of responses $p$ obtained during $E$.

#### 4.1.3.2 Approach II: modifying $\Delta$ so that it depends on the partial order of operations.

For simplicity, we first consider only one-shot objects. To represent the order in which operations are performed on a one-shot object, we can use a partial order on the set of processes. Specifically, if the operation by process $p$ finishes before the operation by process $q$ begins, then $p \prec q$. If the operation by process $p$ is concurrent with the operation by process $q$, then $p$ and $q$ are incomparable.

Here a one-shot object is represented by a tuple $(I, \Pi, O, \Delta)$. $I$ is the input complex. Each vertex in $I$ is of the form $(id, op)$, where $id$ is a process identifier and $op$ is an operation to be performed by this process. $\Pi$ is the set of all possible partial orders on the set of processes. $O$ is the output complex. Each vertex in $O$ is of the form $(id, op, resp)$, where

$(id, op)$ is a vertex in $I$ and $resp$ is a possible response to the operation $op$ performed by the process with identifier $id$. $\Delta : I \times \Pi \to 2^O$ is a chromatic map which maps a simplex $\sigma \in I$ and a partial order $\pi \in \Pi$ to a subcomplex of $O$. If the object has a sequential specification, then each simplex in this subcomplex describes an output configuration resulting from a sequential execution whose order of operations respects the given partial order. An algorithm correctly implements an object if, for every execution $E$, there exists a simplex of $\Delta(\sigma, \pi)$ that describes the response to each operation performed during $E$, where $\sigma$ is the simplex of $I$ describing the operation each process performs during the execution and $\pi \in \Pi$ describes the partial order of these operations.

Now, consider any execution in which each process may access the object multiple times. For each process $p$, let $n_p$ denote the number of operations performed by $p$ during this execution. We can use a partial order on the set $\{(p, i) \mid p \text{ is a process and } 1 \leq i \leq n_p\}$ to describe the relevant information about the operations performed during the execution. Specifically, if the $i$th operation by process $p$ finishes before the $j$th operation by process $q$ begins, then $(p, i) \prec (q, j)$. If the $i$th operation by process $p$ is concurrent with the $j$th operation by process $q$, then $(p, i)$ and $(q, j)$ are incomparable.

Vertices in the input complex $I$ are now of the form $(p, seq_p)$, where $seq_p$ is a sequence of operations to be performed by the process $p$. Likewise, vertices in the output complex $O$ are now of the form $(p, seq_p, resp)$, where $resp$ is a sequence of possible responses to the operations in $seq_p$. $\Delta$ now maps a simplex in $I$ and a partial order on the set $\{(p, i) \mid p \text{ is a process and } 1 \leq i \leq |seq_p|\}$ to a subcomplex of $O$.

Finally, it was also briefly discussed that the order of operations in the two proposed approaches could be captured through *higher-dimensional automata* (HDA) [53], which have an inherent geometric nature that may facilitate the use of combinatorial topology techniques.

## 4.2 Interpreted Systems and Dynamic Epistemic Logic

*Alexandru Baltag, Hans van Ditmarsch, Krisztina Fruzsa, Joseph Y. Halpern, Yoram Moses, Hugo Rincón Galeana, Ulrich Schmid, Sonja Smets, Diego A. Velázquez*

**License** Creative Commons BY 4.0 International license
© Alexandru Baltag, Hans van Ditmarsch, Krisztina Fruzsa, Joseph Y. Halpern, Yoram Moses, Hugo Rincón Galeana, Ulrich Schmid, Sonja Smets, Diego A. Velázquez

The Interpreted Systems (IS) approach ([31]) has, over the last four decades, been used to study distributed systems and has served as a tool in the design and analysis of distributed and multi-agent protocols and systems. The discussion began in this with the question of whether and how Dynamic Epistemic Logic (DEL) can contribute in a similar fashion.

Broadly speaking, we thus considered the following specific questions:

- What are the main differences between IS and DEL modeling?
- What could DEL add to IS modeling?
- What could IS add to DEL modeling?

The (partial) answers developed in this discussion group will be summarized in dedicated subsections below.

### 4.2.1 What are the main differences between the IS and DEL modeling?

While both DEL and IS provide means to track the manner in which knowledge and belief change over time in a multi-agent system, the two approaches differ dramatically in the way that they model the effect that actions have. In DEL, actions operate directly on the epistemic state of the system (this epistemic state is represented by a Kripke structure). For every action, the framework must provide an *action model*, which maps a Kripke structure and an action to a new Kripke structure, thereby describing the manner in which the epistemic state changes when the action is performed. In IS actions are taken to modify the physical configuration (or global state) of the system. As in automata theory, this is modeled by way of a transition function, mapping a configuration and an action to the configuration that results from performing the action. The epistemic state of the system in IS is also modeled using a Kripke structure, but this structure is induced by the system configurations, where each agent's indistinguishability relation is determined by its local state.

In the IS way of modeling of distributed systems, based on the runs-and-systems framework [31], agent protocols and the environment are outside of the formal logic used for epistemic reasoning. In fact, they are represented using *transition systems* (TS) which are the primary objects of study. Since the executions of a distributed protocol can typically be described in terms of a transition system (see [47]), the IS approach can readily be applied to a broad range of distributed systems problems.

In DEL [26], agent protocols (as well as the environment) can be expressed in the formal logic used for epistemic reasoning. In particular, action models enable one to model many forms of communication between the agents while focusing on capturing the resulting knowledge evolution of the involved agents. Moreover, precondition formulas can be used to tie communication actions to particular epistemic states of the agents (that are necessary to hold when communicating). Works such as [37], which even connect DEL and combinatorial topology, show that the idea of using DEL to help with topological reasoning about distributed systems is promising.

DEL and IS have been used for the analysis of communication protocols. Specific features of the protocols themselves indicate which framework fits best. The features we are looking at are a) temporal aspects (synchronous versus asynchronous communications) and b) the private versus public nature of the information contained in the events or actions that happen as part of the protocol. Here, we note that in DEL, time is not explicitly represented. Embeddings from DEL into IS (or rather into temporal epistemic logics) have been proposed. This includes both synchronous embeddings [10], wherein the execution of an action corresponds to a clock tick, and asynchronous embeddings [21], wherein histories of actions of variable length correspond to time in some way. While some believe that DEL and IS can be extended with ingredients so they can model the same kind of protocols, the starting base of these frameworks does differ. The built-in features of the starting base (such as the temporal aspects in IS and the action models in DEL) will make the application of one or the other framework a more straightforward choice for a given type of protocol. For instance when we focus on cryptographic protocols with complex cryptographic primitives as part of the actions or events that happen in the protocol, we see an added advantage of starting with DEL which offers us the expressive power of modeling the security aspects in the actions separately and allows us to automatically compute the state-transitions via its product update mechanism. DEL is used in the study of cryptographic protocols in works like [44, 20, 23, 30].

An advantageous feature of DEL is the fact that its built-in model update mechanism makes it immediately well-suited for automated model checking tasks. The implementation was done in different programming languages, but the most elaborate tool uses Haskell in DEMO [29]. The power of DEMO for model checking tasks has been illustrated on a small number of well-known sample protocols including comparison to temporal epistemic model checkers like MCK [27], which is based on IS modeling.

A disadvantage of DEL in general is the potential blowup of the model (or the precondition formulas). There are various reasons for such a blowup. As the models used in DEL are small (updates, and their temporal interpretation, are computable from an initial model and the sequence of actions), complexities with respect to model size are high [3]. Another reason is the iteration of action models, which leads to exponential growth of the size with each model update. This can be addressed with symmetry reductions [34], restriction to point-generated submodels [12], dedicated tools for specific frame classes (such as models with equivalence relations) instead of arbitrary binary relations [29], and bisimulation contraction [12].

A feature of DEL often forgotten is that such logics and frameworks are essentially logics of *observation*, not of action. One can investigate the consequence of actions if they were to happen, but one cannot state "let this action happen now". This is of course very different from IS. Once we add factual change to the DEL framework, this drawback is somewhat overcome. Now there are actions such as turning on and off the light, setting values of keys, etcetera. But agency remains lacking.

### 4.2.2 What could DEL add to IS modeling?

Operationalizing the implicit Kripke model construction resulting from executing the system primarily requires the construction of action models for typical communication protocols considered in epistemic analysis via IS. The challenge here is to develop action models that avoid state explosion. One promising starting point could be communication patterns, updates that compare in interesting ways to action models [55, 15, 14]. The availability of such action models would open up some interesting avenues such as:

- Could correctness proofs of protocols formulated via suitable precondition formulas of such action models be automated, via analyzing the Kripke model obtained by the product updates w.r.t. satisfying the desired properties? This would replace the translation of knowledge-based protocols obtained via IS to TS, and the obligatory proof that the resulting protocol guarantees that the agents attain their respective necessary level of knowledge.

  Good target practice will be to model the consensus protocol for processes with send-omission faults from [1] in DEL, by (i) defining an action model for send-omission faults, (ii) modeling the protocol implementations via suitable precondition formulas, and (iii) generating and analyzing the Kripke model after suitably many rounds. We were optimistically thinking of including these results already in the report, however that was a bridge too far. Still, we are committed to this and other benchmarks to compare DEL and IS.
- An area in DEL that is highly targeted towards protocols and automation is that of epistemic planning [13, 6, 48], and in particular concerning decidability of epistemic planning problems [2, 46]. Could epistemic planning techniques be used for synthesizing protocols in a distributed setting?

### 4.2.3 What could IS add to DEL modeling?

The IS approach has been used to provide insights into the interaction between knowledge and communication in distributed systems ([39, 19, 7]), has allowed to generalize existing protocols ([40]), and to design optimal and state-of-the-art protocols solving well-known problems ([28, 50, 35, 16, 1]). These works provide insight into the role of synchrony and partial synchrony, refine our analysis of achievable types of optimality in different models of distributed computation, and into various aspects of fault-tolerance. The Knowledge of

Preconditions principle [49], which is naturally formulated in IS, provides model-independent theorem capturing the connection between knowledge and action in multi-agent systems. It has been shown to facilitate the design and analysis of efficient distributed systems protocols. Although protocols are not explicitly represented in the IS framework, it is possible to consider high-level "knowledge-based programs" in this setting, in which actions depend on tests for knowledge. In general, however, such programs are specifications and do not always uniquely determine the agents' behavior. Finding ways to implement them may be challenging. The more developed literature applying IS for the study of distributed computing can serve as an example and inspiration for directions in which DEL may be expanded. Fundamentally, DEL can be viewed as a subclass of IS. Its added declarative nature described above comes at a cost, however. Modeling and studying aspects of timing, including the role of synchrony, asynchrony and partial synchrony have been carried out successfully using the IS framework, while handling asynchrony, and partial synchrony appear to pose significant challenges for DEL. Similarly, incorporating failures, which have been studied in many IS settings, seems likely to give rise to cumbersome DEL models.

In order to be able to apply the IS framework, one only needs a description of the set of possible histories of a system, and a clear definition of the local states of agents in every configuration. As mentioned above, the fact that systems are often described by way of a transition system obviates the need to translate a given system into the formalism. Treating, say, an existing internet protocol using DEL would typically require a significant reformulation effort. Making DEL more flexible in this sense seems to be a considerable albeit worthwhile challenge. Finally, while the IS framework has given rise to new and improved protocols as well as tight bounds and impossibility results in different models of distributed computing, this has not been a focus of DEL works. It will be very interesting to see how well DEL can match such tasks, in light of the results obtained using IS.

## 4.3 Representing Epistemic Attitudes via Simplicial Complexes

*Alexandru Baltag, Henning Basold, Hans van Ditmarsch, Roman Kniazev, Sophia Knight, Roman Kuznets, Jérémy Ledent, David Lehnherr, Rojo Randrianomentsoa, Sonja Smets, Thomas Studer*

**License** 🅭 Creative Commons BY 4.0 International license
© Alexandru Baltag, Henning Basold, Hans van Ditmarsch, Roman Kniazev, Sophia Knight, Roman Kuznets, Jérémy Ledent, David Lehnherr, Rojo Randrianomentsoa, Sonja Smets, Thomas Studer

The general header 'Epistemic Attitudes via Simplicial Complexes' permits a wealth of topics, also involving group epistemics such as distributed knowledge and common knowledge, and generalizations of simplicial complexes such as simplicial sets. However, the overall focus of the discussions turned towards a single issue: how to model *belief* on simplicial complexes, and, to a somewhat lesser extent, how to model dynamic epistemic attitudes such as deceit and Byzantine behaviour in general. In this report we focus on belief, the structures needed to encode it, their further generalization to polychromatic simplicial complexes (where simplexes can contain multiple vertices with the same colour), and the use of those for neighbourhood (epistemic) semantics on complexes.

The discussion was very much example driven. We therefore give a minimum of formal precision concerning complexes and their epistemic semantics. A *simplicial complex* $C$ is a set of subsets of a set $V$ of vertices that is downward closed, and a simplicial model is a triple $(C, \chi, \ell)$ where $\mathcal{C}$ is a simplicial complex, $\chi$ is a *chromatic map* assigning an agent (or a colour) to each vertex and such that different vertices in a face are assigned different colours, and $\ell$ is a *valuation* of propositional variables to vertices (or, alternatively, to faces). A maximal face is a facet. Intuitively a vertex represents the local state of an agent and a facet represents a global state. In a *simplicial set* there can be multiple occurrences of a simplex (thus it is a multi-set of subsets and not a set).

The semantics of knowledge on simplicial complexes replaces the usual Kripke model semantics with equivalence classes for agents $a$, given binary (equivalence) relations, with adjacency of facets (or faces) to a vertex with colour $a$. The basic definition is as follows, where we refer to works like [45, 37, 24, 36] for details. Let $X$ be a facet in $C$ ($X \in \mathcal{F}(C)$) of simplicial model $\mathcal{C} = (C, \chi, \ell)$. We consider formulas in a standard epistemic logical language. Then:

$$\mathcal{C}, X \models K_a\varphi \text{ iff } \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ such that } a \in \chi(X \cap Y).$$

In fact we can interpret formulas in arbitrary faces of a complex[1] where the case of a vertex, a singleton face, captures the knowledge semantics as follows, where $\chi(v) = a$:

$$\mathcal{C}, v \models K_a\varphi \text{ iff } \mathcal{C}, Y \models \varphi \text{ for all } Y \in C \text{ with } v \in Y$$

We will use a running example to demonstrate all semantics for knowledge, belief, and other epistemic attitudes. And we start by explaining the knowledge semantics. Hopefully, this will suffice for the reader to prepare the ground for the subsequent different belief semantics.

Consider a simplicial model $\mathcal{C}$ as above for three agents (colours) $a, b, c$, where for convenience the vertices are named $xy$ for $x = a, b, c$ and $y = 0, 1, 2$ and where these numbers simultaneously represent local values, formalized as propositional variables $y_x$ only true there. (So in facet $X$ variables $0_a$, $0_c$, and $1_b$ are true and $1_a, 2_a, 1_c, 0_b, 2_b$ are false.) We also assume that variables are local for vertices but that does not seem to be of prime importance (one could think of assigning variables to facets).

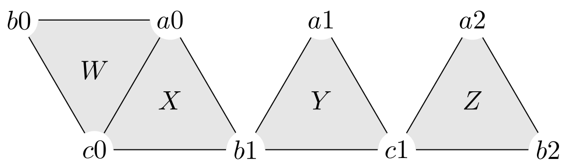


$\mathcal{C}^i$

We now have that in facet $X$ agent $a$ knows that the value of $c$ is 0, because in the facets $W$ and $X$ intersecting with $X$ in vertex $a0$ the value of $c$ is 0. However, agent $b$ is uncertain there whether the value of $c$ is 0 or 1 (where we use $\widehat{K}_b$ to abbreviate $\neg K_b \neg$):

$$\mathcal{C}, X \models K_a 0_c \qquad \mathcal{C}, X \models \widehat{K}_b 0_c \wedge \widehat{K}_b 1_c$$

### 4.3.1 Irrevocable Belief

Let us now model belief ($B_a$) instead of knowledge ($K_a$). Unlike knowledge, belief may be incorrect. A common way to model belief in modal logic, interpreted on Kripke models, is by way of relations that are serial, transitive, and Euclidean, for which the logic is KD45. That is, it satisfies:

[1] To interpret formulas in arbitrary faces we use the multi-pointed semantics of [24], section 'Local semantics for simplicial complexes', wherein $\mathcal{C}, \boldsymbol{X} \models \varphi$ for a set of faces $\boldsymbol{X}$ iff $\mathcal{C}, X \models \varphi$ for all $X \in \boldsymbol{X}$. In particular, $\mathcal{C}, v \models \varphi$ then means that $\mathcal{C}, \mathsf{star}(v) \models \varphi$, where $\mathsf{star}(v) = \{Y \in C \mid v \in Y\}$.

- $B_a\varphi \to \widehat{B}_a\varphi$ (consistency)
- $B_a \to B_a B_a\varphi$ (positive introspection)
- $\widehat{B}_a\varphi \to B_a\widehat{B}_a\varphi$ (negative introspection)

This is known as consistent belief [25] and also as irrevocable belief (or as conviction) [54]. Recalling the KD45 "balloons" (clusters of indistinguishable worlds) with "ropes" (asymmetric pairs in the accessibility relation) from isolated worlds pointing to these clusters, a way to model belief on complexes is with a *belief function* $f : \mathcal{V}(C) \to \mathcal{V}(C)$ that is an idempotent and color-preserving vertex map. It is not a chromatic map as it may not preserve other simplices. Let $X_a$ be the vertex of color $a$ that belongs to a facet $X$. Belief is now defined as

$$\mathcal{C}, X \models B_a\varphi \text{ iff } \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ with } f(X_a) \in Y$$

If, for a world $v$, $f(v) = v$, belief becomes knowledge. The map needs to be idempotent to ensure that the properties of belief hold, we can see the map $f$ as "pointing" vertex $v$ to a cluster of $f(v)$ adjoining facets, in words referring to a corresponding Kripke model, to a set of $b$-indistinguishable worlds.

As an example, we reconsider the model above but now enriched with a belief function that maps vertex $a0$ to vertex $a2$ and otherwise maps all vertices to themselves, as below.

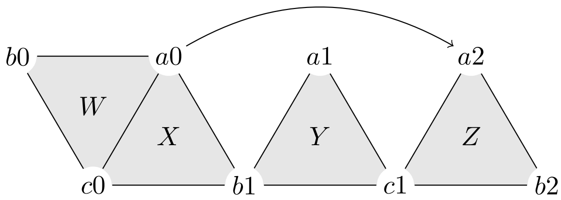


$\mathcal{C}^{ii}$

This represents that in local state $a0$ agent $a$ believes that her local state is $a2$ and that the global state is $Z$. Observe that this also implies that $a$ incorrectly believes that the value of her local state is 2 and not 0. However, a reasonable assumption may be that an agent correctly knows its local state, and may only be mistaken about what it believes to be the local states of other agents. (This assumption underlies some of our later epistemic explorations, wherein different vertices of the same colour jointly represent an agent's beliefs.) We now have that, for example,

| | |
|---|---|
| $\mathcal{C}, X \models 1_b \land \neg\widehat{B}_a 1_b \land B_a 2_b$ | agent $a$ incorrectly believes that the local state of $b$ is 2 |
| $\mathcal{C}, Z \models 2_b \land B_a 2_b$ | but $a$ *knows* that the local state of $b$ is 2 in global state $Z$ |
| $\mathcal{C}, X \models B_b 1_b \land B_b B_a 2_b$ | agent $b$ knows that agent $a$ has incorrect beliefs about him |

and so forth ... Note that we only have one modality in the logical language here, the informal "agent $a$ knows" and "agent $b$ knows" above are only written because in the facets (global states) where these formulas are interpreted the agents' beliefs are correct.

As a side issue, let us now consider the belief function $g$ such that $g(a0) = a2$, $g(c0) = c1$, and otherwise the identity. In the model $\mathcal{C}^i$, the distributed knowledge of agents $a$ and $c$ in facet $W$ is represented by the edge $\{c0, a0\}$. Whereas in the model $\mathcal{C}^{ii}$ but with belief function $g$, agents $a$ and $c$ have distributed belief embodied in the edge $\{c1, a2\}$. Group epistemic notions other than for knowledge were not discussed, but might worth to be investigated later.

The belief function is an enrichment of the simplicial model. It does not, on first sight, seem very topological in nature, and therefore not very simplicial. It rather looks like an accessibility relation between the worlds of a Kripke model. In fact, taken as a relation, a belief function is serial, transitive, and Euclidean, just as required for belief. Let us see how we can stretch this analogy a bit further before we return to more topological interpretations, in Subsection 4.3.3.

### 4.3.2 Revocable Belief, Safe Belief, and Knowledge

Now assume any function that is a transitive relation. Consider the belief function $f$ that maps vertex $a0$ to $a1$ and vertex $a1$ to $a2$, and therefore, by transitivity, $a0$ to $a2$. As follows (transitivity is implicit):

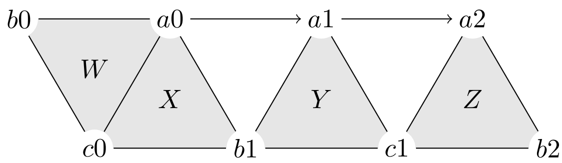


$\mathcal{C}^{iii}$

This induces a (strict) total order on the $a$ vertices such that $a0 > a1 > a2$, with an associated weak total order $\geq$. We can now interpret this total order as sorting comparable $a$-vertices into more and less plausible, where in this case $a2$ is most plausible, $a1$ somewhat less, and $a0$ least, by analogy of the plausibility models of [4] (and not dissimilar to [8, 22], see also [11]). Let us additionally write $a0 \geqq a2$ to indicate that $a2$ is the most plausible local state for $a$ in this total order (or, dually, $\leq$ and $\leqq$), $\equiv$ for equally plausible, and $\sim$ for comparable.

We now can assume a multi-modal language and define notions of belief (revocable/defeasible belief), safe belief (strong belief), and knowledge. Depending on one's preferred dialect this can be done in different ways. In the semantics for $a$-adjoining facets we would get, where $X \in \mathcal{F}(C)$:

$$\begin{array}{lll}
\mathcal{C}, X \models B_a\varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ and } v \leqq X_a \text{ with } v \in X \\
\mathcal{C}, X \models \Box_a\varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ and } v \leq X_a \text{ with } v \in X \\
\mathcal{C}, X \models K_a\varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ and } v \sim X_a \text{ with } v \in X
\end{array}$$

The relation $\geq$ between $a$-vertices induces a relation $\geq_a$ between facets, namely such that $X \geq_a Y$ if $X_a \geq Y_a$, and analogously we can induce relations $\geqq_a$, $\equiv_a$ and $\sim_a$. Note that as a consequence all facets containing a given local state for agent $a$ are equally plausible for her. Using this relation between facets, for the semantics of defeasible belief $B_a$ we would now get (where those for $\Box_a$ and $K_a$ are analogous) the following equivalent semantics, where for good measure we have added two more equivalent formulations. Which one looks most simplicial?

$$\begin{array}{lll}
\mathcal{C}, X \models B_a\varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in \mathcal{F}(C) \text{ with } Y \leqq_a X \\
\mathcal{C}, x \models B_a\varphi & \text{iff} & \mathcal{C}, Y \models \varphi \text{ for all } Y \in C \text{ with } v \leqq x \text{ and } v \in Y \\
\mathcal{C}, x \models B_a\varphi & \text{iff} & \mathcal{C}, v \models \varphi \text{ for all } v \in \mathcal{V}(C) \text{ with } v \leqq x
\end{array}$$

As an example of the semantics above we can easily determine in the simplicial model $\mathcal{C}^{iii}$ that

$$\begin{array}{ll}
\mathcal{C}^{iii}, X \models B_a 2_b & \text{Note } a2 \leqq a0 = X_a \text{, and } Z \text{ is the only facet containing } a2. \\
\mathcal{C}^{iii}, Y \models \Box_a \widehat{B}_c 1_b & \text{Belief in } \widehat{B}_c 1_b \text{ is safe for } a \text{ in } Y \\
 & \quad \text{as } Y, Z \leq Y \text{ and both satisfy } \widehat{B}_c 1_b. \\
\mathcal{C}^{iii}, Y \not\models \Box_a 2_b & \text{But belief in } 2_b \text{ is not safe for } a \text{ as } Y \not\models 2_b. \\
\mathcal{C}^{iii}, X \not\models K_a(0_b \vee 1_b) & \text{This is because } Z \leq_a X \text{ and } Z \not\models 0_b \vee 1_b.
\end{array}$$

This notion of knowledge is different from the standard notion of simplicial knowledge applied in model $\mathcal{C}^i$. We recall that $\mathcal{C}^i, X \models K_a(0_b \vee 1_b)$, as knowledge in $X$ means truth in $W$ and $X$. Whereas $\mathcal{C}^{iii}, X \not\models K_a(0_b \vee 1_b)$, as this 'novel' knowledge in $X$ means truth in $W, X, Y, Z$, and $0_b \vee 1_b$ is false in $Z$.

This can be adjusted by ***tentatively*** proposing a *belief quotient* of a simplicial model with respect to a belief function inducing a total order $\leq$ as above. As we will see this may not always be a simplicial model, but it seems it always will be a simplicial set. The belief quotient of a complex $\mathcal{C} = (C, \chi, \ell)$ would be $\mathcal{C}_\leq = (C_\leq, \chi_\leq, \ell_\leq)$ as follows. Let for any $v \in \mathcal{V}(C)$ vertex $v_\leq$ be defined as $\{w \in \mathcal{V}(C) \mid w \sim v\}$, then $\mathcal{V}(C_\leq) := \{v_\leq \mid v \in \mathcal{V}(C)\}$. Note that the choice of vertex determines its colour, so $v_\leq$ is the equivalence class of all vertices $w$ that are more or less plausible then $v$ for agent $\chi(v)$, so there is no need for an '$a$' as in $\leq_a$. The faces of $C_\leq$ now are $Y$ for which there is a $Z \in C$ with $Y = \{v_\leq \mid v \in Z\}$. As comparable vertices may have different labelings of propositional variables, proposing $\ell_\leq(v_\leq) := \bigcup\{\ell(v) \mid w \sim_a v\}$ may not be very useful. However, we recall that requiring labelings to be the same for comparable vertices was already mentioned as a reasonable assumption for belief functions. Let us therefore assume that all $\leq$-comparable vertices have the same labeling. We now can produce some nice examples.

In $\mathcal{C}^i$, the variables $0_a, 1_a, 2_a$ were only true in $a0, a1, a2$, respectively. Let us now, instead, assume that they are true in all three $a0, a1, a2$ so that $0_a \leftrightarrow 1_a$ and $1_a \leftrightarrow 2_a$ on the model. We now simply take $a0_\leq$ to be $a0$ itself. Given that, and some visual trickery, the belief quotient of $\mathcal{C}^{ii}$ is $\mathcal{C}^{iv}$ below.

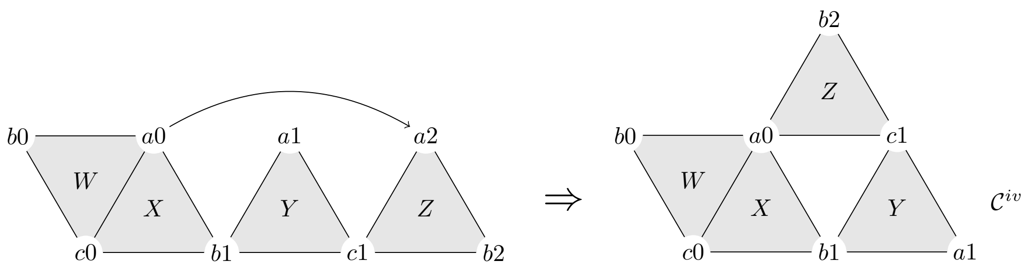


Similarly, the belief quotient of $\mathcal{C}^{iii}$ is $\mathcal{C}^v$ below, where we note that this is a simplicial set, as the edges $\{a0, b1\}$ and $\{a0, c1\}$ occur twice in the complex (denoted by the double-lined edge).

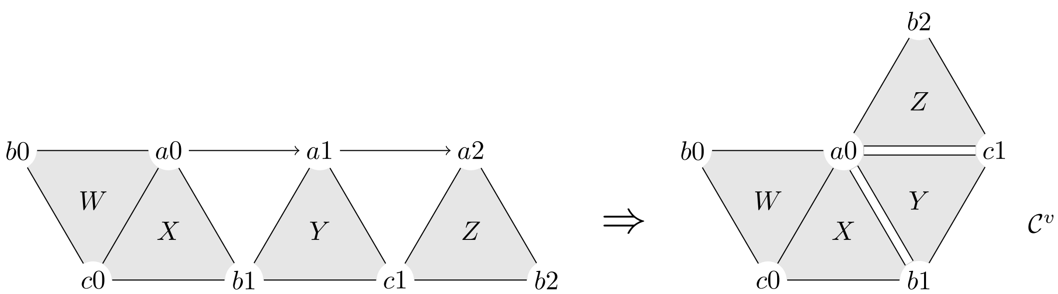


This representation does not make any difference for the semantics, but we are now back to a possible more pleasing, and familiar, notation for more and less plausible facets adjoining a vertex. In $\mathcal{C}^{iv}$ we have that vertex $a0$ intersects with $W, X, Z$ where $W$ and $X$ are equally plausible and $Z$ is more (and most) plausible. In $\mathcal{C}^v$ we have that vertex $a0$ intersects with $W, X, Y, Z$ where $W$ and $X$ are equally plausible, both are more plausible than $Y$, which is more plausible than $Z$.

### 4.3.3 Polychromatic Simplicial Complexes

A simple belief function for irrevocable belief, or a more complex belief function to model defeasible or strong belief, or even knowledge in a different way, is not a very topological tool. Nor is ordering the different facets intersecting in a vertex in the belief quotient version. It was widely discussed how to make such ideas into something more topological. Reconsidering the simplicial model $\mathcal{C}^{ii}$ once more, instead of considering pairs $(a0, a1), (a1, a2)$ (and $(a0, a2)$) in a relation, we can also consider such a pair as an edge, which is topology.

A structure where both vertices of an edge have the same colour would no longer be a chromatic simplicial complex, where in any simplex all vertices must have different colours. We tentatively call complexes where the $\chi$ map may assign the same colour to more than one vertex in a simplex a *polychromatic simplicial complex.*

Starting out with edges $\{a0, a1\}$ and $\{a0, a1\}$ in the model $\mathcal{C}^{iii}$ we would get the model $\mathcal{C}^{vi}$ as below.

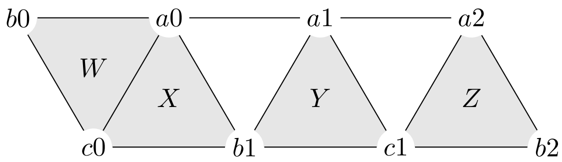


$\mathcal{C}^{vi}$

This model $\mathcal{C}^{vi}$ allows for additional epistemic notions, using the novel topological information. For example, reminiscent of the plausibility order (but also essentially different, as we will see) consider the shortest $a$-path linking vertices. Given $a0$, $a1$ is one step away, and $a2$ is two steps away. Based on that, and somewhat dual to the prior $\mathcal{C}^{iii}$, we can interpret this as agent $a$ considering $a0$ its most plausible local state, $a1$ less, and $a2$ least (the dual where $a2$ is most and $a0$ least plausible from the perspective of $a0$ seems less intuitive in this interpretation). Similarly we can order facets that way, from the perspective of a given agent $a$, where the distance between facets is the length of the shortest $a$-path between them. We now interpret $B_a\varphi$ as before, as truth in the most plausible facet or vertex.

Note that there is a difference with the previous semantics for total $\leq$ (and $\leq_a$) orders: the truth of $B_a\varphi$ depends on the actual facet. This is unlike the notions of belief and knowledge employing plausibility models in the previous subsection. These are subjective notions that do not depend on the actual point of evaluation (however, unlike strong belief, that is an 'objective' notion as well, that is, a notion depending on the point of evaluation given an equivalence class for an agent).

In model $\mathcal{C}^{vi}$ we would now have that $\mathcal{C}^{vi}, X \models B_a 0_c$, whereas $\mathcal{C}^{vi}, Y \models B_a 1_b$, and $\mathcal{C}^{vi}, Z \models B_a 2_b$ (where all these are distinguishing formulas in the model, that is, only true in those facets and in none other).

If we let pair $(a0, a2)$ in (relation induced by) the belief function count as well, we get this model $\mathcal{C}^{vii}$ and there seems nothing against considering them as the three edges of a triangle, as in $\mathcal{C}^{viii}$:

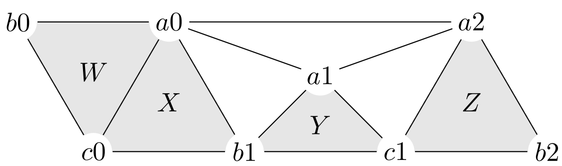


$\mathcal{C}^{vii}$

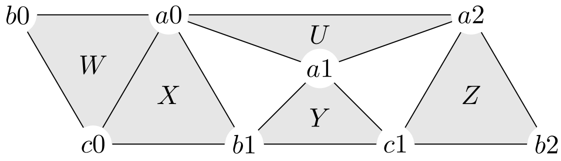


$\mathcal{C}^{viii}$

In this more topological interpretation involving simplexes (edges or triangles, here) of vertices with the same colour, we can maybe somehow interpret this as different mental states of the same agent, a structured form of the beliefs the agent considers possible. The representation clearly allows agent $a$ to reason about what her beliefs are in vertex $a0$ as different from what her beliefs are in vertex $a1$, or $a2$. Relations and generalizations to (semi-simplicial sets and) simplicial sets are again conceivable. So far, only semi-simplicial sets have been studied.

It would also be interesting to investigate what happens if we allow *higher-dimensional paths*, such as the chaining in order to obtain *common distributed knowledge* [5, 24].

Finally, how can we interpret formulas on simplexes where some but not all vertices are of the same colour, such as the one below?

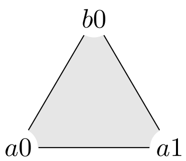


$\mathcal{C}^{ix}$

Another surely fertile direction of research would be to consider models like $\mathcal{C}^{vii}$ and $\mathcal{C}^{viii}$ in a neighbourhood semantics [41, 51], a common framework in modal logical semantics although infrequently applied to epistemic notions [9]. Exploring neighbourhood semantics might lead to new insights on *coalition logic* which found recent application to the specification and verification of smart contracts [52].

In neighbourhood semantics, going into simplicial mode, we would have that $\mathcal{C}, X \models K_a\varphi$ if there is a *neighbourhood* (a set of facets) of $X$ for agent $a$ such that $\varphi$ is true in all the facets in that neighbourhood. Or, possibly more intuitively, given a vertex $v$ with colour $a$, $K_a\varphi$ is true in $v$ if there is an $a$-neighbourhood of $v$ where $\varphi$ is true.

In $\mathcal{C}^{vii}$ (or $\mathcal{C}^{viii}$) facet $W$ has three neighbourhoods for $a$: $\{W, X\}$, $\{Z\}$, and $\{Y\}$. Agents $b$ and $c$ only have a singleton neighbourhood in these structures, we revert to the standard knowledge semantics here. Therefore true in $W$ are: $K_a 0_c$ (in $\{W, X\}$), $K_a 1_b$, and $K_a 2_b$. Observe that these three neighbourhoods do not intersect (there are not any two such that their intersection contains a facet), which, in Kripke semantics, is common fare for neighbourhoods. (In this sort of interpretation, it seems even impossible.) Although all this is crying for an intuitive interpretation, we will leave this as such in this report and instead relay it to further research.

### 4.3.4 Simplicial Structures for Byzantine Actions

It was agreed upon that before modelling Byzantine agents (that is, unreliable information or malicious intentions), we needed to have a fitting notion of belief. And indeed that occupied the main part of this report. It is clear that, as usual in action model logic, similar structures as presented above for the static notion of belief can be employed to encode dynamic notions involving deceit, error, and Byzantine conduct in general. Indeed, in simplicial semantics the "Kripke-style action models" have natural counterparts as *simplicial action models* [45, 24].

However, asymmetric versions of that needed for Byzantine communication have not yet been investigated. Additionally, novel update mechanisms such as communication patterns and related [5, 14] could also be contemplated to model Byzantine conduct. The update of simplicial models with communication patterns has been described in [14], including summary suggestions how to represent them as "simplicial communication patterns". All these horizons seem currently pretty far away, although mouth-watering in their presumed modelling features.

Let us close this far too short subsection on dynamics with a simplicial action model $\mathcal{C}^x$ representing that agent $c$ dies (that process $c$ crashes), although agents $a$ and $b$ only partially observe that, they remain uncertain about this. The preconditions of all three vertices are $\top$ (the always true proposition). As before, there are two $a$–$b$ edges, the lower edge is a maximal simplex (the complex is an impure simplicial set). If we update a simplicial model with this simplicial action model, any simplex is duplicated into one where $c$ is alive and a for $a$ and $b$ indistinguishable one where $c$ is dead. Even more elementary, simplicial action model $\mathcal{C}^{xi}$ consists of the $a$–$b$ edge part of $\mathcal{C}^x$ and models that $c$ is dying, pure and simple. And then, finally, to have at least some Byzantine in the picture, the action $\mathcal{C}^{xii}$ where $a$ incorrectly believes $c$ to be dying, to the despair of $b$ and, obviously, $c$ who both observe that.

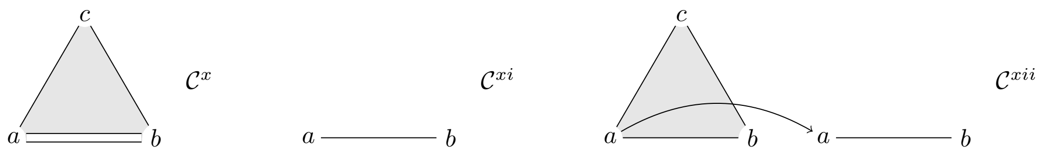


## References


**1** Kaya Alpturer, Joseph Y. Halpern, and Ron van der Meyden. Optimal eventual Byzantine agreement protocols with omission failures. In *PODC'23: Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing*, pages 244–252. Association for Computing Machinery, 2023. `doi:10.1145/3583668.3594573`.

**2** Guillaume Aucher and Thomas Bolander. Undecidability in epistemic planning. In Francesca Rossi, editor, *IJCAI-13, Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence: Beijing, China, 3–9 August 2013*, pages 27–33. AAAI Press, 2013. Available from: `https://www.ijcai.org/Abstract/13/015`.

**3** Guillaume Aucher and François Schwarzentruber. On the complexity of dynamic epistemic logic. In Burkhard C. Schipper, editor, *TARK 2013: Theoretical Aspects of Rationality and Knowledge, Proceedings of the 14th Conference – Chennai, India*, pages 19–28. University of California, Davis, 2013. `doi:10.48550/arXiv.1310.6406`.

**4** Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 11–58. Amsterdam University Press, 2008. Available from: `https://www.jstor.org/stable/j.ctt46mz4h.4`.

**5** Alexandru Baltag and Sonja Smets. Learning what others know. In Elvira Albert and Laura Kovács, editors, *LPAR23. LPAR-23: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, volume 73 of *EPiC Series in Computing*, pages 90–119. EasyChair, 2020. `doi:10.29007/plm4`.

**6** Vaishak Belle, Thomas Bolander, Andreas Herzig, and Bernhard Nebel. Epistemic planning: Perspectives on the special issue. *Artificial Intelligence*, 316:103842, March 2023. `doi:10.1016/j.artint.2022.103842`.

7 Ido Ben-Zvi and Yoram Moses. Beyond Lamport's *Happened-before*: On time bounds and the ordering of events in distributed systems. *Journal of the ACM*, 61(2):13:1–13:26, April 2014. doi:10.1145/2542181.

8 Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007. doi:10.3166/jancl.17.129-155.

9 Johan van Benthem, David Fernández-Duque, and Eric Pacuit. Evidence and plausibility in neighborhood structures. *Annals of Pure and Applied Logic*, 165(1):106–133, January 2014. doi:10.1016/j.apal.2013.07.007.

10 Johan van Benthem, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38(5):491–526, October 2009. doi:10.1007/s10992-008-9099-x.

11 Johan van Benthem and Sonja Smets. Dynamic logics of belief change. In Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Epistemic Logic*, pages 313–393. College Publications, 2015.

12 Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001. doi:10.1017/CBO9781107050884.

13 Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-classical Logics*, 21(1):9–34, 2011. doi:10.3166/jancl.21.9-34.

14 Armando Castañeda, Hans van Ditmarsch, David A. Rosenblueth, and Diego A. Velázquez. Communication pattern logic: Epistemic and topological views. *Journal of Philosophical Logic*, 2023. doi:10.1007/s10992-023-09713-8.

15 Armando Castañeda, Hans van Ditmarsch, David A. Rosenblueth, and Diego A. Velázquez. Comparing the update expressivity of communication patterns and action models. In Rineke Verbrugge, editor, *Proceedings Nineteenth conference on Theoretical Aspects of Rationality and Knowledge, Oxford, United Kingdom, 28–30th June 2023*, volume 379 of *Electronic Proceedings in Theoretical Computer Science*, pages 157–172. Open Publishing Association, 2023. doi:10.4204/EPTCS.379.14.

16 Armando Castañeda, Yannai A. Gonczarowski, and Yoram Moses. Unbeatable consensus. *Distributed Computing*, 35(2):123–143, April 2022. doi:10.1007/s00446-021-00417-3.

17 Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. Long-lived tasks. In Amr El Abbadi and Benoît Garbinato, editors, *Networked Systems: 5th International Conference, NETYS 2017, Marrakech, Morocco, May 17–19, 2017, Proceedings*, volume 10299 of *Lecture Notes in Computer Science*, pages 439–454. Springer, 2017. doi:10.1007/978-3-319-59647-1_32.

18 Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. Unifying concurrent objects and distributed tasks: Interval-linearizability. *Journal of the ACM*, 65(6):45:1–45:42, December 2018. doi:10.1145/3266457.

19 K. Mani Chandy and Jayadev Misra. How processes learn. *Distributed Computing*, 1(1):40–52, March 1986. doi:10.1007/BF01843569.

20 Francien Dechesne and Yanjing Wang. To know or not to know: epistemic approaches to security protocol verification. *Synthese*, 177(Supplement 1):51–76, 2010. doi:10.1007/s11229-010-9765-8.

21 Cédric Dégremont, Benedikt Löwe, and Andreas Witzel. The synchronicity of dynamic epistemic logic. In Krzysztof R. Apt, editor, *TARK XIII, Theoretical Aspects of Rationality and Knowledge: Proceedings of the Thirteenth Conference (TARK 2011)*, pages 145–152. Association for Computing Machinery, 2011. doi:10.1145/2000378.2000395.

22 Hans van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese*, 147(2):229–275, November 2005. doi:10.1007/s11229-005-1349-7.

**23** Hans van Ditmarsch, Jan van Eijck, Ignacio Hernández-Antón, Floor Sietsma, Sunil Simon, and Fernando Soler-Toscano. Modelling cryptographic keys in dynamic epistemic logic with DEMO. In Javier Bajo Pérez et al., editor, *Highlights on Practical Applications of Agents and Multi-Agent Systems: 10th International Conference on Practical Applications of Agents and Multi-Agent Systems*, volume 156 of *Advances in Intelligent and Soft Computing*, pages 155–162. Springer, 2012. doi:10.1007/978-3-642-28762-6_19.

**24** Hans van Ditmarsch, Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. Knowledge and simplicial complexes. In Björn Lundgren and Nancy Abigail Nuñez Hernández, editors, *Philosophy of Computing: Themes from IACAP 2019*, volume 143 of *Philosophical Studies Series*, pages 1–50. Springer, 2022. doi:10.1007/978-3-030-75267-5_1.

**25** Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi. An introduction to logics of knowledge and belief. In Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Epistemic Logic*, pages 1–51. College Publications, 2015.

**26** Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007. doi:10.1007/978-1-4020-5839-4.

**27** Hans van Ditmarsch, Wiebe van der Hoek, Ron van der Meyden, and Ji Ruan. Model checking Russian Cards. In *Proceedings of the Third Workshop on Model Checking and Artificial Intelligence (MoChArt 2005)*, volume 149(2) of *Electronic Notes in Theoretical Computer Science*, pages 105–123. Elsevier, 2006. doi:10.1016/j.entcs.2005.07.029.

**28** Cynthia Dwork and Yoram Moses. Knowledge and common knowledge in a Byzantine environment: Crash failures. *Information and Computation*, 88(2):156–186, October 1990. doi:10.1016/0890-5401(90)90014-9.

**29** Jan van Eijck. DEMO – A demo of epistemic modelling. In Johan van Benthem, Dov Gabbay, and Benedikt Löwe, editors, *Interactive Logic: Selected Papers from the 7th Augustus de Morgan Workshop, London*, volume 1 of *Texts in Logic and Games*, pages 303–362. Amsterdam University Press, 2007. Available from: https://www.jstor.org/stable/j.ctt45kdbf.15.

**30** Jan van Eijck and Malvin Gattinger. Elements of epistemic crypto logic (extended abstract). In *AAMAS'15: Proceedings of the 2015 International Conference on Autonomous Agents & Multiagent Systems*, pages 1795–1796. Association for Computing Machinery, 2015. Available from: https://www.ifaamas.org/Proceedings/aamas2015/aamas/p1795.pdf.

**31** Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995. doi:10.7551/mitpress/5803.001.0001.

**32** Krisztina Fruzsa, Roman Kuznets, and Ulrich Schmid. Fire! In Joseph Halpern and Andrés Perea, editors, *Proceedings Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, Beijing, China, June 25–27, 2021*, volume 335 of *Electronic Proceedings in Theoretical Computer Science*, pages 139–153. Open Publishing Association, 2021. doi:10.4204/EPTCS.335.13.

**33** Eli Gafni. Snapshot for time: The one-shot case. Eprint 1408.3432, arXiv, 2014. doi:10.48550/arXiv.1408.3432.

**34** Nina Gierasimczuk and Jakub Szymanik. A note on a generalization of the Muddy Children puzzle. In Krzysztof R. Apt, editor, *TARK XIII, Theoretical Aspects of Rationality and Knowledge: Proceedings of the Thirteenth Conference (TARK 2011)*, pages 257–264. Association for Computing Machinery, 2011. doi:10.1145/2000378.2000409.

**35** Guy Goren and Yoram Moses. Silence. *Journal of the ACM*, 67(1):3:1–3:26, February 2020. doi:10.1145/3377883.

**36** Éric Goubault, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. Semi-simplicial set models for distributed knowledge. In *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, 2023. doi:10.1109/LICS56636.2023.10175737.

37 Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task computability. *Information and Computation*, 278:104597, June 2021. doi:10.1016/j.ic.2020.104597.

38 Michell Guzmán, Sophia Knight, Santiago Quintero, Sergio Ramírez, Camilo Rueda, and Frank Valencia. Reasoning about distributed knowledge of groups with infinitely many agents. In Wan Fokkink and Rob van Glabbeek, editors, *30th International Conference on Concurrency Theory: CONCUR 2019, August 27–30, 2019, Amsterdam, the Netherlands*, volume 140 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 29:1–29:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.CONCUR.2019.29.

39 Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, July 1990. doi:10.1145/79147.79161.

40 Joseph Y. Halpern and Lenore D. Zuck. A little knowledge goes a long way: Knowledge-based derivations and correctness proofs for a family of protocols. *Journal of the ACM*, 39(3):449–478, July 1992. doi:10.1145/146637.146638.

41 Helle Hvid Hansen, Clemens Kupke, and Eric Pacuit. Neighbourhood structures: Bisimilarity and basic model theory. *Logical Methods in Computer Science*, 5(2):2:1–2:38, April 2009. doi:10.2168/LMCS-5(2:2)2009.

42 Maurice P. Herlihy, Dmitry Kozlov, and Sergio Rajsbaum. *Distributed Computing through Combinatorial Topology*. Morgan Kaufmann, 2014. doi:10.1016/C2011-0-07032-1.

43 Maurice P. Herlihy and Jeannette M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, July 1990. doi:10.1145/78969.78972.

44 Arjen Hommersom, John-Jules Meyer, and Erik de Vink. Update semantics of security protocols. *Synthese*, 142(2):229–267, September 2004. doi:10.1007/s11229-004-2247-0.

45 Jérémy Ledent. *Geometric semantics for asynchronous computability*. PhD thesis, Paris-Saclay University, Palaiseau, France, 2019. Prepared at École polytechnique. Available from: https://theses.hal.science/tel-02445180.

46 Qiang Liu and Yongmei Liu. Multi-agent epistemic planning with common knowledge. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, Stockholm, 13–19 July 2018*, pages 1912–1920. International Joint Conferences on Artificial Intelligence, 2018. doi:10.24963/ijcai.2018/264.

47 Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.

48 Bastien Maubert, Sophie Pinchinat, François Schwarzentruber, and Silvia Stranieri. Concurrent games in Dynamic Epistemic Logic. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, Yokohama, January 2021*, pages 1877–1883. International Joint Conferences on Artificial Intelligence, 2020. doi:10.24963/ijcai.2020/260.

49 Yoram Moses. Relating knowledge and coordinated action: The knowledge of preconditions principle. In R. Ramanujam, editor, *Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, Carnegie Mellon University, Pittsburgh, USA, June 4–6, 2015*, volume 215 of *Electronic Proceedings in Theoretical Computer Science*, pages 231–245. Open Publishing Association, 2016. doi:10.4204/EPTCS.215.17.

50 Yoram Moses and Mark R. Tuttle. Programming simultaneous actions using common knowledge. *Algorithmica*, 3(1–4):121–169, November 1988. doi:10.1007/BF01762112.

51 Eric Pacuit. *Neighborhood Semantics for Modal Logic*. Short Textbooks in Logic. Springer, 2017. doi:10.1007/978-3-319-67149-9.

52 Marc Pauly and Rohit Parikh. Game Logic – An overview. *Studia Logica*, 75(2):165–182, November 2003. doi:10.1023/A:1027354826364.

53 Vaughan Pratt. Modeling concurrency with geometry. In *Conference Record of the Eighteenth Annual ACM Symposium on Principles of Programming Languages*, pages 311–322. Association for Computing Machinery, 1991. doi:10.1145/99583.99625.

54 Krister Segerberg. Irrevocable belief revision in dynamic doxastic logic. *Notre Dame Journal of Formal Logic*, 39(3):287–306, summer 1998. doi:10.1305/ndjfl/1039182247.

55 Diego A. Velázquez, Armando Castañeda, and David A. Rosenblueth. Communication pattern models: An extension of action models for dynamic-network distributed systems. In Joseph Halpern and Andrés Perea, editors, *Proceedings Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, Beijing, China, June 25–27, 2021*, volume 335 of *Electronic Proceedings in Theoretical Computer Science*, pages 307–321. Open Publishing Association, 2021. doi:10.4204/EPTCS.335.29.

## Participants


- Alexandru Baltag
University of Amsterdam, NL
- Henning Basold
Leiden University, NL
- Armando Castañeda
National Autonomous University of Mexico, MX
- Hans van Ditmarsch
CNRS – Toulouse, FR &
Université de Toulouse, FR
- Faith Ellen
University of Toronto, CA
- Pierre Fraigniaud
CNRS – Paris, FR &
Université Paris Cité, FR
- Krisztina Fruzsa
TU Wien, AT
- Murdoch James Gabbay
Heriot-Watt University – Edinburgh, GB
- Guy Goren
Protocol Labs – Kibbutz Nahsholim, IL
- Joseph Y. Halpern
Cornell University – Ithaca, NY, US
- Roman Kniazev
École Polytechnique – Palaiseau, FR
- Sophia Knight
University of Minnesota – Duluth, MN, US
- Roman Kuznets
TU Wien, AT
- Jérémy Ledent
University of Strathclyde – Glasgow, GB
- David Lehnherr
Universität Bern, CH
- Shihao (Jason) Liu
University of Toronto, CA
- Yoram Moses
Technion – Haifa, IL
- Susumu Nishimura
Kyoto University, JP
- Thomas Nowak
ENS – Gif-sur-Yvette, FR
- Ami Paz
CNRS – Gif-sur-Yvette, FR
- Sergio Rajsbaum
National Autonomous University of Mexico, MX
- Rojo Randrianomentsoa
TU Wien, AT
- Hugo Rincón Galeana
TU Wien, AT
- David A. Rosenblueth
National Autonomous University of Mexico, MX
- Ulrich Schmid
TU Wien, AT
- Sonja Smets
University of Amsterdam, NL
- Thomas Studer
Universität Bern, CH
- Diego A. Velázquez
National Autonomous University of Mexico, MX


Report from Dagstuhl Seminar 23281

# Theoretical Advances and Emerging Applications in Abstract Interpretation

**Arie Gurfinkel*[1], Isabella Mastroeni*[2], Antoine Miné*[3], Peter Müller*[4], and Anna Becchi†[5]**

1 **University of Waterloo, CA.** arie.gurfinkel@uwaterloo.ca
2 **University of Verona, IT.** isabella.mastroeni@univr.it
3 **Sorbonne University – Paris, FR.** antoine.mine@lip6.fr
4 **ETH Zürich, CH.** peter.mueller@inf.ethz.ch
5 **Bruno Kessler Foundation – Trento, IT.** abecchi@fbk.eu

**Abstract**

This report documents the program and the outcomes of Dagstuhl Seminar 23281 "Theoretical Advances and Emerging Applications in Abstract Interpretation."

Abstract Interpretation (AI) is a theory of the approximation of program semantics. Since its introduction in the 70s, it lead to insights into theoretical research in semantics, a rich and robust mathematical framework to discuss about semantic approximation and program analysis, and the design of effective program analysis tools that are now routinely used in this industry. The seminar brought together academic and industrial partners to assess the state of the art in AI as well as discuss its future. It considered its foundational aspects, connections with other formal methods, emergent applications, user needs in program verification, tool design and evaluation, as well as educational aspects and community management. Its goal was to collect new ideas and new perspectives on all these aspects of AI in order to pave the way for new applications.

**Seminar** July 9–14, 2023 – https://www.dagstuhl.de/23281

**2012 ACM Subject Classification** Software and its engineering → Automated static analysis; Software and its engineering → Completeness; Software and its engineering → Correctness; Software and its engineering → Formal methods; Software and its engineering; Software and its engineering → Software functional properties; Software and its engineering → Software safety; Software and its engineering → Software verification and validation

**Keywords and phrases** abstract domains, abstract interpretation, program semantics, program verification, static program analysis

**Digital Object Identifier** 10.4230/DagRep.13.7.66

## 1 Executive Summary

*Antoine Miné (Sorbonne University – Paris, FR)*
*Arie Gurfinkel (University of Waterloo, CA)*
*Isabella Mastroeni (University of Verona, IT)*
*Peter Müller (ETH Zürich, CH)*

**License** Creative Commons BY 4.0 International license
© Antoine Miné, Arie Gurfinkel, Isabella Mastroeni, and Peter Müller

Abstract Interpretation (AI) is a theory of the approximation of possible program behaviors. Since its introduction in the late 70s, it has evolved into a very general theory to describe and compare formal semantics of programs and systems. As a more practical aspect, it provides a

* Editor / Organizer
† Editorial Assistant / Collector

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

formal framework to design effective static analyzers that are automatic, sound, and efficient. The last two decades have seen the emergence of practical tools now used routinely in the industry, starting with the embedded critical software industry and now being also applied to more mainstream software.

Despite its strength, designing new static analyses with AI remains a challenging task, involving both theoretical research and engineering efforts. The limited diffusion of AI knowledge in Universities, in Engineering Schools, and in the industry may hinder the development and more widespread use of AI-based tools. Moreover, the early focus on the verification of run-time errors in embedded software, while it provided a simpler and more controlled context than the general problem of verifying consumer software, and resulted in industrial successes, also could give the false impression that AI is only suitable for this task, while its theory is in fact much more general. Finally, the AI community has little interactions with that of other formal methods.

Based on these early observations, we set out to organize a Dagstuhl seminar on "Theoretical Advances and Emerging Applications in Abstract Interpretation" to discuss the state of the art in AI, identify its key challenges, and plan for its future. The seminar brought together 37 international experts in static analysis, from 10 countries, and from both academic and industrial background, covering a wide spectrum from pure Abstract Interpretation (AI) theory, to tool providers, to industrial users.

To provide a structure to the discussion, the seminar was organized as a series of topical days focusing on identified aspects (but not excluding other topics) of the state of the art and perceived challenges in AI: static safety verification, tools and applications, verification beyond safety analysis, and education. We proposed three invited talks and an invited tutorial of extended length (1h) and from key people in the community to bootstrap discussions on a variety of topics, ranging from theoretical AI to industrial tools, and providing historical perspectives. These were complemented with 23 shorter talks proposed spontaneously throughout the seminar. These talks, from 10min to 30min long, discussed a large variety of topics, including new research results, theoretical advances, experience reports, technical realizations, and reports on teaching activities. A significant number of presentations discussed connections of AI with other formal methods, including SMT solving, types, program logic.

During the seminar, we also organized several breakout sessions, where smaller working groups discussed a selection of topics: soundness requirements for static analysis tools, expressive domains, community infrastructure, education, connections with other formal methods, connections with machine learning, and tackling the verification of hyper-properties. As a conclusion of these working groups and the overall seminar, several tasks were started and a number of action items were proposed to advance further:

- The group on soundness requirements proposed a first list of requirements that an analyzer should fulfill. It identified the need to discuss these findings with members of the soundiness manifesto, which raised awareness on the lack of proper requirements.
- The group on community infrastructure proposed a series of practical actions to build the community on AI. It suggested the creation of a Special Interest Group in AI to coordinate the efforts.
- The group on teaching started an on-going list (to be completed) of educational materials on AI. It presented the need to develop materials targeting undergraduate students and practitioners, to make efforts to better share available teaching resources, and to provide introductory courses on AI on MOOC platforms.
- The group on the interaction between AI and other formal methods presented relevant connections with deductive verification, dynamic techniques, and model checking. In particular, it suggested the organization of a seminar on AI and deductive verification.

In general, this seminar expressed the interest to continue the discussion on the future of AI in further seminars focusing on specific challenges and opportunities as uncovered during the seminar, and adapting the list of participants in consequence.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Code Reuse Vulnerabilities in Modern Web Applications

*Musard Balliu (KTH Royal Institute of Technology – Stockholm, SE)*

**License** 🅭 Creative Commons BY 4.0 International license
© Musard Balliu
**Joint work of** Mikhail Shcherbakov, Musard Balliu, Cristian-Alexandru Staicu
**Main reference** Mikhail Shcherbakov, Musard Balliu, Cristian-Alexandru Staicu: "Silent Spring: Prototype Pollution Leads to Remote Code Execution in Node.js", in Proc. of the 32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023, pp. 5521–5538, USENIX Association, 2023.
**URL** https://www.usenix.org/conference/usenixsecurity23/presentation/shcherbakov

We study code reuse vulnerabilities in modern web application. Prototype pollution is a dangerous vulnerability affecting prototype-based languages like JavaScript and the Node.js platform. It refers to the ability of an attacker to inject properties into an object's root prototype at runtime and subsequently trigger the execution of legitimate code gadgets that access these properties on the object's prototype, leading to attacks such as Denial of Service (DoS), privilege escalation, and Remote Code Execution (RCE).

In this work, we set out to study the problem in a holistic way, from the detection of prototype pollution to detection of gadgets, with the goal of finding end-to-end exploits beyond DoS, in full-fledged Node.js applications. We build a multi-staged framework that uses multi-label static taint analysis to identify prototype pollution in Node.js libraries and applications, as well as a hybrid approach to detect universal gadgets, notably, by analyzing the Node.js source code. We implement our framework on top of GitHub's static analysis framework CodeQL to find 11 universal gadgets in core Node.js APIs, leading to code execution. Furthermore, we use our methodology in a study of 15 popular Node.js applications to identify prototype pollutions and gadgets. We manually exploit eight RCE vulnerabilities in three high-profile applications such as NPM CLI, Parse Server, and Rocket.Chat.

## 3.2 Cooperative Verification

*Dirk Beyer (LMU München, DE)*

**License** 🅭 Creative Commons BY 4.0 International license
© Dirk Beyer
**Joint work of** Dirk Beyer, Heike Wehrheim
**Main reference** Dirk Beyer, Heike Wehrheim: "Verification Artifacts in Cooperative Verification: Survey and Unifying Component Framework", in Proc. of the Leveraging Applications of Formal Methods, Verification and Validation: Verification Principles – 9th International Symposium on Leveraging Applications of Formal Methods, ISoLA 2020, Rhodes, Greece, October 20-30, 2020, Proceedings, Part I, Lecture Notes in Computer Science, Vol. 12476, pp. 143–167, Springer, 2020.
**URL** https://doi.org//10.1007/978-3-030-61362-4_8

Cooperative verification is an approach in which several verifiers help each other solving the verification problem by sharing artifacts about the verification process. There are many verification tools available, and they have different strengths. While the tools continuously increase their individual capabilities, the potential of cooperation is largely unused. The problem is that in order to use verifiers 'off-the-shelf', we need clear interfaces to invoke the tools and to pass information. Part of the interfacing problem is to define standard artifacts to be exchanged between verifiers. We explain a few recent approaches for cooperative combinations that are based on verification witnesses as exchange format, including witness validation, component-based CEGAR, and exchanging invariants between automatic and interactive verifiers. We also give a brief overview of CoVeriTeam, a tool for composing verification systems from existing off-the-shelf components.

**References**

1 Beyer, D., Dangl, M., Wendler, P.: A unifying view on SMT-based software verification. J. Autom. Reasoning 60(3), 299–335 (2018). https://doi.org/10.1007/s10817-017-9432-6
2 Beyer, D., Haltermann, J., Lemberger, T., Wehrheim, H.: Decomposing Software Verification into Off-the-Shelf Components: An Application to CEGAR. In: Proc. ICSE. pp. 536–548. ACM (2022). https://doi.org/10.1145/3510003.3510064
3 Beyer, D., Kanav, S.: CoVeriTeam: On-demand composition of cooperative verification systems. In: Proc. TACAS. pp. 561–579. LNCS 13243, Springer (2022). https://doi.org/10.1007/978-3-030-99524-9_31
4 Beyer, D., Keremoglu, M.E.: CPAchecker: A tool for configurable software verification. In: Proc. CAV. pp. 184–190. LNCS 6806, Springer (2011). https://doi.org/10.1007/978-3-642-22110-1_16
5 Beyer, D., Spiessl, M., Umbricht, S.: Cooperation between automatic and interactive software verifiers. In: Proc. SEFM. p. 111–128. LNCS 13550, Springer (2022). https://doi.org/10.1007/978-3-031-17108-6_7
6 Beyer, D., Wehrheim, H.: Verification artifacts in cooperative verification: Survey and unifying component framework. In: Proc. ISoLA (1). pp. 143–167. LNCS 12476, Springer (2020). https://doi.org/10.1007/978-3-030-61362-4_8

## 3.3 Interactive Abstract Interpretation

*Bor-Yuh Evan Chang (University of Colorado – Boulder, US)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Bor-Yuh Evan Chang
**Joint work of** Benno Stein, Bor-Yuh Evan Chang, Manu Sridharan, David Flores
**Main reference** Benno Stein, Bor-Yuh Evan Chang, Manu Sridharan: "Demanded abstract interpretation", in Proc. of the PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021, pp. 282–295, ACM, 2021.
**URL** https://doi.org//10.1145/3453483.3454044

We consider the problem of making expressive static analyzers interactive. Formal static analysis is seeing increasingly widespread adoption as a tool for verification and bug-finding, but even with powerful cloud infrastructure it can take minutes or hours to get batch analysis results after a code change. While existing techniques offer some demand-driven or incremental aspects for certain classes of analysis, the fundamental challenge we tackle is doing both for arbitrary abstract interpreters. Our technique, demanded abstract interpretation, lifts program syntax and analysis state to a dynamically evolving graph structure, in which program edits, client-issued queries, and evaluation of abstract semantics are all treated uniformly.

## 3.4 Formal Verification of Avionics Software

*David Delmas (Airbus S.A.S. – Toulouse, FR)*

**License** Creative Commons BY 4.0 International license
© David Delmas
**Joint work of** David Delmas, Antoine Miné, Abdelraouf Ouadjaout, Famantanantsoa Randimbivololona, Abderrahmane Brahmi
**Main reference** David Delmas, Abdelraouf Ouadjaout, Antoine Miné: "Static Analysis of Endian Portability by Abstract Interpretation", in Proc. of the Static Analysis – 28th International Symposium, SAS 2021, Chicago, IL, USA, October 17-19, 2021, Proceedings, Lecture Notes in Computer Science, Vol. 12913, pp. 102–123, Springer, 2021.
**URL** https://doi.org//10.1007/978-3-030-88806-0_5
**Main reference** Abdellatif Atki, Abderrahmane Brahmi, David Delmas, Mohamed Habib Essoussi, Famantanantsoa Randimbivololona, Thomas Marie: "Formalise to automate: deployment of a safe and cost-efficient process for avionics software".In ERTSS: Proc. of the 9th European Congress on Embedded Real Time Software and Systems, Jan 2018, Toulouse, France.
**URL** https://www.di.ens.fr/ delmas/erts18/
**Main reference** David Delmas: "Static analysis of program portability by abstract interpretation", 2022.
**URL** https://theses.hal.science/tel-04028096

The size and complexity of avionics software have grown exponentially from one aircraft generation to the next in the past 4 decades. Traditional software development processes leveraging informal verification techniques fail to scale within reasonable costs. In particular, verification is liable for a steadily growing share of the overall development costs. The 2015 current status was about 70.

To address this issue, Airbus have been transforming internal development processes since 2016. Internal domain-specific languages have been developed to enable the formalization of design artifacts, and automate part of verification activities. Automation is enabled by the interoperation of tools relying on sound formal techniques. For instance, Frama-C/WP and SMT-solvers are used to automate unit verification with deductive methods. Most so-called Unit Proofs are automatic, assuming high-level memory and numerical models, as well as some preconditions. Such assumptions are verified by other tools, such as the Astrée static analyzer, which leverages Abstract Interpretation to prove the absence of run-time errors and check assumed non-aliasing properties. We rely on the CompCert formally verified compiler to enable that most formal verification activities may be conducted on source code.

Beyond safety properties and currently established processes, we also develop internally static analyses by Abstract Interpretation to automate regression verification and portability verification. In particular, our portability analysis is able to prove without false alarms the portability of low-level C avionics software up to 1 million lines of C across platforms with opposite byte-orders (endianness).

## 3.5 Calculating Equational Laws over ADTs

*Gidon Ernst (LMU München, DE)*

**License** Creative Commons BY 4.0 International license
© Gidon Ernst
**Joint work of** Gidon Ernst, Grigory Fedyukovich, Robin Sögtrop

We motivate the use of program transformation, traditionally used for the optimization of functional programs, as basic building blocks for lemma synthesis for recursive functions over algebraic data types. The key idea is to calculate at the function level instead of at the formula level, which allows one to work with intermediate syntactic functions. An important step therefore is to complement existing techniques (fixpoint fusion, deaccumulation) with capabilities to recognize when two functions are identical resp. share commonalities wrt. preconditions. The approach, work in progress, discovers typical distributive laws of list/tree functions like length, append, remove, which are similar in shape to those lemmas often required for data refinement proofs.

## 3.6 Teaching Abstract Interpretation with LiSA

*Pietro Ferrara (University of Venice, IT)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Pietro Ferrara
**Joint work of** Pietro Ferrara, Luca Negrini, Vincenzo Arceri

LiSA (LIbrary for Static Analysis – `https://github.com/lisa-analyzer/lisa`) is a Java library that implements the most common components of abstract interpretation-based static analyses. In this talk, we report our experience when adopting LiSA during courses at the master level focused on the theory of abstract interpretation. LiSA allowed students to implement, execute, and practice the theoretical concepts formalized throughout the course. However, it turned out that proper implementation of non-trivial abstract domains is beyond the capabilities of standard CS master students.

## 3.7 Alpha from Below over Quantified First-order Formulas

*Eden Frenkel (Tel Aviv University, IL)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Eden Frenkel
**Joint work of** Eden Frenkel, Sharon Shoham Buchbinder, Oded Padon, Tej Chajed

Often, verification of infinite-state and other complex systems, such as the Paxos consensus protocol, must work for unbounded domains. Quantified first-order logic allows us to reason over these unbounded domains, describe the desired behavior of these systems, and specify correctness properties. This is done by constructing formulas that describe the initial states of the system and its possible transitions, as well as formalizing the notion of illegal or unwanted states. A proof for the safety of a system then becomes an unreachability proof for the bad states, which can be provided via another formula, a safety invariant, which holds on the initial states, is invariant with respect to the transitions, and doesn't hold on the bad states.

In this work we tackle invariant inference though the framework of abstract interpretation. We propose an algorithm that computes the strongest over-approximation of reachable states by iteratively sampling counter-examples to induction, but which is highly infeasible due to the magnitude of the abstract domain and SMT solver limitations. We proceed by presenting techniques based on syntactic subsumption that manage to avoid redundant formulas and explore an identically expressive search space exponentially more efficiently, and generalize to the abstract domain of first-order formulas.

### 3.8 Fast Approximations of Quantifier Elimination

*Isabel Garcia-Contreras (University of Waterloo, CA), Arie Gurfinkel (University of Waterloo, CA), Hari Govind V K, and Sharon Shoham Buchbinder (Tel Aviv University, IL)*

**License** Creative Commons BY 4.0 International license
© Isabel Garcia-Contreras, Arie Gurfinkel, Hari Govind V K, and Sharon Shoham Buchbinder
**Main reference** Isabel Garcia-Contreras, Hari Govind V. K., Sharon Shoham, Arie Gurfinkel: "Fast Approximations of Quantifier Elimination", CoRR, Vol. abs/2306.10009, 2023.
**URL** https://doi.org//10.48550/ARXIV.2306.10009

Quantifier elimination (qelim) is used in many automated reasoning tasks including program synthesis, exist-forall solving, quantified SMT, Model Checking, and solving Constrained Horn Clauses (CHCs). Exact qelim is computationally expensive. Hence, it is often approximated. For example, Z3 uses "light" pre-processing to reduce the number of quantified variables. CHC-solver Spacer uses model-based projection (MBP) to under-approximate qelim relative to a given model, and over-approximations of qelim can be used as abstractions. In this talk, we present the QEL framework for fast approximations of qelim. QEL provides a uniform interface for both quantifier reduction and model-based projection. QEL builds on the egraph data structure – the core of the EUF decision procedure in SMT – by casting quantifier reduction as a problem of choosing ground (i.e., variable-free) representatives for equivalence classes. We have used QEL to implement MBP for the theories of Arrays and Algebraic Data Types (ADTs). We integrated QEL and our new MBP in Z3 and evaluated it within several tasks that rely on quantifier approximations, outperforming state-of-the-art.

### 3.9 Uniform Interpolation for Efficient Domain Reduction

*Isabel Garcia-Contreras (University of Waterloo, CA), Arie Gurfinkel (University of Waterloo, CA), and Jorge Navas (Certora – Seattle, US)*

**License** Creative Commons BY 4.0 International license
© Isabel Garcia-Contreras, Arie Gurfinkel, and Jorge Navas

Handling precise abstract values over large sets of program variables is costly. A solution is to split such "monolithic" values (over the whole set of program variables) into several subvalues over smaller subsets of variables. Applying abstract transformers separately to each subvalue is then more efficient than to the monolithic one, but typically incurs precision loss. To address this, one can transfer information between subvalues, a process known as reduction. Reduction is done iteratively by refining a subvalue using information from other subvalues until no values are further refined. Information transfer can be stopped at any time, guaranteeing soundness. Convergence is not ensured. In this talk, we study termination and precision properties of reduction from the perspective of logic. We define the notion of refutational equivalence between a monolithic and split value as a practical way to understand when a split domain is "precise enough". Our main result is that if the theories used in the abstract domain admit and are closed uniform interpolation, reduction can be done in one step – by computing the uniform interpolant of the formula – guaranteeing refutational equivalence. For theories (or their combination) that do not admit uniform interpolation, or are not closed under uniform interpolation, we show that if a uniform interpolant is found, the reduction procedure can be immediately stopped, guaranteeing the best precision for the partitioned domain.

## 3.10 An incomplete journey in Completeness

*Roberto Giacobazzi (University of Verona, IT)*

**License** Creative Commons BY 4.0 International license
© Roberto Giacobazzi

In this talk I will present some results concerning the bridge between the theory of computation and abstract interpretation. The two theories differ profoundly in the way programs are interpreted: In the standard theory of computation program equivalence is extensional, we have compositionally and referential transparency. Abstract interpretation instead is deeply intensional, non compositional (in the sense that by composition we may lose precision). However, it is possible to view the theory of abstract interpretation from the perspective of the theory of computation. In this case a different notion of program equivalence is considered – the one defined by the equivalence of the abstract interpretations, and push the theory of abstract interpretation to its limits by studying the properties of abstraction that make an abstract interpretation closer to the standard notion of interpretation as originally defined by A. Turing. Interestingly, a analogous of Rice Theorem can be stated for abstract interpretation and a number of results can be obtained for classes of programs for which an abstract interpreter is precise (complete).

## 3.11 Abstract interpretation based under approximations and Sufficient Incorrectness Logic

*Roberta Gori (University of Pisa, IT)*

**License** Creative Commons BY 4.0 International license
© Roberta Gori

**Joint work of** Flavio Ascari, Roberto Bruni, Roberta Gori, Francesco Logozzo
**Main reference** Flavio Ascari, Roberto Bruni, Roberta Gori: "Limits and difficulties in the design of under-approximation abstract domains", in Proc. of the Foundations of Software Science and Computation Structures – 25th International Conference, FOSSACS 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings, Lecture Notes in Computer Science, Vol. 13242, pp. 21–39, Springer, 2022.
**URL** https://doi.org//10.1007/978-3-030-99253-8_2

To address bug finding rather than correctness, Incorrectness Logic has been recently proposed by O'Hearn: it is based on under-approximations, thus it only reports *true* alarms. In principle, Abstract Interpretation techniques can handle under-approximations as well as over-approximations, but, in practice, few attempts were developed for the former, notwithstanding the wide literature on the latter. We aim to answer the following open question raised by O'Hearn: which role can Abstract Interpretation play in the development of under-approximate tools for bug catching? Our findings clarify, for the first time, why over- and under-approximation analysers exhibited such a different development and outline the limits of under-approximation Abstract Interpretation based analyses. Our key argument is the practical difficulty to design an effective under-approximation abstract domain able to deal with common program statements. For this reasons we investigate logics for underapproximations. We introduce Sufficient Incorrectness Logic (SIL), a new under-approximating, triple-based program logic to reason about program errors. SIL is designed to set apart the initial states leading to errors. We formally compare SIL to existing triple-based program logics.

## 3.12 On the fly verification with (incremental) interactive abstract interpretation

*Manuel Hermenegildo (IMDEA Software Institute – Pozuelo de Alarcón, ES & UPM – Madrid, ES)*

**License** Creative Commons BY 4.0 International license
© Manuel Hermenegildo
**Joint work of** Manuel Hermenegildo, Isabel Garcia-Contreras, José F. Morales, Pedro López-García, Louis Rustenholz, Daniela Ferreiro, Daniel Jurjo
**Main reference** Miguel A. Sanchez-Ordaz, Isabel Garcia-Contreras, Victor Perez-Carrasco, José F. Morales, Pedro López-García, Manuel V. Hermenegildo: "VeriFly: On-the-fly Assertion Checking via Incrementality", CoRR, Vol. abs/2106.07045, 2021.
**URL** https://arxiv.org/abs/2106.07045

We demonstrated how the integration of the Ciao abstract interpretation framework within different IDEs takes advantage of our efficient and incremental fixpoint to achieve effective verification on-the-fly, as the program is developed. We also demonstrated an embedding of this framework within the browser, and how it can be used to build interactive tutorials for teaching abstract interpretation.

### Further reading

The work and demo presented builds on several specific components of the Ciao abstract interpretation framework:

- **The "top-down" algorithm:** (a.k.a. the PLAI algorithm) is the fundamental component of our approach: see [4] and [5]. The latter is a step by step tutorial on how the algorithm was derived, the reasons for the different optimizations, etc. and is probably the best single reference for the original "top-down algorithm." This first algorithm already achieves precision and efficiency through the use of memo tables; inferring call-answer pairs (aka summaries, precondition-postcondition pairs, etc.) which can be several per each procedure / block / program point (a.k.a. multivariance, path/context/call-site sensitivity, cloning, etc.); abstraction of the paths and recursions/loops in the program as graphs / regular trees; detection of SCCs; dependency tracking between memo table entries for accelerating the fixpoint (keeping track of what has to be recomputed when something changes, which is later instrumental for incremental analysis); handles procedures (projection, extension), including (mutual) recursion; etc. It is also generic in the sense that it takes care on one hand of abstracting the control: paths, dynamic CFG, procedure call and return (projection, extension), including (mutual) recursion, etc. Then, multiple 'abstract domains' are available as plug-ins to abstract the data: recursive heap data structures, pointer aliasing, numerical domains, etc.
- **Incremental Analysis:** After developing the PLAI algorithm we developed and benchmarked the incremental version. It was first described in [2], which was later published in longer form as [3], and an improved incremental algorithm was later described in [6]. Finally, we recently extended the '96 algorithm to deal at the same time with modules (coarse-grained incrementality) and fine-grained incrementality within each module [1].
- **Interactive verification:** The final component is the integration of the analyer/verifier-/optimizer CiaoPP in the IDE. The latest version of this integration ("Verifly" [8, 7]), put in practice the idea of interactive verification with abstract interpretation.

**References**

1 I. Garcia-Contreras, J. F. Morales, and M. V. Hermenegildo. Incremental and Modular Context-sensitive Analysis. *Theory and Practice of Logic Programming*, 21(2):196–243, January 2021. https://arxiv.org/abs/1804.01839

2 M. V. Hermenegildo, G. Puebla, K. Marriott, and P. Stuckey. Incremental Analysis of Logic Programs. In *International Conference on Logic Programming*, pages 797–811. MIT Press, June 1995. https://cliplab.org/papers/incanal-iclp95_bitmap.pdf

3 M. V. Hermenegildo, G. Puebla, K. Marriott, and P. Stuckey. Incremental Analysis of Constraint Logic Programs. *ACM Transactions on Programming Languages and Systems*, 22(2):187–223, March 2000. https://cliplab.org/papers/incanal-toplas.pdf

4 K. Muthukumar and M. Hermenegildo. Determination of Variable Dependence Information at Compile-Time Through Abstract Interpretation. In *1989 North American Conference on Logic Programming*, pages 166–189. MIT Press, October 1989. https://cliplab.org/papers/abs-int-naclp89.pdf

5 K. Muthukumar and M. Hermenegildo. Deriving A Fixpoint Computation Algorithm for Top-down Abstract Interpretation of Logic Programs. Technical Report ACT-DC-153-90, Microelectronics and Computer Technology Corporation (MCC), Austin, TX 78759, April 1990. http://cliplab.org/papers/mcctr-fixpt.pdf

6 G. Puebla and M. V. Hermenegildo. Optimized Algorithms for the Incremental Analysis of Logic Programs. In *International Static Analysis Symposium (SAS 1996)*, number 1145 in Lecture Notes in Computer Science, pages 270–284. Springer-Verlag, September 1996. https://cliplab.org/papers/inc-fixp-sas_bitmap.pdf

7 M. A. Sanchez-Ordaz, I. Garcia-Contreras, V. Perez-Carrasco, J. F. Morales, P. Lopez-Garcia, and M.V. Hermenegildo. VeriFly: On-the-fly Assertion Checking with CiaoPP. In *6th Workshop on Formal Integrated Development Environment (F-IDE 2021, part of NASA NFM'21)*, Electronic Proceedings in Theoretical Computer Science (EPTCS), pages 1–5. Open Publishing Association (OPA), May 2021. Co-located with ETAPS 2021. https://cister-labs.pt/f-ide2021/images/preprints/F-IDE_2021_paper_7.pdf

8 M.A. Sanchez-Ordaz, I. Garcia-Contreras, V. Perez-Carrasco, J. F. Morales, P. Lopez-Garcia, and M. V. Hermenegildo. Verifly: On-the-fly Assertion Checking via Incrementality. *Theory and Practice of Logic Programming*, 21(6):768–784, September 2021. Special Issue on ICLP'21. http://arxiv.org/abs/2106.07045

## 3.13 Automated Reasoning for Privacy

*Temesghen Kahsai (Amazon Lab 126, US)*

**License** Creative Commons BY 4.0 International license
© Temesghen Kahsai

Automated reasoning techniques offer an exceptional avenue to attain the utmost assurance in safeguarding data privacy. In this presentation, we will explore our adaptations of these techniques in addressing vital inquiries concerning code and cloud infrastructure, aimed at identifying potentially harmful misconfigurations. We will delve into the deployment of highly scalable static analysis for sensitive data tracking, the utilization of a provably correct differential privacy library to guarantee the safety of shared aggregated data, and the implementation of diverse pseudo-anonymization methods to safeguard sensitive information.

## 3.14 Abstract Interpretation in Industry – Practical Experience with Astrée

*Daniel Kästner (AbsInt – Saarbrücken, DE)*

**License** Creative Commons BY 4.0 International license
© Daniel Kästner

The presentation gives a brief overview of the verification goals addressed by the Astrée analyzer, and then focuses on the development history from the ready-for-market academic version from 2009 till today's state. Enhancements are grouped into different categories, usability, compliance to formal requirements, new capabilities, domain-specific extensions, precision improvement and optimizations for scalability. The current state is briefly summarized by experimental results on automotive integration analysis projects, taken from our 2023 SAE conference paper. We then briefly summarize "selling points" of verification tools to industry users with an emphasis on the role of safety norms: they define the minimum state of the art for system development and play a fundamental role in the adoption of tools in industrial development processes. We give examples how abstract interpretation is addressed in DO-178C and ISO 26262, observe that it is entirely missing in others, and conclude that it is under-representated in today's safety norms. The presentation ends with listing some verification challenges we see upcoming due to current market trends.

### References

1 D. Kästner, C. Mallon, L. Mauborgne, S. Schank, S. Wilhelm, C. Ferdinand. *Automatic Sound Static Analysis for Integration Verification of AUTOSAR Software.* SAE Technical Paper 2023-01-0591, SAE World Congress 2023, Detroit, April 2023. DOI: `https://doi.org/10.4271/2023-01-0591`

## 3.15 SSA Translation Is an Abstract Interpretation, and its Application to Machine Code Analysis

*Matthieu Lemerre (CEA LIST – Gif-sur-Yvette, FR)*

**License** Creative Commons BY 4.0 International license
© Matthieu Lemerre
**Joint work of** Matthieu Lemerre, Olivier Nicole, Xavier Rival, Sébastien Bardin
**Main reference** Matthieu Lemerre: "SSA Translation Is an Abstract Interpretation", Proc. ACM Program. Lang., Vol. 7(POPL), pp. 1895–1924, 2023.
**URL** https://doi.org//10.1145/3571258

Conversion to Static Single Assignment (SSA) form is usually viewed as a syntactic transformation algorithm that gives unique names to program variables, and reconciles these names using "phi" functions based on a notion of domination. We instead propose a semantic approach, where SSA translation is performed using a simple dataflow analysis. Based on a new technique to use cyclic terms in abstract domains, we propose a Symbolic Expression abstract domain that performs a Global Value Numbering analysis, upon which we build our SSA translation. This implies a shift in perspective, as global value numbering becomes a prerequisite of SSA translation, instead of depending on SSA.

One application to performing SSA Translation by Abstract Interpretation is that SSA optimizations passes can be implemented as a combination of abstract domains, allowing to perform several optimizations simultaneously to solve the usual phase ordering problem and avoiding tedious maintenance of SSA invariants.

Our main motivation for this research is an analyser for machine code which uses SSA as its main intermediate representation. Machine code is too low-level to allow SSA translation without a prior semantic analysis, while SSA is an intermediate representation that makes static analysis easier than direct analysis of machine code. Viewing SSA translation as a semantic analysis solves this chicken-and-egg problem, allowing to simultaneously decompile machine code to SSA and use the SSA representation to perform the other semantic analyses (value analysis, memory analysis, and control-flow analysis). We illustrate the use of such an analysis on an embedded OS kernel where we prove security properties directly from its executable.

**References**

1 *Matthieu Lemerre*: SSA Translation Is an Abstract Interpretation. In Principle of Programming Languages (POPL), 2023.

2 *O. Nicole, M. Lemerre, S. Bardin, X. Rival*: No Crash, No Exploit : Automated Verification of Embedded Kernels . In Real-time systems and applications (RTAS), 2021

## 3.16 A Multilanguage Static Analysis of Python/C Programs with Mopsa

*Raphaël Monat (INRIA Lille, FR) and Antoine Miné (Sorbonne University – Paris, FR)*

**License** Creative Commons BY 4.0 International license
© Raphaël Monat and Antoine Miné
**Joint work of** Ouadjaout, Abdelraouf; Miné, Antoine
**Main reference** Raphaël Monat, Abdelraouf Ouadjaout, Antoine Miné: A Multilanguage Static Analysis of Python Programs with Native C Extensions. SAS 2021: 323-345
**URL** https://link.springer.com/chapter/10.1007/978-3-030-88806-0_16

Mopsa is a conservative static analysis platform, independent of language and abstraction choices. Developers are free to add arbitrary abstractions (numeric, pointer, memory, etc.) and syntax iterators for new languages. Mopsa encourages the development of independent abstractions which can cooperate or be combined to improve precision. In this talk, we will show how Mopsa analyses Python programs calling C libraries. It analyses directly and fully automatically both the Python and the C source codes. It reports runtime errors that may happen in Python, in C, and at the interface. We implemented our analysis in a modular fashion: it reuses off-the-shelf C and Python analyses written in the same analyzer. Our analyzer can tackle tests of real-world libraries a few thousand lines of C and Python long.

This talk is based on our SAS'21 paper.

## 3.17 Crab: A library for building abstract-interpretation-based analyses

*Jorge Navas (Certora – Seattle, US)*

**License** Creative Commons BY 4.0 International license
© Jorge Navas

Crab is an open-source library that helps to develop static analyses based on the theory of Abstract Interpretation. It provides a variety of software components such as inter-procedural forward and backward analyses, fixpoint solvers and a rich set of abstract domains, including interfaces with numerical abstract domain libraries such as Apron, Elina or PPLite.

This tutorial focuses on demonstrating how to use Crab from the perspective of two different kinds of users: (1) the ones who want to quick prototype new abstractions by combining existing abstract domains, and (2) those who want to develop new static analyses using Crab as a black-box. We show several interactive demos on how to derive new abstractions based on the combination of an array domain with relational numerical domains to prove memory safety of eBPF programs and how to implement a new LLVM-based static analyzer to infer upper-bounds for dynamic memory allocation in less than 100 lines of C++ code.

## 3.18 Calculational Design of Program Logics by Abstract Interpretation

*Patrick Cousot*

**License** 🅭 Creative Commons BY 4.0 International license
© Patrick Cousot
**Main reference** Submitted

We study transformational program logics for correctness and incorrectness that we extend to explicitly handle both termination and nontermination. We show that the logics are abstract interpretations of the right image transformer for a natural relational semantics covering both finite and infinite executions. This understanding of logics as abstractions of a semantics facilitates their comparisons through their respective abstractions of the semantics (rather that the much more difficult comparison through their formal deductive systems). More importantly, the formalization provides a calculational method for constructively designing the sound and complete formal deductive system by abstraction of the semantics. As an example, we extend Hoare logic to cover all possible behaviors of nondeterministic programs and design a new precondition (in)correctness logics. This logic can be used to prove that false alarms in static analysis are due to the over approximation of nonterminating behaviors by terminating over approximation in the analysis which cannot be done by incorrectness or outcome logic is unreachable in the concrete (although it is reachable in the abstract).

## 3.19 Mentorship for Formal Methods

*Ruzica Piskac (Yale University – New Haven, US)*

**License** 🅭 Creative Commons BY 4.0 International license
© Ruzica Piskac
**Joint work of** Mark Santolucito, Ruzica Piskac
**Main reference** Mark Santolucito, Ruzica Piskac: "Formal Methods and Computing Identity-based Mentorship for Early Stage Researchers", in Proc. of the 51st ACM Technical Symposium on Computer Science Education, SIGCSE 2020, Portland, OR, USA, March 11-14, 2020, pp. 135–141, ACM, 2020.
**URL** https://doi.org//10.1145/3328778.3366957

The field of formal methods relies on a large body of background knowledge that can dissuade researchers from engaging with younger students, such as undergraduates or high school students. However, we have found that formal methods can be an excellent entry point to computer science research – especially in the framing of Computing Identity-based Mentorship. In this talk, we report on our experience in using a cascading mentorship model to involve early stage researchers in formal methods, covering the process with these students from recruitment to publication. We present case studies and how we were able to integrate formal methods research with the students' own interests. We outline some key strategies that have led to success and reflect on strategies that have been, in our experience, inefficient.

## 3.20 Abstract Interpretation-based Program (Analysis) Logics

*Francesco Ranzato (University of Padova, IT)*

**License** Creative Commons BY 4.0 International license
© Francesco Ranzato
**Joint work of** Roberto Bruni, Roberto Giacobazzi, Roberta Gori, Francesco Ranzato
**Main reference** Roberto Bruni, Roberto Giacobazzi, Roberta Gori, Francesco Ranzato: "A Correctness and Incorrectness Program Logic", J. ACM, Vol. 70(2), pp. 15:1–15:45, 2023.
**URL** https://doi.org//10.1145/3582267

We introduce a program logic, called Local Completeness Logic for an abstract domain $A$ ($\mathrm{LCL}_A$), for proving both the correctness and incorrectness of program specifications. This proof system, which is parametrized by the abstraction $A$, combines over- and under-approximating reasoning: in a provable triple $\vdash_A [p]c[q]$ for the program $c$, $q$ is an under-approximation of the strongest post-condition of $c$ on input pre-condition $p$, such that their abstractions in $A$ coincide. If $A$ is the straightforward abstraction making all program properties equivalent, then the logic $\mathrm{LCL}_A$ coincides with O'Hearn's incorrectness logic. We discuss the pitfalls of this logic $\mathrm{LCL}_A$ and why it is hard to make proofs within it. We, therefore, advocate designing a weakening of $\mathrm{LCL}_A$ which should be able to prove properties of program analyses rather than program behaviors, in particular for proving that a program analysis is the best possible one in the underlying domain $A$.

## 3.21 VeriCode: Correct Translation of Abstract Specifications to C-Code

*Gerhard Schellhorn (Universität Augsburg, DE)*

**License** Creative Commons BY 4.0 International license
© Gerhard Schellhorn

The talk presented the new VeriCode project (funded by German Research Foundation DFG) we just started. The projects aims at the generation of efficient code from abstract specification. Such specifications are typical when using higher-order logic and abstract programs to specify functions. They typically use abstract datatypes which are easily translated to functional code. However the resulting code is typically not very efficient and requires garbage collection. The project was motivated by an earlier DFG project called Flashix that produced a verified file system for flash memory. A simple code generator is already implemented that produces C- and Scala-Code, based on the principle that in contrast to a functional implementation data structure should not share. The approach enables destructive updates and allows to avoid garbage collection. The approach still causes too much copying and the talk showed some first optimizations. Compared to a native implementation of a flash filesystem (UBIFS) in C that we used as a blueprint, our code is still some factors slower and the project aims to close the gap. The talk also showed the overall approach of the project, which is not just to implement an efficient code generator, but to verify that it is correct again using specifications of the relevant functionality and semantics. Ultimately this should allow to bootstrap the code generator by again generating code from the specification of its functionality.

## 3.22 Data Race Repair using Static Analysis Summaries

*Ilya Sergey (National University of Singapore, SG)*

**License** Creative Commons BY 4.0 International license
© Ilya Sergey
**Joint work of** Andreea Costea, Abhishek Tiwari, Sigmund Chianasta, Kishore R, Abhik Roychoudhury, Ilya Sergey
**Main reference** Andreea Costea, Abhishek Tiwari, Sigmund Chianasta, Kishore R, Abhik Roychoudhury, Ilya Sergey: "Hippodrome: Data Race Repair Using Static Analysis Summaries", ACM Trans. Softw. Eng. Methodol., Vol. 32(2), pp. 41:1–41:33, 2023.
**URL** https://doi.org//10.1145/3546942

Implementing bug-free concurrent programs is a challenging task in modern software development. State-of-the-art static analyses find hundreds of concurrency bugs in production code, scaling to large codebases. Yet, fixing these bugs in constantly changing codebases represents a daunting effort for programmers, particularly because a fix in the concurrent code can introduce other bugs in a subtle way.

In this talk, I will show how to harness compositional static analysis for concurrency bug detection, to enable a new Automated Program Repair (APR) technique for data races in large concurrent Java codebases. The key innovation of our work is an algorithm that translates procedure summaries inferred by the analysis tool for the purpose of bug reporting into small local patches that fix concurrency bugs (without introducing new ones). This synergy makes it possible to extend the virtues of compositional static concurrency analysis to APR, making our approach effective (it can detect and fix many more bugs than existing tools for data race repair), scalable (it takes seconds to analyse and suggest fixes for sizeable codebases), and usable (generally, it does not require annotations from the users and can perform continuous automated repair). Our study conducted on popular open-source projects has confirmed that our tool automatically produces concurrency fixes similar to those proposed by the developers in the past.

## 3.23 Property-Directed Reachability as Abstract Interpretation in the Monotone Theory

*Sharon Shoham Buchbinder (Tel Aviv University, IL)*

**License** Creative Commons BY 4.0 International license
© Sharon Shoham Buchbinder
**Joint work of** Yotam M. Y. Feldman, Mooly Sagiv, Sharon Shoham Buchbinder, Mooly Sagiv
**Main reference** Yotam M. Y. Feldman, Mooly Sagiv, Sharon Shoham, James R. Wilcox: "Property-directed reachability as abstract interpretation in the monotone theory", Proc. ACM Program. Lang., Vol. 6(POPL), pp. 1–31, 2022.
**URL** https://doi.org//10.1145/3498676
**Main reference** Yotam M. Y. Feldman, Sharon Shoham: "Invariant Inference with Provable Complexity from the Monotone Theory", in Proc. of the Static Analysis – 29th International Symposium, SAS 2022, Auckland, New Zealand, December 5-7, 2022, Proceedings, Lecture Notes in Computer Science, Vol. 13790, pp. 201–226, Springer, 2022.
**URL** https://doi.org//10.1007/978-3-031-22308-2_10

Inferring inductive invariants is one of the main challenges of formal verification. One of the latest breakthroughs in invariant inference is property-directed reachability (IC3/PDR). In this talk, we utilize the rich theory of abstract interpretation to shed light on the overapproximation of the reachable states performed by PDR's frames. Namely, we define an eager version of PDR, called Lambda-PDR, in which all generalizations of counterexamples are used to strengthen a frame, and show that its frames can be formulated as an abstract interpretation algorithm in a logical domain based on Bshouty's monotone theory. Since the frames of Lambda-PDR are tighter than the frames of PDR, the same overapproximation, and more, is present in PDR's frames. We demonstrate that this overapproximation can result in an exponential gap compared to exact forward reachability.

## 3.24 Exploiting Pointer Analysis in Memory Models for Deductive Verification

*Mihaela Sighireanu (ENS Paris-Saclay – Gif-sur-Yvette, FR)*

**License** Creative Commons BY 4.0 International license
© Mihaela Sighireanu
**Joint work of** Quentin Bouillaguet, François Bobot, Mihaela Sighireanu, Boris Yakobowski
**Main reference** Quentin Bouillaguet, François Bobot, Mihaela Sighireanu, Boris Yakobowski: "Exploiting Pointer Analysis in Memory Models for Deductive Verification", in Proc. of the Verification, Model Checking, and Abstract Interpretation – 20th International Conference, VMCAI 2019, Cascais, Portugal, January 13-15, 2019, Proceedings, Lecture Notes in Computer Science, Vol. 11388, pp. 160–182, Springer, 2019.
**URL** https://doi.org//10.1007/978-3-030-11245-5_8

Cooperation between verification methods is crucial to tackle the challenging problem of software verification. The paper focuses on the verification of C programs using pointers and it formalizes a cooperation between static analyzers doing pointer analysis and a deductive verification tool based on first order logic. We propose a framework based on memory models that captures the partitioning of memory inferred by pointer analyses, and complies with the memory models used to generate verification conditions. The framework guided us to propose a pointer analysis that accommodates to various low-level operations on pointers while providing precise information about memory partitioning to the deductive verification. We implemented this cooperation inside the Frama-C platform and we show its effectiveness in reducing the task of deductive verification on a complex case study.

## 3.25 (Un-)Realizability of Condition Synthesis as CHC-SAT

*Yakir Vizel (Technion – Haifa, IL)*

**License** Creative Commons BY 4.0 International license
© Yakir Vizel
**Joint work of** Yakir Vizel, Bat-Chen Rothenberg, Orna Grumberg

*Condition synthesis* takes a program in which some of the conditions in conditional branches are missing, and a specification, and automatically infers conditions to fill-in the holes such that the program meets the specification.

In this talk, we present CoSyn [1], an algorithm for determining the realizability of a condition synthesis problem, with an emphasis on proving unrealizability efficiently. CoSyn is baed on a reduction of the condition synthesis problem to satisfiability of Constrained Horn Clauses (CHCs). In order to allow this reduction to CHCs, we use the novel concept of a *doomed* initial state, which is an initial state that can reach an error state along *every* run of the program. For a doomed initial state $\sigma$, there is no way to make the program safe by forcing $\sigma$ (via conditions) to follow one computation or another. CoSyn encodes the existence of a doomed initial state as CHCs.

CoSyn is implemented in SeaHorn using Spacer as the CHC solver and evaluated it on multiple examples. The evaluation shows that CoSyn outperforms the state-of-the-art syntax-guided tool Cvc5 in proving both realizability and unrealizability. Evaluation also shows that joining forces of CoSyn and Cvc5 outperforms Cvc5 alone, allowing to solve more instances, faster.

**References**

1 B. Rothenberg et.al., *Condition Synthesis Realizability via Constrained Horn Clauses.* NASA Formal Methods – 15th International Symposium, Houston, TX, USA, May 16-18, 2023.

## 3.26 Dataflow Refinement Type Inference

*Thomas Wies (New York University, US)*

**License** Creative Commons BY 4.0 International license
© Thomas Wies
**Joint work of** Zvonimir Pavlinovic, Yusen Su, Thomas Wies
**Main reference** Zvonimir Pavlinovic, Yusen Su, Thomas Wies: "Data flow refinement type inference", Proc. ACM Program. Lang., Vol. 5(POPL), pp. 1–31, 2021.
**URL** https://doi.org//10.1145/3434300

Refinement types enable lightweight verification of functional programs. Algorithms for statically inferring refinement types typically work by reduction to solving systems of constrained Horn clauses extracted from typing derivations. An example is Liquid type inference, which solves the extracted constraints using predicate abstraction. However, the reduction to constraint solving in itself already signifies an abstraction of the program semantics that affects the precision of the overall static analysis. To better understand this issue, we study the type inference problem in its entirety through the lens of abstract interpretation. We propose a new refinement type system that is parametric with the choice of the abstract domain of type refinements as well as the degree to which it tracks context-sensitive control flow information. We then derive an accompanying parametric inference algorithm as an abstract interpretation of a novel data flow semantics of functional programs. We further show that the type system is sound and complete with respect to the constructed abstract semantics. Our theoretical development reveals the key abstraction steps inherent in refinement type inference algorithms. The trade-off between precision and efficiency of these abstraction steps is controlled by the parameters of the type system. Existing refinement type systems and their respective inference algorithms, such as Liquid types, are captured by concrete parameter instantiations.

## 3.27 Timing Analysis by Abstract Interpretation

*Reinhard Wilhelm (Universität des Saarlandes – Saarbrücken, DE)*

**License** Creative Commons BY 4.0 International license
© Reinhard Wilhelm
**Main reference** Reinhard Wilhelm: "Real time spent on real time", Commun. ACM, Vol. 63(10), pp. 54–60, 2020.
**URL** https://doi.org//10.1145/3375545

Hard real-time systems need a proof that they keep their deadlines. This proof is produced by a code-level WCET analysis and a schedulability analysis for a set of tasks to be executed on the same platform. A code-level WCET analysis computes a safe upper bound for all execution times of a task. The only sound WCET analysis used widely in industry, developed by my group in Saarbrücken, uses several abstract interpretations to safely and efficiently derive such upper bounds for tasks to be executed on single-core platforms. Central for the practicality of our approach are adequate abstractions of the architecture of the execution platform. Thus, the architecture is an integral part of the semantics from which abstract interpretations of real-time programs are derived. This entails a number of peculiarities compared to more conventional abstract interpretations. They are more concerned with the occupancy of platform resources than with the values contained in those resources. The occupancy of resources, e.g. the cache contents or the usage of bus bandwidth influence the timing behavior. Iteration over loops, as part of the fixed-point iteration, needs to consider machine parameters to achieve accuracy. The first iteration of a loop typically loads the cache, and later iterations profit from this cache loading. Therefore, first and non-first iterations may have vastly different execution times. In order to obtain accurate execution times for loops, the iteration of the analysis first needs to stabilize the execution time of the loop body. In my talk I gave some such examples of peculiarities [2]. The main part was concerned with the development history of our WCET-analysis technology [1].

### References


1 Reinhard Wilhelm. *Real time spent on real time*. Commun. ACM 63(10): 54-60 (2020)
2 Jan Reineke and Reinhard Wilhelm. *Static Timing Analysis – What is Special?*. Semantics, Logics, and Calculi 2016: 74-87


# 4 Working groups

## 4.1 Soundness requirements, transparency of assumptions

*Bor-Yuh Evan Chang (University of Colorado – Boulder, US) and Raphaël Monat (INRIA Lille, FR)*

**License** Creative Commons BY 4.0 International license
© Bor-Yuh Evan Chang and Raphaël Monat

We held a discussion session around soundness requirements, and the transparency of assumptions made to ensure an analysis or a tool is sound.

The starting point of this discussion was the Soundiness manifesto [1]. The goal of this manifesto is to raise awareness and highlight that there likely are – at least small and sometimes intentional – discrepancies between the formalized concrete semantics and the actual implementation. This might be due to the concrete semantics being formalized in a research paper not matching the whole semantics of the targetted language – either to simplify the presentation, make a trade-off in the analysis, or because modern programming languages are large and complex. However, the soundiness paper is sometimes misunderstood and used as an excuse to give up on establishing soundness results. Thus, we argue (and believe it to be in line with the original intent of the soundiness manifesto) for transparency in soundness claims: researchers should state and make explicit the limitations of their soundness claims, and in particular the assumptions that are being made. The reasonableness of those assumptions should be empirically evaluated whenever possible.

In order to ensure this transparency at the tool level, we suggest tools should report the assumptions they made alongside the properties they have been able – or unable – to prove.

For example, let us consider a C program where an extern function is called. For early works, it might be sufficient to abort the analyses when encountering such cases. Then, a reasonable approach to analyze this program is to assume the function does not have side-effects and that it returned a value abiding by its return type signature. Instead of silently performing this assumption and continuing the analysis, we argue that for transparency, this assumption must be stored and displayed during the analyzer's report.

As another example, consider in Java, a program can modify the fields of an object using reflection. Because perhaps such uses of reflection are rare in the code of interest or that it is considered bad practice, a concrete semantics that models arbitrary field updates at any point using reflection would lead to unrealistically imprecise and pessimistic analyses. Thus, a reasonable approach is to make the assumption there is no reflective field update and exclude it from the concrete semantics. Transparency means checking for the potential reflective field update on the runs of the analysis are reporting it if it is encountered. And transparency means empiricially evaluating the prevalence of reflective field update on a representative corpus.

This approach is being prototyped within the Mopsa static analyzer for its C, Python and multilanguage analyses. We believe it may also help developers better understand the outputs of static analyzers.

Thanks to this approach, the assumptions made by a static analyzer are now clearly defined in the implementation. In turn, this simplifies establishing theoretical soundness claims on paper: it is easier to list the current limitations, and soundness assumptions of a given analysis in practice.

Of course, developers of commercial static analyzers that are used to certify safety-critical systems with respect to regulations such as DO-178C are deeply aware of this transparency-based approach. They internally have documents describing the trusted computing base, and the exact soundness theorem, with all potential assumptions. While this is too time-consuming and stringent for proof-of-concepts and early academic software, we believe both that it is an ideal and a must for usable tools.

Our next step is to discuss with the authors of the original soundiness manifesto. It was suggested in our plenary recap and discussion to push for an evaluation of the transparency of static analysis tools in selected conferences performing artifact evaluation.

#### References


1 Benjamin Livshits, Manu Sridharan, Yannis Smaragdakis, Ondrej Lhoták, José Nelson Amaral, Bor-Yuh Evan Chang, Samuel Z. Guyer, Uday P. Khedker, Anders Møller, Dimitrios Vardoulakis: In defense of soundiness: a manifesto. Commun. ACM 58(2): 44-46 (2015)


## 4.2 Expressive Domains

*Gidon Ernst (LMU München, DE)*

**License** 🅭 Creative Commons BY 4.0 International license
© Gidon Ernst

We discussed abstract domains for data structures and domains that can capture complex properties. We observed that a large variety of techniques have been described in the literature, e.g., for the data structures sequences/strings, sets, multisets, arrays, pointer-structures, and trees/general ADTs, and which e.g., can express properties like sortedness or initialization. However, while for numerical domains, we have nice open-source libraries that can easily be embedded into larger use-cases, it was noted that this is hardly the case for such expressive domains despite the fact that there are many implementations. We then discussed the challenges in the design of both the domains itself as well as a uniform interface that could capture their features and commonalities: The first challenge is to define a common interface in the first place. Since domains are usually designed to tackle specific properties, the operations they offer are tuned to their respective use cases and therefore not uniform. As an example, it seemed unclear whether it is possible to represent the features offered by domains as part of a CHC solver. As possible future work, a first step towards usable off-the-shelf libraries was to investigate the design of a suitable API. A second challenge is the combined use of abstract domains, for example when nesting domains for element types inside container types. There is no canonical choice to do so many, and which is appropriate strongly depends on the application. This kind of integration often relies on the tight integration of the data structures representing the abstract domains at the implementation level. This kind of compositionality, however, is hard to achieve general because domains – even for the same purpose – may rely on rather different techniques. A related problem is how to provide on-line transfer from and to symbolic representations. We idientified as a question for the efficient combination of domains, whether some of the higher-level domains benefit from some specialized operators of the underlying abstract domains. The third challenge is that it is hard to design abstract domains that are robust. A lot of knowledge goes into selecting the right abstraction for a given (sub-)problem and choosing the an inappropriate one incurs high computational cost. Unfortunately, this is hard to determine automatically, and moreover may rely on user-provided partial specifications as guidance.

## 4.3 Education: Teaching Abstract Interpretation to the Masses

*Pietro Ferrara (University of Venice, IT)*

**License** Creative Commons BY 4.0 International license
© Pietro Ferrara

During this working group, the discussion was divided into two main themes: teaching to graduate students, and to undergraduates or, generally speaking, practitioners. In the first case, a lot of good material is available. A non-comprehensive list (that should be further integrated) to the best of the participants' knowledge is the following one: books [1, 2]; slides [3, 4, 5, 6, 7, 8]; tools [9, 10].

For undergrads and/or practitioners there is less material available. After an open discussion with all the seminar participants, we were able to assemble the following list: [11, 12, 13].

Generally speaking, the first idea would be to augment undergraduate courses on the implementation of interpreters with the idea of abstraction (e.g., [7]). Another proposal was to design a generic interface to interact (ranging from the application to the extension) with static analyzers through a unique library, avoiding having analyzers that require being bound to specific programming languages. This is a rather standard approach adopted by our communities that over the years achieved better visibility and popularity (e.g., the machine learning community with libraries such as Scikit-learn and Pytorch, or the theorem-proving community with libraries such as SMT-lib).

Overall, the participants agreed that the community is missing a place where educational material about static analysis can be shared, as well as introduction courses in popular MOOC platforms.

**References**

1 Patrick Cousot: Principles of Abstract Interpretation, https://mitpress.mit.edu/9780262044905/principles-of-abstract-interpretation/
2 Rival and Yi: Introduction to Static Analysis, https://mitpress.mit.edu/9780262043410/introduction-to-static-analysis/
3 Patrick Cousot's course at MIT: http://web.mit.edu/afs/athena.mit.edu/course/16/16.399/www/
4 Feret, Giet, Rival course at ENS Paris: https://www.di.ens.fr/~rival/semverif-2023/
5 Miné, Urban, Feret, Rival in Paris (MPRI): https://www-apr.lip6.fr/~mine/enseignement/mpri/current/
6 Ferrara (old – 2012) course at ETH: https://ethz.ch/content/dam/ethz/special-interest/infk/chair-program-method/pm/documents/Education/Courses/SS2012/SPA/Lectures.zip
7 Evan course at the University of Colorado Boulder: https://csci3155.cs.colorado.edu/f22/ and https://github.com/csci3155/
8 Jan Midtgaard: Abstract Interpretation (2015 Winter School) https://janmidtgaard.dk/aiws15/
9 LiSA (Ca' Foscari University of Venice): https://github.com/lisa-analyzer/lisa
10 MOPSA (Sorbonne Université): https://gitlab.com/mopsa/mopsa-analyzer
11 Anders Møller and Michael I. Schwartzbach: Static Program Analysis https://cs.au.dk/~amoeller/spa/
12 Manuel's older intro to AI, and AI for (C)LP (a tutorial from the early 90's. but we still use it sometimes) https://cliplab.org/logalg/slides/B_ai.pdf
13 Manuel: Some tutorials on PLAI/CiaoPP using our analyzers embedded in web pages (work in progress) https://ciao-lang.org/ciao/build/doc/ciaopp_tutorials.html/

## 4.4 Community Infrastructure – Interest Group in Static Analysis

*Pietro Ferrara (University of Venice, IT)*

**License** Creative Commons BY 4.0 International license
© Pietro Ferrara

The main outcome of this working group was a proposal to build up a community infrastructure to facilitate the promotion of events in our community, stimulate interactions, and provide better visibility of the main outcomes in our field. In particular, currently, we have various groups on social networks (such as Facebook and Linkedin) that are mostly silent. However, every year we have a series of regular events that help the networking activities of our community, such as conferences like SAS and VMCAI, workshops like SOAP, and events organized by various institutions such as the "Dependable and Secure Software Systems" workshop at ETH or the Challenges of Software Verification Symposium at Ca' Foscari University of Venice (just to name a few the participants were aware of, but such a list should be quite expanded).

The final proposal of the working group was to: (i) build up a website that contains information about events and materials about static analysis (taking inspiration from https://microservices.community), (ii) open a mailing list about announcements regarding scientific activities and opportunities in static analysis, and (iii) establish an Interest Group in Static Analysis (IGSA) with an advisory board that supervise the overall process. Potentially this might become an ACM Special Interest Group, an IFIP Working Group, or something else, based on the success of the initiative.

## 4.5 Abstract Interpretation and Other Formal Methods

*Arie Gurfinkel (University of Waterloo, CA)*

License Creative Commons BY 4.0 International license
© Arie Gurfinkel

This working group discussed potential new connections between Abstract Interpretation and other Formal Methods. Based on the interests and expertise of the participants, the other formal methods that the group focused on were: Deductive Verification, Dynamic Techniques, Deduction, Model Checking, and Practical and Industrial applications.

For deductive verification, there is a desire for Abstract Interpretation to infer complex specifications, especially in the context of memory analysis. Currently, most developed Abstract Interpretation techniques are geared towards numeric domains that are not sufficiently structurally complex (i.e., no quantifiers, no memory separation, etc.). Deductive verification also requires very reliable tools that reliably provide an answer since such tools involve a direct and active interaction with a user.

For dynamic techniques, there is a desire for better abstract domains for strings. Perhaps based on the deep connection between automata and regular languages. There is also a potential in using dynamic analysis to infer context for Abstract Interpretation-based static analysis. This, for example, is already used in debloading projects such OCCAM at SRI. The group felt that this combination might also create new challenges to maintaining soundness of the analysis and/or articulating the soundness conditions clearly.

For deduction, it was discussed that the main challenge is to identify an insight for why a proof works or fails. In this context, a complex but hard-to-understand specification, such as an automatically generated inductive invariant, is not very helpful. Perhaps the results of Abstract Interpretation can be presented in a more readable form by using some pre-defined set of predicates.

For model checking, there is an interest in building Abstract Interpretation models that explain core model checking algorithms such as IC3 and PDR. There is recent work on this subject, but the group agrees that better understanding is required to capture the many nuances of these algorithms. There is also a potential application in using Abstract Interpretation to infer temporal specifications.

Group members from industry, were interested in Abstract Interpretation for probabilistic programs, with application to differential privacy. There is also interest in Abstract Interpretation in aid of property-based testing.

Overall, the working group concluded that there is a lot of interest in both applications of Abstract Interpretation in different domains, and in research on challenges that are posed by these domains. The group members from Deductive Verification have been most vocal, and, perhaps the combination of Abstract Interpretation and Deductive Verification is most suited for a follow up seminar.

## 4.6 Tools and Applications for Abstract Interpretation

*Falk Howar (TU Dortmund, DE)*

License Creative Commons BY 4.0 International license
© Falk Howar

We had a discussion on challenges that developers of research tools must address to enable industrial collaborations and on possible working directions for the research community to foster the development of robust and scalable research tools. Among the attendees were researchers, academic tools developers, as well as developers of commercial tools and industrial researchers from companies that use formal methods tools in their product development, presenting both, academic perspectives, and insights from industry.

The initial focus of the discussion was on the potential and challenges of using research tools effectively in real-world scenarios and started with the observation that it is often difficult to use research tools in industrial collaborations.

Lessons learned by academic developers included the following points:

Usefulness: It is important to have findings and to make these findings accessible to project partners, who typically are not experts in a particular formal method or a particular research tool.

Robustness: It is important that tools can be used in industrial settings. This often requires operation in an automated analysis pipeline, e.g., a build system. In such a setting, an analysis will not be run manually executed one target but is executed as part of a build process.

Participants from companies complemented these observations with the following points:

Helpful features: To support robustness, certain features help applicability that are (usually) not important in an academic context: since industrial codebases in many cases cannot be extended with annotations for a tool, it is important to support declarative configurations and external annotations. Usefulness can be increased by support for problem extraction (i.e., interpretation / mapping of results to code) and by providing information on relationships between alarms (i.e., root causes).

Project formats: Analyzing an industrial codebase is not an easy task, especially when done by someone who is not a contributor to that codebase. Typically, success hinges on support from developers, access to architecture documentation, and on communication between maintainers and researchers. Project formats should accommodate these success factors.

Licenses: Copyleft licenses are a big obstacle in industry if code is under such a license and the industrial partner must look at the code (e.g., to understand how a tool work).

Eventual Payoff: It is important for industrial partners to understand if and what the eventual payoff can be when using a technique or tool. As a concrete successful example, the technique of unit proving was mentioned. In the experience of one participant, maintainability and conformance to low level specs seems to be much easier with unit proving than with testing once it is established in the development process.

After the initial collection of lessons learned, discussion evolved around dealing with theoretical limitations of techniques and (potential) usability requirements.

Dealing with Limitations: Commercial tool developers (AbsInt and Astree tools) reported that in their experience mostly scalability and precision are bottlenecks and that proofs tend to work on slices, but do not work on whole programs. It was discussed if and to what extend a target can be modified in order to scale or work around limitations (e.g., loop unrolling, transformations at the IR level). Consensus was that, while it is usually not possible to change a customer's code, transformations typically happen during analysis (with some limitations in safety-critical domains). It was deemed important that the result/verdict of an analysis can be explained to the user and in terms of the original code.

Usability: There was no consensus on or shared understanding of usability requirements. Most participants agreed that a user interface is not a strong requirement for a research tool, but that usability is still important in the sense that users need to be able to understand the findings of a tool (e.g., based on CFGs).

The discussion concluded with plans for actions the research community could take to recognize and facilitate the development of robust research tools. Case studies were identified as a relevant format: It shows potential users how a tool can be used on a realistic example and forces developers to invest in robustness and applicability. Examples of existing case studies and publicly available systems included the effort to verify the Linux kernel, medical systems (e.g., infusion pumps, a pacemaker), and a wheel-break system. Participants agreed that it will be important to promote case study papers in academic conferences. Published case studies should include a public repository with the problem, corresponding artifacts (e.g., documentation), and an experience report.

## 4.7 Hyperproperties verification

*Isabella Mastroeni (University of Verona, IT)*

**License** Creative Commons BY 4.0 International license
© Isabella Mastroeni

Hyperproperties are intended as sets of sets of traces, for a more general point of view, we can see them as sets of properties. A standard example of hyperproperty is non-interference, a 2-safety property that requires to compare pairs of executions. The question for stimulating the discussion was about the right implication formalization for hyperproperties. In the literature we mainly find two different approaches: (1) The first one allows to add new elements in the hyper set, but all the elements in the stronger property must be in the weaker one, more formally, let $A, B, C, D, E$ sets, then $\{A, B, C\}$ implies $\{A, B, C, D, E\}$; (2) the second one consists in allowing to approximate the inner elements without adding any new element, formally let $A, B, C, A', B', C'$ be sets such that $A \subseteq A'$, $B \subseteq B'$ and $C \subseteq C'$, then $\{A, B, C\}$ implies $\{A', B', C'\}$ but $\{A, B, C\}$ does not imply $\{A', B', C', D'\}$ where $D'$ does not contain any of $A, B, C$. Hence, which of these orders properly capture the approximation order between hyperproperties.

There are several approaches to noninterference that can help in understanding the issue, in particular those based on hyper analysis of programs. But also there are other relevant formal approaches to hyper properties verification, such as the hyper Hoare logic, useful for understanding the issue.

By reasoning on both these formal approaches, the discussion ended up understanding that probably there is no a right or better order, in particular (2) is necessary for computing fixpoints programs/properties in an hyper levels both in the analysis and in the logic, while (1) is used for deciding whether a property imply another one, again both in the analysis and in the logic (consequence rule). Hence, the conclusion was that (1) is the approximation order while (2) is the computational order, which maybe must co-exists when dealing with hyper property verification. We believe that this final observation, even if it may appear quite "simple", may represent anyway an important first step for really understanding the connection between standard approaches to static analysis and the emerging hyper property verification issue. All participants agree that surely further understanding of the problem is necessary.

## 4.8 AI for AI

*Antoine Miné (Sorbonne University – Paris, FR)*

**License** 🅭 Creative Commons BY 4.0 International license
© Antoine Miné

A small working group entitled "AI for AI" discussed existing and possible future interactions between Abstract Interpretation (AI) and Machine Learning (ML). Both directions, AI to help ML and ML to help AI, were considered.

We started with the observation that, with the progress and increasing use of ML systems, ensuring their correctness also became an increasing concern. Since 2015, the field of formal verification techniques applied to ML systems has been growing steadily. A large amount of work targets the verification of trained neural networks, but other models are considered as well (e.g., support vector machines, ensemble trees). Moreover, many kinds of formal methods have been adapted to work on ML systems, including: SMT solving, constraint solving, optimisation, model-checking, Abstract Interpretation. Notable works on Abstract Interpretation for ML models include the developments of dedicated numeric abstract domains able to prove local robustness properties and fairness properties. A more detailed review of recent work can be found in Urban et al. 2021 [1]. However, given the pace of the research in this area, such reviews become outdated very quickly.

We then discussed what we perceived as some of the current limitations of AI methods for ML and possible area of improvements.

- The verification of robustness has been successful for models of moderate size, and for local robustness against perturbation and adversarial attacks. Future challenge concern more global robustness properties. Numeric abstractions and partitioning techniques from AI could be key to achieve scalability.
- ML research has recently focused on interpretability and explainability of ML models. This constitutes an interesting opportunity to apply formal verification and obtain sound and automatic guarantees on these properties. We postulate that AI-specific techniques, such as relational abstractions, might help in evaluating out inputs influence the outcome of a ML model.
- In general, however, a key difficulty for formal verification (including Abstract Interpretation) is the lack of formal specification of the properties to prove. This is an area that must be improved in order to unlock further progress.
- Most works in ML ignore the effect of floating-point rounding errors. AI techniques exist to soundly verify floating-point programs and to evaluate the effect of rounding errors formally. Applying such techniques to ML models could prove useful.

The working group also considered how ML could help AI. An important aspect is to keep the formal guarantees of the combined AI and ML method by maintaining soundness, even in the absence of soundness guarantees for the ML part.

- One possible use of ML in AI is the automatic parameterization of static analyzers to optimize the efficiency and precision of the analysis within the parameter space (e.g., choice of abstract domains, level of relationality and sensitivity, etc.).

- A specific aspect of AI is the use of widening operators to automatically infer complex invariants. The heuristic nature of widenings makes them an interesting target for ML techniques: the heuristics "guessing" candidate invariants do not need strong soundness guarantees, given that candidate invariants can be checked with a classic, sound stability tests. This idea could be extended to the inference of contract pre/post-conditions to unlock efficient modular analyses.

As concluding remark, we observed that early research in formal verification of ML was mainly conducted by formal method researchers, adapting on ML models the techniques they were familiar with for the verification of software. Given the substantial differences between ML models and programs, and the need to extract and formalize the properties we are interested to verify on ML models, we believe that progress in this field requires a tighter collaboration of experts in ML and experts in AI.

### References

1 C. Urban and A. Miné. A review of formal methods applied to machine Learning. Technical report, Computing Research Repository (arXiv) (CoRR), Apr. 2021. `http://www-apr.lip6.fr/~mine/publi/article-urban-mine-mlfm2021.pdf`.

## Participants

- Vincenzo Arceri
University of Parma, IT
- Musard Balliu
KTH Royal Institute of Technology – Stockholm, SE
- Anna Becchi
Bruno Kessler Foundation – Trento, IT
- Dirk Beyer
LMU München, DE
- Bor-Yuh Evan Chang
University of Colorado – Boulder, US
- Patrick Cousot
New York University, US
- David Delmas
Airbus S.A.S. – Toulouse, FR
- Gidon Ernst
LMU München, DE
- Pietro Ferrara
University of Venice, IT
- Eden Frenkel
Tel Aviv University, IL
- Isabel Garcia-Contreras
University of Waterloo, CA
- Roberto Giacobazzi
University of Verona, IT
- Roberta Gori
University of Pisa, IT
- Arie Gurfinkel
University of Waterloo, CA
- Reiner Hähnle
TU Darmstadt, DE
- Ben Hermann
TU Dortmund, DE
- Manuel Hermenegildo
IMDEA Software Institute – Pozuelo de Alarcón, ES & UPM – Madrid, ES
- Falk Howar
TU Dortmund, DE
- Daniel Kästner
AbsInt – Saarbrücken, DE
- Temesghen Kahsai
Amazon Lab 126, US
- Matthieu Lemerre
CEA LIST – Gif-sur-Yvette, FR
- Isabella Mastroeni
University of Verona, IT
- Antoine Miné
Sorbonne University – Paris, FR
- Raphaël Monat
INRIA Lille, FR
- Peter Müller
ETH Zürich, CH
- Jorge Navas
Certora – Seattle, US
- Marie Pelleau
Université Côte d'Azur – Sophia Antipolis, FR
- Ruzica Piskac
Yale University – New Haven, US
- Francesco Ranzato
University of Padova, IT
- Gerhard Schellhorn
Universität Augsburg, DE
- Ilya Sergey
National University of Singapore, SG
- Sharon Shoham Buchbinder
Tel Aviv University, IL
- Mihaela Sighireanu
ENS Paris-Saclay – Gif-sur-Yvette, FR
- Yakir Vizel
Technion – Haifa, IL
- Thomas Wies
New York University, US
- Reinhard Wilhelm
Universität des Saarlandes – Saarbrücken, DE
- Enea Zaffanella
University of Parma, IT

Report from Dagstuhl Seminar 23291

# Parameterized Approximation: Algorithms and Hardness

**Karthik C. S.*[1], Parinya Chalermsook*[2], Joachim Spoerhase*[3], Meirav Zehavi*[4], and Martin Herold†[5]**

1 **Rutgers University – New Brunswick, US.** karthik.0112358@gmail.com
2 **Aalto University, FI.** parinya.chalermsook@aalto.fi
3 **University of Sheffield, GB.** j.spoerhase@sheffield.ac.uk
4 **Ben Gurion University – Beer Sheva, IL.** zehavimeirav@gmail.com
5 **MPI für Informatik – Saarbrücken, DE.** mherold@mpi-inf.mpg.de

**Abstract**

Parameterization and approximation are two established approaches of coping with intractability in combinatorial optimization. In this Dagstuhl Seminar, we studied parameterized approximation as a relatively new algorithmic paradigm that combines these two popular research areas. In particular, we analyzed the solution quality (approximation ratio) as well as the running time of an algorithm in terms of a parameter that captures the "complexity" of a problem instance.

While the field has grown and yielded some promising results, our understanding of the area is rather ad-hoc compared to our knowledge in approximation or parameterized algorithms alone. In this seminar, we brought together researchers from both communities in order to bridge this gap by accommodating the exchange and unification of scientific knowledge.

**Seminar** July 16–21, 2023 – https://www.dagstuhl.de/23291
**2012 ACM Subject Classification** Theory of computation → Computational complexity and cryptography; Theory of computation → Design and analysis of algorithms
**Keywords and phrases** approximation algorithms, Hardness of approximation, Parameterized algorithms
**Digital Object Identifier** 10.4230/DagRep.13.7.96

## 1 Executive Summary

*Joachim Spoerhase (University of Sheffield, GB)*
*Karthik C. S. (Rutgers University – New Brunswick, US)*
*Parinya Chalermsook (Aalto University, FI)*
*Meirav Zehavi (Ben Gurion University – Beer Sheva, IL)*

**License** Creative Commons BY 4.0 International license
© Joachim Spoerhase, Karthik C. S., Parinya Chalermsook, and Meirav Zehavi

Parameterization and approximation are two popular approaches of coping with intractability in combinatorial optimization. They have gained substantial maturity over the last decades, leading to tight bounds for many fundamental computational problems and beautiful algorithmic techniques that show surprising interplay between algorithms and various areas of mathematics. In this Dagstuhl Seminar, we studied parameterized approximation as a relatively new algorithmic paradigm that combines these two popular research areas. In particular, we analyzed the solution quality (approximation ratio) as well as the running

* Editor / Organizer
† Editorial Assistant / Collector

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

time of an algorithm in terms of a parameter that captures the "complexity" of a problem instance.

While the field has grown and yielded some promising results, our understanding of the area is rather ad-hoc compared to our knowledge in approximation or parameterized algorithms alone. In this seminar, we brought together researchers from both communities in order to bridge this gap by accommodating the exchange and unification of scientific knowledge.

Our first goal was to foster a transfer of techniques between the classic fields of approximation and parameterization. We discussed how recent developments in one research area can be transferred to another. Towards this, we organized five invited one-hour tutorials (one on each morning) that were delivered by leading experts from the fields and which we believe helped exchange of state-of-the-art techniques. The tutorial topics covered parameterized or polynomial time approximation algorithms for graph edit and network design problems, for clustering problems, matroid constrained maximization problems, as well as the hardness of parameterized approximation of problems arising in error correcting codes.

Our second goal was to systematically identify important research directions and concrete open problems in research areas that are relevant to parameterized approximation. Towards this, we organized a panel discussion on the first seminar day led by experts from parameterized algorithms, approximation algorithms, hardness of approximation, but also from neighboring areas such as fine-grained complexity and coding theory. A vibrant discussion ensued between the moderator, panelists, but also the other participants. Moreover, we organized two open discussion sessions, which were open to any participant to bring up a topic they would like to discuss with the other participants: This could be, for example, concrete open problems, more general research directions, or highlights. Besides several concrete open problems suggested by particpants, there were two contributions to the sessions that gave overviews over a coherent collection of open questions on hardness of parameterized approximation under Gap-ETH and beyond, as well as on the parameterized (in-)approximability of clustering problems.

Our third goal was to bolster the creation of new collaborations between researchers in the two communities by encouraging the participants to actively discuss the suggested directions for open problems. It was therefore a particular priority for us to provide sufficient time for collaboration. Towards this, we reserved time slots for six collaboration sessions and one session for progress reports.

We thank Martin Herold for collecting the abstracts from the participants and for his assistance with creating and editing this report.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Hardness of Approximation in P via Short Cycle Removal: Cycle Detection, Distance Oracles, and Beyond

*Amir Aboud (Weizmann Institute – Rehovot, IL)*

**License** Creative Commons BY 4.0 International license
© Amir Aboud
**Joint work of** Amir Aboud, Karl Bringmann, Seri Khoury, Or Zamir
**Main reference** Amir Aboud, Karl Bringmann, Seri Khoury and Or Zamir: Hardness of Approximation in P via Short Cycle Removal: Cycle Detection, Distance Oracles, and Beyond. STOC 2022.
**URL** https://doi.org/10.1145/3519935.3520066

This talk will overview a new technique for gap amplification called "short cycle removal" and it applications for hardness of approximation for polynomial time problems. In particular, we will present lower bounds on the approximation factor of distance oracles that preprocess a graph in almost-linear time and answer distance queries in almost-constant time. Based on joint works with Karl Bringmann, Nick Fischer, Seri Khoury, and Or Zamir.

## 3.2 Baby PIH: Parameterized Inapproximability of Min CSP

*Venkatesan Guruswami (University of California – Berkeley, US)*

**License** Creative Commons BY 4.0 International license
© Venkatesan Guruswami

The Parameterized Inapproximability Hypothesis (PIH) is the analog of the PCP theorem in the world of parameterized complexity. It asserts that no FPT algorithm can distinguish a satisfiable 2CSP instance from one which is only $(1-\varepsilon)$-satisfiable (where the parameter is the number of variables) for some constant $0 < \varepsilon < 1$.

We consider a minimization version of CSPs (Min-CSP), where one may assign $r$ values to each variable, and the goal is to ensure that every constraint is satisfied by some choice among the $r \times r$ pairs of values assigned to its variables (call such a CSP instance $r$-list-satisfiable). We prove the following strong parameterized inapproximability for Min CSP: For every $r \geq 1$, it is W[1]-hard to tell if a 2CSP instance is satisfiable or is not even $r$-list-satisfiable. We refer to this statement as "Baby PIH", following the recently proved Baby PCP Theorem. Our proof adapts the combinatorial arguments underlying the Baby PCP theorem, overcoming some significant obstacles that arise in the parameterized setting.

An extension of our result to an average-version of Baby PIH would prove the inapproximability of parameterized $k$-ExactCover, a notorious open problem.

## 3.3 A $(3/2 + \epsilon)$-Approximation for Multiple TSP with a Variable Number of Depots

*Matthias Kaul (TU Hamburg, DE)*

**License** Creative Commons BY 4.0 International license
© Matthias Kaul
**Joint work of** Max Deppert, Matthias Kaul, Matthias Mnich
**Main reference** Max Deppert, Matthias Kaul, Matthias Mnich: "A $(3/2 + \epsilon)$-Approximation for Multiple TSP with a Variable Number of Depots", in Proc. of the 31st Annual European Symposium on Algorithms, ESA 2023, September 4-6, 2023, Amsterdam, The Netherlands, LIPIcs, Vol. 274, pp. 39:1–39:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
**URL** https://doi.org//10.4230/LIPICS.ESA.2023.39

One of the most studied extensions of the famous Traveling Salesperson Problem (TSP) is the Multiple TSP: a set of $m \geq 1$ salespersons collectively traverses a set of $n$ cities by $m$ non-trivial tours, to minimize the total length of their tours. This problem can also be considered to be a variant of Uncapacitated Vehicle Routing, where the objective is to minimize the sum of all tour lengths. When all $m$ tours start from and end at a single common depot $v_0$, then the metric Multiple TSP can be approximated equally well as the standard metric TSP, as shown by Frieze (1983)[1]. The metric Multiple TSP becomes significantly harder to approximate when there is a set $D$ of $d \geq 1$ depots that form the starting and end points of the $m$ tours. For this case, only a $(2 - 1/d)$-approximation in polynomial time is known, as well as a 3/2-approximation for constant $d$ which requires a prohibitive run time of $n^{\Theta(d)}$ (Xu and Rodrigues, INFORMS J. Comput., 2015)[3]. A recent work of Traub, Vygen and Zenklusen (STOC 2020)[2] gives another approximation algorithm for metric Multiple TSP with run time $n^{\Theta(d)}$, which reduces the problem to approximating metric TSP. In this paper we overcome the $n^{\Theta(d)}$ time barrier: we give the first efficient approximation algorithm for Multiple TSP with a variable number d of depots that yields a better-than-2 approximation. Our algorithm runs in time $(1/\epsilon)^{O(d \log d)} \cdot n^{O(1)}$, and produces a $(3/2 + \epsilon)$-approximation with constant probability. For the graphic case, we obtain a deterministic 3/2-approximation in time $2^d \cdot n^{O(1)}$.

**References**

1 Alan M. Frieze. An extension of christofides heuristic to the k-person travelling salesman problem. *Discret. Appl. Math.*, 6(1):79–83, 1983.
2 Vera Traub, Jens Vygen, and Rico Zenklusen. Reducing path TSP to TSP. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 14–27. ACM, 2020.
3 Zhou Xu and Brian Rodrigues. A 3/2-approximation algorithm for the multiple TSP with a fixed number of depots. *INFORMS J. Comput.*, 27(4):636–645, 2015.

## 3.4 Parameterized Approximability of $F$-Deletion Problems

*Euiwoong Lee (University of Michigan – Ann Arbor, US)*

**License** Creative Commons BY 4.0 International license
© Euiwoong Lee

For a fixed class $F$ of graphs, the $F$-Deletion problem, given a graph $G$, asks to remove the minimum number of vertices so that the resulting graph belongs to the class $F$. The study of various $F$-Deletion problems has led to interesting connections between approximation algorithms and parameterized algorithms, ultimately leading to parameterized approximation algorithms. In this talk, we survey known results for subgraph, induced subgraph, and minor deletion problems.

## 3.5 The Cut Covering Lemma

*Jason Li (University of California – Berkeley, US)*

**License** Creative Commons BY 4.0 International license
© Jason Li
**Main reference** Stefan Kratsch, Magnus Wahlström: "Compression via Matroids: A Randomized Polynomial Kernel for Odd Cycle Transversal", ACM Trans. Algorithms, Vol. 10(4), pp. 20:1–20:15, 2014.
**URL** https://doi.org//10.1145/2635810

We present the cut covering lemma, a classic result by Kratsch and Wahlstrom (2012)[1] in FPT literature with many applications to kernelization. Loosely speaking, the lemma states that on any graph with k terminal vertices, there is a smaller graph on $O(k^3)$ vertices that preserves the complete cut structure of the graph. We discuss connections to the field of fast graph algorithms, in particular the implications of an approximate version of the cut covering lemma with $O(k)$ vertices.

### References

1 Stefan Kratsch and Magnus Wahlström. Compression via matroids: a randomized polynomial kernel for odd cycle transversal. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 94–103. SIAM, 2012.

## 3.6 Parameterized Approximation Schemes for Clustering with General Norm Objectives

*Dániel Marx (CISPA – Saarbrücken, DE)*

**License** Creative Commons BY 4.0 International license
© Dániel Marx
**Joint work of** Fateme Abbasi, Sandip Banerjee, Jaroslaw Byrka; Parinya Chalermsook; Ameet Gadekar; Kamyar Khodamoradi; Roohani Sharma; Joachim Spoerhase
**Main reference** Fateme Abbasi, Sandip Banerjee, Jaroslaw Byrka, Parinya Chalermsook, Ameet Gadekar, Kamyar Khodamoradi, Dániel Marx, Roohani Sharma, Joachim Spoerhase: "Parameterized Approximation Schemes for Clustering with General Norm Objectives", CoRR, Vol. abs/2304.03146, 2023.
**URL** https://doi.org//10.48550/ARXIV.2304.03146

We consider the well-studied algorithmic regime of designing a $(1+\epsilon)$-approximation algorithm for a $k$-clustering problem that runs in time $f(k, \epsilon)\text{poly}(n)$. Our main contribution is a clean and simple EPAS that settles more than ten clustering problems (across multiple well-studied objectives as well as metric spaces) and unifies well-known EPASes. Our algorithm gives EPASes for a large variety of clustering objectives (for example, $k$-means, $k$-center, $k$-median, priority $k$-center, $\ell$-centrum, ordered $k$-median, socially fair $k$-median aka robust $k$-median, or more generally monotone norm $k$-clustering) and metric spaces (for example, continuous high-dimensional Euclidean spaces, metrics of bounded doubling dimension, bounded treewidth metrics, and planar metrics). Key to our approach is a new concept that we call bounded $\epsilon$-scatter dimension – an intrinsic complexity measure of a metric space that is a relaxation of the standard notion of bounded doubling dimension.

## 3.7 FPT Approximation Schemes for Matroid-constrained Problems

*Hadas Shachnai (Technion – Haifa, IL)*

**License** Creative Commons BY 4.0 International license
© Hadas Shachnai
**Joint work of** Ilan Doron Arad, Ariel Kulik, Hadas Shachnai
**Main reference** Ilan Doron Arad, Ariel Kulik, Hadas Shachnai: "Budgeted Matroid Maximization: a Parameterized Viewpoint", CoRR, Vol. abs/2307.04173, 2023.
**URL** https://doi.org//10.48550/ARXIV.2307.04173

We study budgeted variants of well known maximization problems with multiple matroid constraints. Given an $\ell$-matchoid $\mathcal{M}$ on a ground set $E$, a profit function $p$ and a cost function $c$ on $E$, and a budget $B$, the goal is to find in the $\ell$-matchoid a feasible set $S$ of maximum profit $p(S)$ subject to the budget constraint, i.e., $c(S) \leq B$. The *budgeted $\ell$-matchoid* (BM) problem includes as special cases budgeted $\ell$-dimensional matching and budgeted $\ell$-matroid intersection. A strong motivation for studying BM from parameterized viewpoint comes from the APX-hardness of unbudgeted $\ell$-dimensional matching (i.e., $B = \infty$) already for $\ell = 3$. Nevertheless, while there are known FPT algorithms for the unbudgeted variants of the above problems, the *budgeted* variants are studied here for the first time through the lens of parameterized complexity.

We show that BM parametrized by solution size is $W[1]$-hard, already with a degenerate single matroid constraint. Thus, an exact parameterized algorithm is unlikely to exist, motivating the study of *FPT-approximation schemes* (FPAS). Our main result is an FPAS for BM (implying an FPAS for $\ell$-dimensional matching and budgeted $\ell$-matroid intersection), relying on the notion of representative set $-$ a small cardinality subset of elements which preserves the optimum up to a small factor. We also give a lower bound on the minimum possible size of a representative set which can be computed in polynomial time.

## 3.8 Approximating Weighted Connectivity Augmentation below Factor 2

*Vera Traub (Universität Bonn, DE)*

**License** Creative Commons BY 4.0 International license
© Vera Traub
**Joint work of** Vera Traub, Rico Zenklusen
**Main reference** Vera Traub, Rico Zenklusen: "A (1.5+$\epsilon$)-Approximation Algorithm for Weighted Connectivity Augmentation", in Proc. of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023, pp. 1820–1833, ACM, 2023.
**URL** https://doi.org//10.1145/3564246.3585122

The Weighted Connectivity Augmentation Problem (WCAP) asks to increase the edge-connectivity of a graph in the cheapest possible way by adding edges from a given set. It is one of the most elementary network design problems for which no better-than-2 approximation algorithm has been known, whereas 2-approximations can be easily obtained through a variety of well-known techniques.

In this talk, I will discuss an approach showing that approximation factors below 2 are achievable for WCAP, ultimately leading to a $(1.5 + \epsilon)$-approximation algorithm. Our approach is based on a highly structured directed simplification of WCAP with planar optimal solutions. We show how one can successively improve solutions of this directed simplification by moving to mixed-solutions, consisting of both directed and undirected edges. These insights can be leveraged in local search and relative greedy strategies, inspired by recent advances on the Weighted Tree Augmentation Problem, to obtain a $(1.5 + \epsilon)$-approximation algorithm for WCAP.

# 4 Panel discussions

## 4.1 Panel Discussion

*Frances A. Rosamond (University of Bergen, NO), Amir Abboud (Weizmann Institute – Rehovot, IL), Michael R. Fellows (University of Bergen, NO), Venkatesan Guruswami (University of California – Berkeley, US), Bingkai Lin (Nanjing University, CN), and Saket Saurabh (The Institute of Mathematical Sciences – Chennai, IN)*

**License** Creative Commons BY 4.0 International license
© Frances A. Rosamond, Amir Abboud, Michael R. Fellows, Venkatesan Guruswami, Bingkai Lin, and Saket Saurabh

The aim of the panel was fostering exchange between the various research communities (eg, parameterized and approximation algorithms as well as hardness of approximation) by identifying (i) general key challenges (meta research questions in a conceptual or technical sense) that are important to advance and promote the field, (ii) suitable taxonomy that allows to classify the possible algorithmic results, (iii) advance systematic understanding (in contrast to ad-hoc results), (iv) and concrete open research questions.

In particular the panelists discussed (under the moderation of Frances Rosamond) the following questions.

- In which direction would you like to see the field grow?
- What is a distinctive technical challenge in parameterized approximation? By "distinctive challenge" we mean an (exciting) technical challenge that (may) require an approach that goes beyond combining techniques that were previously already used in parameterization or approximation separately. We believe that it is crucial for the field to gain momentum and attract researchers that there are such unique challenges.
- What is your favorite result in the field?

After the panelists discussed the three questions, the floor was opened to all the participants of the Dagstuhl Seminar to share their ideas.

# 5 Open problems

## 5.1 TSP with line-neighborhoods in $\mathbb{R}^3$

*Antonios Antoniadis (University of Twente, NL)*

**License** Creative Commons BY 4.0 International license
© Antonios Antoniadis

The traveling salesperson problem (TSP) with line neighborhoods given a set of $n$ lines in $\mathbb{R}^3$ one seeks a shortest tour (closed curve) $\mathcal{C}$ that visits each line. A line $L$ is visited by $\mathcal{C}$ if and only if $\mathcal{C} \cap L$ is non-empty. In [1] an $O(\log^3 n)$-approximation algorithm was presented that is based on a reduction from TSP with line neighborhoods to Group Steiner Tree (at the loss of a constant factor in the approximation ratio). The setting where the lines are parallel is equivalent to solving a classical Euclidean instance in $\mathbb{R}^2$ and thus the problem is NP-hard. It was also shown among other results in [2], that the problem is actually APX-hard and admits an $O(\log^2 n)$-approximation algorithm, albeit with a running time of $n^{O(\log \log n)}$. Given the large gap between the respective upper and lower bounds, the most important open question with respect to the problem is whether or not it admits a constant-approximation algorithm.

### References

1 Adrian Dumitrescu and Csaba D. Tóth. The traveling salesman problem for lines, balls, and planes. *ACM Trans. Algorithms*, 12(3):43:1–43:29, 2016.

2 Antonios Antoniadis, Sándor Kisfaludi-Bak, Bundit Laekhanukit, and Daniel Vaz. On the approximability of the traveling salesman problem with line neighborhoods. In Artur Czumaj and Qin Xin, editors, *18th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2022, June 27-29, 2022, Tórshavn, Faroe Islands*, volume 227 of *LIPIcs*, pages 10:1–10:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.


## 5.2 Parameterized approximability of clustering problems

*Vincent Cohen-Addad (Google Paris, FR)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Vincent Cohen-Addad

Given a set of points (clients) $C$ and a set of facilities $F$ in a metric space $(C \cup F, \text{dist})$, the classic $k$-Median problem asks to find a subset $S \subseteq F$ of size $k$ (the centers), such that the total distance $d(S) := \sum_{p \in C} \min_{c \in S} \text{dist}(p, c)$ of points in $C$ to the closest facility is minimized.

1. Find a $(1 + 2/e)$-approximation that runs in time $2^{O(k)}\text{poly}(n)$ (where $(C \cup F, \text{dist})$ is an arbitrary metric space).
2. Consider the continuous setting of $k$-median with $L_\infty$ metrics (i.e.: $C$ is a subset of $\mathbb{R}^d$ and $F$ is $\mathbb{R}^d$). Find a 2-Approximation that runs in $2^{O(k)}\text{poly}(nd)$ time.

## 5.3 Parameterized Approximation for the Santa Claus Problem

*Fabrizio Grandoni (SUPSI – Lugano, CH)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Fabrizio Grandoni

In the Santa Claus problem we are given a collection of presents and a collection of children. Each present i has a value $v_{ij} \geq 0$ for child $j$. The happiness of a child $j$ is the sum of the values of the presents that (s)he receives. Our goal is to assign the presents so as to maximize the minimum happiness of any child. Santa Claus turns out to be an extremely challenging problem in terms of approximation algorithms. The best known lower bound on the (polynomial-time) approximation ratio is 2, while the best known upper bound on the same ratio is polynomial in the number $n$ of items.

Given the difficulty of this problem, it makes sense to consider FPT approximation algorithms. One natural parameter is the number $k$ of children. I did ask if a constant approximation (or better) is possible in FPT time. Andreas Wiese made me notice that a parameterized approximation scheme (PAS) for this problem is implied by [1] (described in the form of an EPTAS for constant $k$)

An interesting open question that remains is to define alternative parameters that make sense in practice, and design FPT approximation algorithms with respect to them.


### References

1 Klaus Jansen and Marten Maack. An EPTAS for scheduling on unrelated machines of few different types. *Algorithmica*, 81(10):4134–4164, 2019.

## 5.4 FPT Inapproximability Results Beyond Gap-ETH

*Pasin Manurangsi (Google Thailand – Bangkok, TH)*

**License** Creative Commons BY 4.0 International license
© Pasin Manurangsi

While a number of parameterized problems are known to be hard to approximate under the Gap Exponential Time Hypothesis (Gap-ETH), certain results are still out of reach of the current techniques even under Gap-ETH. We discuss a few such questions, including:

1. Strong FPT Inapproximability of $k$-Set Cover. While it is known that approximating $k$-Set Cover to within $g(k)$-factor is W[1]-hard for any function $g$[6]. The situation is less clear when we allow dependency on $n$ (the number of elements in the universe) in the approximation factor; the greedy algorithm yields $O(\log n)$-approximation while the best known ETH-hardness result is only $\tilde{\Omega}(\log^{1/k} n)$[4]. Open Question: Can we improve this hardness to, say, $\Omega(\log^{0.99} n)$ under Gap-ETH?
2. Total FPT Inapproximability of **Exact** $k$-Set Cover. Exact $k$-Set Cover is a special case of the Set Cover problem where we are promised that there exists a set cover of size $k$ such that the subsets in the solution are all disjoint. (Note here that the output solution only needs to cover the space but needs not be disjoint.) This version of Set Cover is often useful in subsequent reductions, e.g. to Coding-theoretic and Lattice problems. While the hardness of Exact Set Cover is achieved for free in the NP-hardness of approximation reduction for Set Cover of Feige[3], the parameterized hardness reductions do **not** give the hardness of such a version. This is due to the so-called "projection" property in Label Cover, which does **not** hold in the parameterized version of Label Cover used in the FPT hardness regime; in fact, it is not hard to see that requiring such projection properties will lead to at most $2^k$ gap. To the best of our knowledge, the best FPT hardness of approximation of Exact Set Cover based on Gap-ETH is only $k^{1/2--o(1)}$, which follows from the result of Manurangsi and Dinur [2]. Open Question: Can we rule out all $g(k)$-approximation FPT algorithm for Exact $k$-Set Cover?
3. Total FPT Inapproximability of Densest $k$-Subgraph (with perfect completeness). The best known FPT hardness of approximation under Gap-ETH of Densest $k$-Subgraph has inapproximability factor of $k^{o(1)}$ (where $o(1)$ can be any function that converges to zero as $k \to \infty$)[1]. On the other hand, under a non-standard "Strongish Planted Clique Hypothesis", this factor can be improved to $o(k)$[5]. Open Question: Can we prove the same hardness under Gap-ETH?

### References


1. Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From gap-exponential time hypothesis to fixed parameter tractable inapproximability: Clique, dominating set, and more. *SIAM J. Comput.*, 49(4):772–810, 2020.
2. Irit Dinur and Pasin Manurangsi. Eth-hardness of approximating 2-csps and directed steiner network. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 36:1–36:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
3. Uriel Feige. A threshold of ln $n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
4. Bingkai Lin. A simple gap-producing reduction for the parameterized set cover problem. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019,*

*July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 81:1–81:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
5 Pasin Manurangsi, Aviad Rubinstein, and Tselil Schramm. The strongish planted clique hypothesis and its consequences. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPIcs*, pages 10:1–10:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
6 Karthik C. S., Bundit Laekhanukit, and Pasin Manurangsi. On the parameterized complexity of approximating dominating set. *J. ACM*, 66(5), aug 2019.

## 5.5 $(2-\varepsilon)$-approximation for the Capacitated Vehicle Routing Problem

*Hang Zhou (Ecole Polytechnique – Palaiseau, FR)*

**License** Creative Commons BY 4.0 International license
© Hang Zhou

In the capacitated vehicle routing problem, we are given a metric space with a vertex called depot and a set of vertices called terminals. The goal is to find a minimum length collection of tours starting and ending at the depot such that each tour visits at most $k$ terminals, and each terminal is visited by some tour. We consider this problem in the Euclidean plane. The best-to-date approximation ratio was $2+\epsilon$ using the iterated tour partitioning technique introduced by Haimovich and Rinnooy Kan. It is an open question whether there is a better-than-2 approximation.

## Participants

- Amir Abboud
Weizmann Institute – Rehovot, IL
- Antonios Antoniadis
University of Twente, NL
- Jarek Byrka
University of Wroclaw, PL
- Karthik C. S.
Rutgers University – New Brunswick, US
- Parinya Chalermsook
Aalto University, FI
- Vincent Cohen-Addad
Google – Paris, FR
- Klim Efremenko
Ben Gurion University – Beer Sheva, IL
- Andreas Emil Feldmann
University of Sheffield, GB
- Michael R. Fellows
University of Bergen, NO
- Ameet Gadekar
Aalto University, FI
- Fabrizio Grandoni
SUPSI – Lugano, CH
- Carla Groenland
Utrecht University, NL
- Venkatesan Guruswami
University of California – Berkeley, US
- Martin Herold
MPI für Informatik – Saarbrücken, DE
- Lawqueen Kanesh
Indian Institute of Techology – Jodhpur, IN
- Matthias Kaul
TU Hamburg, DE
- Madhumita Kundu
University of Bergen, NO
- Euiwoong Lee
University of Michigan – Ann Arbor, US
- Jason Li
University of California – Berkeley, US
- Bingkai Lin
Nanjing University, CN
- Pasin Manurangsi
Google Thailand – Bangkok, TH
- Dániel Marx
CISPA – Saarbrücken, DE
- Pranabendu Misra
Chennai Mathematical Institute, IN
- Tobias Mömke
Universität Augsburg, DE
- Danupon Nanongkai
MPI für Informatik – Saarbrücken, DE
- Marcin Pilipczuk
University of Warsaw, PL
- Nidhi Purohit
University of Bergen, NO
- Rajiv Raman
University Blaise Pascal – Aubiere, FR & IIIT Delhi – New Delhi, IN
- Frances A. Rosamond
University of Bergen, NO
- Saket Saurabh
The Institute of Mathematical Sciences – Chennai, IN
- Chris Schwiegelshohn
Aarhus University, DK
- Hadas Shachnai
Technion – Haifa, IL
- Roohani Sharma
MPI für Informatik – Saarbrücken, DE
- Krzysztof Sornat
SUPSI – Lugano, CH
- Joachim Spoerhase
University of Sheffield, GB
- Vera Traub
Universität Bonn, DE
- Daniel Vaz
ENS – Paris, FR
- Andreas Wiese
TU München, DE
- Meirav Zehavi
Ben Gurion University – Beer Sheva, IL
- Hang Zhou
Ecole Polytechnique – Palaiseau, FR

Report from Dagstuhl Seminar 23292

# SportsHCI

**Florian 'Floyd' Mueller*[1], Carine Lallemand*[2], Dennis Reidsma*[3], Elise van den Hoven*[4], and Maria F. Montoya†[5]**

1 **Exertion Games Lab, Monash University, Melbourne, AU.** floyd@exertiongameslab.org
2 **HCI Research group, University of Luxembourg, LU & Industrial Design Department, TU Eindhoven, NL.** c.e.lallemand@tue.nl
3 **Human Media Interaction, University of Twente, Enschede, NL.** d.reidsma@utwente.nl
4 **Interaction Design Discipline, Faculty of Engineering & IT, University of Technology Sydney, AU & Industrial Design Department, TU Eindhoven, NL.** elise.vandenhoven@uts.edu.au
5 **Monash University – Clayton, AU.** Maria.MontoyaVega@monash.edu

**Abstract**

This report presents the work developed by 22 researchers and academics from across the world gathered for a week in Schloss Dagstuhl, Saarland, Germany, to discuss the future of interactive systems designed to support sport and exercise activity, a field called Sports HCI. Firstly, we present the activities developed day by day, from attendee's presentations to concrete community actions. Secondly, we show in detail the talks presented by the attendees, the interactivity and demo sessions, the discussion sessions, and the implications of the discussed topics to the Sports HCI field. Finally, we present the Sports HCI design pathways that attendees proposed based on the daily activities developed throughout the seminar. Ultimately, we hope this report inspires and motivates other Dagstuhl seminar proposals interested in the exciting field of HCI.

**Seminar** July 16–21, 2023 – https://www.dagstuhl.de/23292
**2012 ACM Subject Classification** Human-centered computing → Interaction paradigms
**Keywords and phrases** SportsHCI, Embodiment, Wearables, Mobile Computing
**Digital Object Identifier** 10.4230/DagRep.13.7.108

## 1 Executive Summary

*Florian 'Floyd' Mueller*
*Carine Lallemand*
*Dennis Reidsma*
*Elise van den Hoven*

**License** Creative Commons BY 4.0 International license
© Florian 'Floyd' Mueller, Carine Lallemand, Dennis Reidsma, and Elise van den Hoven

In July 2023, a seminar took place in which 22 researchers and academics from across the world gathered for a week in Schloß Dagstuhl, Saarland, Germany, to discuss the future of the design of interactive systems to support sport and exercise activity, concerning the emerging field of "SportsHCI". The following report documents the seminar and the efforts of the participants to investigate the underlying gaps in developing interactive technologies for sports by Human-Computer Interaction (HCI) practitioners, as well as knowledge guiding the design of such technologies, and the challenges the field of SportsHCI faces moving forward.

* Editor / Organizer
† Editorial Assistant / Collector

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Interactive systems are increasingly used to support sports and physical exercise activities, not only for functional purposes, such as sensing achievements to determine personal records but also for experiential reasons, such as highlighting the joy of movement, aiming to enrich the sports experience. The result is a surge in human-computer interaction systems for sports, such as sports watches that not only optimize runners’ athletic training plans but also provide beat-adjusted music for entertainment, augmented reality systems in stadiums to support the audience experience, and AI-enhanced capture systems to allow parents to watch their kids play soccer from afar in real-time. We call the resulting field “SportsHCI”.

A growing body of sensors, devices, and systems support these SportsHCI applications, allowing people to engage with athletic performance in novel and interesting ways, extending the field of HCI beyond the mouse and keyboard paradigm. Unfortunately, what is missing, though, is a thorough understanding of how to design such SportsHCI systems in a systematic way. This is in part due to a lack of a theoretical framework that is concerned with the coming together of computing and the physically active human body in a way that incorporates and integrates all aspects of movement, performance, learning, and experience, and a lack of understanding of how this combination has implications for how we approach the design of SportsHCI. It seems reasonable to assume that efforts that aim to address this could help people to engage more in physical activity and make high performance more accessible, and maybe even reach previously unengaged people, motivating them to try physical activity, so that ultimately, more people profit from the many physical, social, and mental health benefits sport and exercise activity can provide.

This Dagstuhl Seminar built on prior work that highlighted that supporting sport and exercise activities with technology should not only be concerned with functional aspects, such as measuring and comparing athletic performance, but it should also consider the experiential aspects, in other words, for example, how do sports people feel about their athletic performance. Therefore, this seminar invited experts from across the world and with different backgrounds, who understand that successful sport and exercise activity-support technology needs to be designed also with this experiential perspective in mind, and seek to develop guidance on how to support it. This seminar invited experts from both industry and academia, including experts from sports science as well as HCI and design, to work towards several large goals, including the identification of grand challenges in SportsHCI, so that a coherent approach could be developed to help more people profit from the benefits of sports and physical exercise.

There were several main outcomes of the Dagstuhl Seminar. The most notable of these are: a draft Grand Challenges paper was outlined, which since then has been finalised and submitted to a major conference in the field; an outline was made for a large Marie Curie grant proposal that accommodates all participants in the seminar; a proposal and call for papers have been drafted for a special issue in one of the major journals in the field; and plans have been laid for setting up, over the course of the coming two years, a separate international conference dedicated to the topic of this seminar. Furthermore, additional papers have been identified that could be developed jointly on the basis of the work done in this seminar. For all these points, concrete community actions were identified, and participants volunteered to take the lead on those actions, thus ensuring followup on the ideas developed during the seminar. Indeed, since finalizing the seminar, several of these action points have already been further taken up.

## 2 Table of Contents

## 3 Organisation of the Seminar

The Seminar was organised in several phases to lead to the definition of Grand Challenges in the field of SportsHCI, as well as to articulate concrete community actions to bring the field forward after completion of the Seminar. First, participants introduced themselves and their work and identified initial challenges that they encountered in their work. These challenges were collected, extended through discussions, and roughly organised into candidate Grand Challenges for the field of SportsHCI. These Grand Challenges were elaborated further in break-out groups. To elaborate our shared vision on these challenges, they were then built upon to articulate novel design ideas in teams of four or five participants, followed by a session in which the teams built on the design ideas to articulate the potential for *scientific* impact for each of those ideas. This further clarified our thoughts on the Grand Challenges as well. In addition, demonstrations of novel SportsHCI systems by various participants as well as a hike in the environment were organised for inspiration and to trigger further exchange of ideas. Finally, the week ended with a session in which everyone worked on the report, on documentation of results, and on identifying (and volunteering for) follow-up actions to be taken after the Seminar. The remainder of this section describes these various activities in more detail in chronological order.

### 3.1 Preparations

Prior to the workshop, the organizers have asked the participants to prepare a pitch (Pecha Kucha format), including the following elements: an introduction of each participant (including their hobbies), a presentation of their relevant SportsHCI work, their expectations for the seminar, their recommended reading for other participants and their rationale behind that choice, and the challenges they encountered in their work in the SportsHCI field. The recommended readings were collected in a shared folder, and more were added throughout the week.

A spreadsheet "who has what to offer" was furthermore shared prior to the seminar, with a focus on sports and activities that could be done during the week. Participants brought equipment and offered joint sessions of the following activities: running, meditation, yoga, gymnastic rings, volleyball, table tennis.

### 3.2 Day 1: Introductions to the Seminar and Participants; Initial Challenges

The seminar took place in one of the Dagstuhl school's conference rooms. Beforehand, the organizers established several work group materials, such as whiteboards, flipover sheets, sticky notes, and markers, to allow participants to record their thoughts and opinions and share them in easy sight inside the room. To begin the seminar, an interactive icebreaker activity was first undertaken to foster quick connections among the participants and establish a playful and enthusiastic group atmosphere. The activity involved forming a circle, where each participant introduced themselves with their name and a corresponding bodily gesture. The rest of the participants would then repeat both the name and gesture. Each time a new participant joined in, all the previous introductions were reiterated in sequence. By the end of the activity, with 22 participants present, a total of 253 introductions were accumulated and everyone was thoroughly familiar with the names of all participants.

After the engaging introductory activity, organizer Florian ‘Floyd’ Mueller provided the opening remarks, which contextualized the seminar and set the basis for the discussion of designing for SportHCI. Floyd encouraged participants to add, while listening through the various introductions, their thoughts on “challenges and opportunities” of designing for SportsHCI to the available flipover sheets at any time during the session.

Each seminar participant then gave a prepared presentation introducing their research (see Figure 1). In order to preserve a spirit of open and spontaneous group ideation, discourse, and collaboration, we strove to avoid a conference-like format of dry, dense lengthy presentations and instead adopted for a “Pecha Kucha” inspired presentation format; a rapid-fire series of short six-minute, visually oriented (picture and video) presentations. In addition to showcasing their work and its relation to SportsHCI, the content of each presentation also included an articulation of what the presenter expected from the seminar, the challenges they encountered in their investigations, and what prior work the group should read and why. An abstract of each presentation can be found at the end of the present report. Finally, during the Pecha Kucha talks, challenges that presenters mentioned and challenges that attendees members associated through the presentation were added as sticky notes to flipover sheets. They were initially grouped according to five categories: technology, users, society, research, and design.


**Figure 1** Pecha Kucha sessions, where all participants introduced their research related to SportsHCI.

## 3.3 Day 2: Interactivity session and group work on the Grand Challenges

An interactivity-style session involved participants trying out each others’ SportsHCI systems and artefacts. These demonstrations are documented in Section 5.

The remainder of the day was spent working further on elaborating Grand Challenges for the field of SportsHCI. This was prepared beforehand by organiser Dennis Reidsma using a shared information board in the Miro platform. All challenges and opportunities identified in Day 1 were added to this information board as separate notes, and lightly grouped into themes prior to the session on Day 2. In a plenary session on Day 2, participants associated additional themes with these notes. A discussion on what constitutes a “grand challenge” was held, and some inclusion and exclusion criteria were defined. Participants were then invited to select a few key challenges for the next stage of the activity.

Then, in breakout groups in two rounds (groups were reshuffled in round 2), challenges were further worked out on the Miro board as the next step towards getting at our grand challenges. After each round, a plenary presentation was given by each group, where they briefly described the current understanding of the high-level description of the challenge, and mentioned several core facets to the challenge. This process is further described in Section 6.

## 3.4 Day 3: Concrete Community Actions; Hike

In the morning of the third day, organiser Elise van den Hoven facilitated a strategic discussion on ways to bring the field forward. The focus was now not on the content and research problems but rather on organisational things that we as a field can pursue to consolidate the field of sportsHCI (special issues, workshops or special venues, etc). Participants first wrote down individual ideas for follow-up actions, which were then collectively sorted and clustered on the digital Miro board. This was followed by a plenary discussion to clean up the categories and clarify the ideas. Subsequently, participants split up into subgroups, selecting the most important action points. Participants were then invited to take concrete action by writing text for five key action points (e.g., draft the call for a journal special issue). These outcomes were documented as separate documents, and shared between participants for future action.

In the afternoon, participants joined the traditional Saarschleife Tafeltour hike in the region. Organizers suggested taking the opportunity of walking together to share career mentoring advice.

**Figure 2** Participants to the SportsHCI seminar hiking in the region.

## 3.5 Day 4: Novel SportsHCI Designs and Potential Scientific Impact

In the morning of the fourth day, organizer Carine Lallemand invited the participants to participate in a design challenge based on the paper "10 lenses to design Sports-HCI" [1]. The goal was to come up with design concepts related to one of the lenses, and to present them in the form of a role playing activity. Participants were randomly assigned to a group of 4 and randomly assigned a lens from the 10 lenses of the paper. After reading the 2-3 pages on that lens, each group could choose their own method to come up with ideas for designs; then they picked one and acted it out to the other groups. A quick debriefing was done with the groups after the roleplay. This was all documented in text (see Section 7) as well as in videos shared in the shared repository maintained by the organisers. This exercise was furthermore the inspiration for one of the outcomes of the seminar (plans for a follow-up paper on the lenses).

In the afternoon, organiser Dennis Reidsma presented slides to introduce the idea of an impact pathway, along with relevant terminologies and resources. These pathways are particularly interesting for funding proposals, as they are the basis for many funding agencies' expectations regarding articulation of expected scientific impact. The groups (same as in the morning) proceeded to write an impact pathway related to their design lens and concept, in the standard pattern used on the slides. These impact pathway descriptions are presented in Section 7.

In the evening, participants were invited by participant Florian Daiber to join a visit to the DFKI lab. Half of the group travelled to Saarbrücken to experience the infinity climbing demo which is included in the overview of the interactivity session in Section 5.

## 3.6 Day 5: Wrap-up, Conclusive Remarks, and Concrete Followup Actions

The last morning of the seminar was devoted to wrapping up all the ideas and insights collected during the week. Based on a report outline prepared by Carine Lallemand and a list of tasks, the participants produced extensive notes and text fragments for all chapters of this report. They also peer-reviewed the concrete text results from the session of Day 3.

Besides the report, it was of utmost importance to define and divide the responsibilities for the follow-up actions. For all concrete "next step actions", and based on personal interests and preferences, the group assigned a main person as the "shaper / lead" and listed all the participants interested to commit to work on the task.

### References

1 Florian Mueller, Damon Young, et al. 10 lenses to design sports-hci. *Foundations and Trends® in Human–Computer Interaction*, 12(3):172–237, 2018.

# 4 Overview of Talks

## 4.1 Embodied interaction

*Elise van den Hoven (University of Technology Sydney – Sydney, AU & Eindhoven University of Technology, NL, elise.vandenhoven@uts.edu.au)*

**License** Creative Commons BY 4.0 International license
© Elise van den Hoven

Within my international research program Materialising Memories (MM)[1] we study embodied interaction in its different meanings. This includes physical technology on or around human bodies as well as giving tangible artefacts agency (according to a definition by Dourish [1]). The typical MM application aims to support remembering practices in the broadest possible sense, including sharing holiday memories, reflection about everyday life, grieving a lost loved one, curating digital photo collections to understanding forgetfulness in older adults, to name a few. One of the projects focuses on motor memory, which is the memory of the muscles in the human body and relevant to SportsHCI. We are currently investigating whether HCI and interaction design can use this as a resource for designing technology, and how people can benefit from technology designed by taking motor memory into account.

### References

1 Paul Dourish. *Where the action is: the foundations of embodied interaction.* MIT press, 2001.

[1] `https://www.materialisingmemories.com/`

## 4.2 Designerly approaches to exercising motivation and injury prevention

*Carine Lallemand (University of Luxembourg, LU & Eindhoven University of Technology, NL, carine.lallemand@uni.lu)*

**License** Creative Commons BY 4.0 International license
© Carine Lallemand

As the leader of an educational community focused on challenge-based learning for sports, wellbeing and preventive health (Vitality Squad, TU Eindhoven), I have been supervising or leading a large variety of projects related to SportsHCI. Being part of a department of industrial design, our vision revolved around designerly ways to trigger healthy lifestyles with solutions contributing to both physical and mental vitality. Relying on aesthetics of interaction principles and the use of data as a creative material, we design interactive artefacts to support motivation and change. My work, done in collaboration with several Ph.D. candidates, covers several application areas and target audiences, with a main focus on exercising motivation (Daphne Menheere), sports injury prevention (Juan Restrepo Villamizar), office vitality through active ways of working (Ida Damen, Roy van den Heuvel), and the design of (inter)active urban environments (Loes van Renswouw [1, 2, 3]). We see an important role for data in the design process, to gain insight into these complex new behavior patterns in everyday life.

In my talk, I presented several projects to illustrate our main design approaches. Based on the Runners' journey [4, 5], projects were developed to trigger exercising motivation using qualitative interfaces [6, 7], data physicalization [8], and aesthetics of friction [9, 10]. We approached injury prevention by on-skin interfaces [11] or the concept of interaction-through-negotiation [12]. I concluded my talk with two challenges for SportsHCI: 1/ the necessity to further develop non-normative views on sports and motivation and to address audiences who are not already intrinsically motivated by physical activity. 2/ the responsible and meaningful use of data.

### References

**1** Loes van Renswouw, Steven Vos, Pieter van Wesemael, and Carine Lallemand. Exploring the design space of interactive urban environments: Triggering physical activity through embedded technology. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, DIS '21, page 955–969, New York, NY, USA, 2021. Association for Computing Machinery.

**2** Loes van Renswouw, Jelle Neerhof, Steven Vos, Pieter van Wesemael, and Carine Lallemand. Sensation: Sonifying the urban running experience. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI EA '21, New York, NY, USA, 2021. Association for Computing Machinery.

**3** Pieter van Wesemael Loes van Renswouw, Carine Lallemand and Steven Vos. Creating active urban environments: insights from expert interviews. *Cities & Health*, 7(3):463–479, 2023.

**4** Daphne Menheere, Carine Lallemand, Erik Van Der Spek, Carl Megens, Andrew Vande Moere, Mathias Funk, and Steven Vos. The runner's journey: Identifying design opportunities for running motivation technology. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, pages 1–14, 2020.

**5** Daphne Menheere, Mark Janssen, Mathias Funk, Erik van der Spek, Carine Lallemand, and Steven Vos. Runner's perceptions of reasons to quit running: Influence of gender, age and running-related characteristics. *International Journal of Environmental Research and Public Health*, 17(17):6046, August 2020.

6 Dan Lockton. Designing qualitative interfaces: Experiences from studio education. In *DRS2022: Bilbao*, DRS2022. Design Research Society, June 2022.
7 Daphne Menheere, Carine Lallemand, Ilse Faber, Jesse Pepping, Bram Monkel, Stella Xu, and Steven Vos. Graceful interactions and social support as motivational design strategies to encourage women in exercising. In *Proceedings of the Halfway to the Future Symposium 2019*, HTTF 2019, New York, NY, USA, 2019. Association for Computing Machinery.
8 Daphne Menheere, Evianne van Hartingsveldt, Mads Birkebæk, Steven Vos, and Carine Lallemand. Laina: Dynamic data physicalization for slow exercising feedback. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, DIS '21, page 1015–1030, New York, NY, USA, 2021. Association for Computing Machinery.
9 Alynne de Haan, Daphne Menheere, Steven Vos, and Carine Lallemand. Aesthetic of friction for exercising motivation: A prototyping journey. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, DIS '21, page 1056–1067, New York, NY, USA, 2021. Association for Computing Machinery.
10 Daphne Menheere, Alynne de Haan, Steven Vos, and Carine Lallemand. Raya: A tangible exercise buddy reminding oneself of the commitment to exercise. In *Human-Computer Interaction – INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30 – September 3, 2021, Proceedings, Part V*, page 471–475, Berlin, Heidelberg, 2021. Springer-Verlag.
11 Juan Restrepo-Villamizar, Steven Vos, Evert Verhagen, and Carine Lallemand. Crafting on-skin interfaces: An embodied prototyping journey. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, DIS '21, page 1129–1142, New York, NY, USA, 2021. Association for Computing Machinery.
12 Juan Restrepo. Hyaku: A qualitative negotiation-through-interaction interface to support runners in achieving balanced training sessions. In *DRS2022: Bilbao*, DRS2022. Design Research Society, June 2022.


## 4.3 Research focus in context

*Dennis Reidsma (University of Twente – Enschede, NL, d.reidsma@utwente.nl)*

**License** Creative Commons BY 4.0 International license
© Dennis Reidsma

My work in Sports Interaction Technology is shaped around Things that Interact for a Purpose against a background of Theory and Domain Knowledge. Each of these may be the focus of my research in some projects, but may as easily serve as the mere context for other projects. The "things" are typically digital technologies, often with expressive qualities. Holographic virtual reality for telepresence; smart objects with sensors and lights on a sports field; game-like settings in immersive virtual reality; robots and embodied agents; music technology; and many more. When the "thing" is the central concern, research questions involve finding good architectures, algorithms, and models, or working on novel creative, and playful designs. Alternatively, the "thing" may just be something used as a vehicle to explore other questions. The "interaction" in my work involves the sense-think-act cycle that is always embedded in the technology that I develop or use. Research questions at this level are about modifying elements in that cycle (what is sensed, what form do responses take, and how are the two mapped in an unfolding interactive dialog with the technology) and asking how this changes the immediate social and physical behaviour and experience of users. The larger "purpose" of my work tends to involve contributions to play, care, and learning –

sports can be said to be all three of those. Research questions that focus on the purpose mostly concern longer-term impact: does the technology indeed change social relations, health, and well-being, learning, etcetera? These questions are typically very intensive to answer and thus less frequently put in the center of my own work. The "theory and domain knowledge" about sports, play, and movement, finally, is present in two ways. On the one hand, it serves as a necessary foundation to inform the other three pillars. On the other hand, sports interaction technology can be used as a lab setup to explore and answer more fundamental questions that contribute to sports and movement theory.

## 4.4 From gymnastics and carpentry to critical sports interaction design

*Lars Elbæk (University of Southern Denmark – Odense M, DK, lelbaek@health.sdu.dk)*

**License** Creative Commons BY 4.0 International license
© Lars Elbæk

As a former elite gymnast, it paved my way into sports science, and I challenged myself to re-train the backward somersault to show off on my 60-year birthday. Despite being a trained carpenter with dyslexia, I pursued building houses, incorporating the double diamond as a window design of my house. With philosophy and a sports-physiology master's, I entered interaction design, creating SportsPlanner without knowing the discipline. This journey led to establishing the first Sports Innovation and Entrepreneurship program in sports science education, encompassing technology. Research has circled around developing games for an interactive back-yard trampoline, making a platform for filming and using visual data to strengthen learning in school physical education. As these two projects we have also in the Kids n' Tweens project developed the iMO-LEARN interactive stool for move-and-learn. In the field of physical activity for special needs we have co-designed the Handi-Wall with the local company PlayAlive inc. We have developed the "Real Man" platform using gamification for social connecting and enhancing behavioral change among blue workers. We have co-designed the feedback elements of the Eye4Talent football coaching tool. Considering health crises and talent development, technology presents ethical dilemmas. We see reality stars being online coaches having no knowledge and compassion and thus coaching towards fertilised eating disorders. It could also be that talent is produced from the many, but too much data may exclude and restrict motivational fellowship, reducing the talent pool. How will AI affect and on which values should we build ethical future HCIsport?

## 4.5 Designing Sports Interaction Technology to create meaningful movement for all

*Dees Postma (University of Twente – Enschede, NL, d.b.w.postma@utwente.nl)*

**License** Creative Commons BY 4.0 International license
© Dees Postma

The field of SportsHCI holds great potential to contribute to a more active and sportful society. Ubiquitous computing, wearable sensors, and artificial intelligence allow for on-the-fly, real-time interactions to support motor learning, skill acquisition, performance, and long-term engagement in sports. To live up to these promises, the field of SportsHCI needs to address a number of challenges to be able to tackle the big societal issues that we face today. The first of these challenges is to join forces within the HCI community and beyond to get a multidisciplinary perspective on wicked problems like: physical inactivity, decreasing physical activity, and impoverished physical literacy. Together, we need to rethink sensing techniques, processing approaches, and actuation paradigms to match the capabilities of interactive technologies to the needs of sports and movement.

## 4.6 Complexity simplified? Understanding and Designing for the Subjective experiences of Objective Sports Measures

*Armağan Karahanoğlu (University of Twente – Enschede, NL, a.karahanoglu@utwente.nl)*

**License** Creative Commons BY 4.0 International license
© Armağan Karahanoğlu

Improving the data accuracy and performance metrics is one of the driving forces of Sports-HCI. Research shows that the introduction of performance and biomechanical measurements altered the act and practice of sports. However, experience of the data collection technology (e.g., sports trackers) changes the way the athletes experience and immerse themselves in sports. I argue that athlete's subjective experience of objective measures is richer and more complex than their looking into performance and biomechanical data. For example, athletes can misinterpret data or may feel frustrated when the meaning derived from sports data does not match with their expectations or feeling of their performance. Furthermore, such mismatch and misinterpretation may negatively influence athletes' perceptions of self, such as self-worth and self-care. Hence, I challenge the sports-HCI with going beyond making data readable, understandable and interpretable, and investigating the practices of when the numbers and the data representation are not aligned with what an athlete feels about their body and performance. In my work, I investigate how sports technology influences athletes' experience and sensemaking of sports-data, their sporting and bodily experiences. I presented an initial framework that can shed light on investigating athlete's data sensemaking practices and concluded with the grand challenge of further investigating the "subjective experiences of objective measures".

## 4.7 SportsHCI: where interactive technologies inspire new uses and users

*Robby van Delden (University of Twente – Enschede, NL, r.w.vandelden@utwente.nl)*

**License** Creative Commons BY 4.0 International license
© Robby van Delden

SportsHCI provides an additional opportunity to engage in sports. Our SportsHCI projects include virtual rowing, tag, skiing and volleyball. These projects showed how VR and projections might visualise elements of our bodily behaviour and might steer us to behave in different ways. It also made us rethink what sports is and what SportsHCI could be. This familiarised us with the notion of sportification amidst a un(-)ification triangle of playification, gamification, and sportification. Adding or removing sportslike elements might be one starting point of what SportsHCI technologies could add in, SportsHCI and beyond. We will continue to explore where technologies can really be of added value and will provide new types of motivations. I see interactions can be aimed at getting people to actively move to exhaustion and excellence, but also as technology which could unlock feelings of awe and enjoyment of moving physically just a little for entrance level athletes. Furthermore, new and not yet existing mixed reality sports, such as GPT4's teamsport SpectraBall where rugby, beatsaber and football are combined, not only provide new experiences for athletes but also open up untapped potential for relatedness through the act of jointly (watching) sports.

## 4.8 Digital Breathing Coach – exploring explicit interaction on the run

*Vincent van Rheden (Salzburg Universit, AT, vincent.vanrheden@plus.ac.at)*

License Ⓒ Creative Commons BY 4.0 International license
© Vincent van Rheden

In this research project I developed a simple instance of breath coach that supports runners with breathing. Specifically we support Locomotor-Respiratory Coupling (LRC), also known as Rhythmic Breathing. Locomotor-Respiratory Coupling (LRC), a breathing technique in which the breath is coupled to steps, can positively impact running experience and economy. Locomotor-Respiratory Coupling (LRC), typically is done in whole-integer ratios, counted in steps per inhalation and exhalation. For example, one could use a 2:2 ratio, implying that the runner performs two steps per inhalation and two steps per exhalation: An Android application picks up the runner's steps and guides the breath through sound feedback. Based on a proof of concept a full fledged Android application is designed and developed to support longitudinal studies. Follow-up steps will explore feeding the system with breath data to further support the runner at individual level when needed. Additionally we aim to explore explicit interactions during the running activity, exploring interactions that do not interfere with the running experience or the running motion.

## 4.9 Designing Superhero Movement Experiences

*Perttu Hämäläinen (Aalto University, FI, perttu.hamalainen@aalto.fi)*

License Ⓒ Creative Commons BY 4.0 International license
© Perttu Hämäläinen

Human-computer interaction has potential to impact physical activity through two key ways: Making it easier to learn movement skills and improving the motivation to keep learning and moving. In my work, I have been mostly focusing on motivation, although my group has also done some learning interventions. I mostly build on intrinsic motivation theories such as Self-Determination Theory and the work of Silvia et al. [1] on curiosity. These highlight the importance of satisfying basic psychological needs such as competence, autonomy, social relatedness, and novelty. In my talk, I present examples of a stream of research originating in my own doctoral research, almost 20 years back: Creating "superhero movement experiences" with both physical and digital manipulations that aim to make one feel more competent as a mover. For instance, one might become the protagonist of a kung-fu game where one's body appears stronger, faster, and more flexible, and where one can execute backflips so that one only does the movement up to a point that is safe, and the game then takes over and guides the virtual body through the rest of the skill. I will also briefly comment on supporting other needs, such as curious exploration of novel movement opportunities in an urban environment, approached through an AI system that recognizes parkour spots and creates spot maps based on Google Street View images.

**References**

1 Paul J Silvia. Curiosity. *The science of interest*, pages 97–107, 2017.

## 4.10 Digital Motion in Sports, Fitness, and Wellbeing

*Lisa Anneke Burr (Salzburg University, AT, lisaanneke.burr@plus.ac.at)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Lisa Anneke Burr

The project "Digital Motion in Sports, Fitness, and Wellbeing" (DiMo) is motivated by the belief that by digitizing the quality of movement and making emotion visible during movement, we can support people in their movement and lead them to optimal performance and motion experiences. The project aims to give people a better understanding of their physical activity by linking motion and emotion, resulting in an overall new (digital) experience. Within this project, I worked on two streams: 1. exploring respiration as an interaction modality while in motion, and 2. working with visual respiration data feedback while running. Here, my focus lay on supporting breath pacing in treadmill runners while also trying to address impacts on breathing awareness. In the seminar, I am really interested in conversations on how technology can be used to build or increase body awareness to support training, recovery, as well as pleasure and enjoyment during sports. I am also keen to dive into discussions about ethical questions that arise in the field of emerging technologies in sports.

## 4.11 Using Human Augmentations to Learn and Enhance Sports Skills

*Don Samitha Elvitigala (Monash University – Melbourne, AU, don.elvitigala@monash.edu)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Don Samitha Elvitigala

Different skills need to be acquired in sports to perform well. Also, following the correct technique will not only enhance performance but also will avoid unnecessary injuries. Hence, my research has focused on integrating sensing and feedback into our bodies for skill learning and enhanced performance. For example, GymSoles allow our body to feel where our centre of pressure is to keep the correct body posture while lifting weight. The interface is integrated with our body taking foot pressure as input and creating vibrio tactile patterns to internalise our centre of pressure allowing the body to keep the correct posture. Similarly, in CricketCoach, we explored how feedback from SMA wires could be used in a cricket training session to render the hand feedback of a cricket coach, allowing a more naturalistic experience in self-cricket learning. In the future, I want to explore how body integrations can automatically be adjusted according to the contextual changes of the player's environment.

## 4.12 Multisensory Wearable Technology to Enhance and Transform Body Perception

*Laia Turmo Vidal (Universidad Carlos III de Madrid, ES, laia.turmo@uc3m.es)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Laia Turmo Vidal

Attuning to, cultivating, and enhancing body perception is crucial in sports. Advances in sensing, actuating, and wireless technology offer new ways to understand body physiology and movements. Further, wearable technology facilitates integration in challenging sporting contexts. My work in SportsHCI focuses on designing and evaluating sensory technologies to enhance and transform body perception. I presented two aspects of my work. First, Intercorporeal Biofeedback, a type of technology that enhances individual and mutual body perception, and that has been proven to improve understanding, communication, performance, and movement learning in yoga, strength training and circus training. Secondly, I presented Body Perception Transformation technologies, which utilize sensory feedback to create illusions of body changes, like feeling faster, lighter, or more fluid. Evaluations in everyday physical activity contexts and with dancers have shown that these technologies yield empowering, creative and motivating experiences for individuals. I concluded presenting current SportsHCI challenges, and perspectives I wish to discuss.

## 4.13 Athlete Experience in Figure Skating and Outdoor Recreation

*Michael Jones (Brigham Young University – Provo, UT, USA, jones@cs.byu.edu)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Michael Jones

Interactive computing is reshaping the athlete experience in sports and in outdoor recreation more broadly defined. This talk will discuss these issues in the context of figure skating and day hiking. In figure skating, the primary issues revolve around injecting data into the coach-athlete relationship where the athlete is a child and the parent is also involved. In the context of day hiking we present self-reported motivations for headphone use and non-use in the United States. These motivations involve safety, mediating social interaction and creating specific experiences. Before the seminar, I uploaded an essay and a research paper for pre-seminar reading. The essay frames cars in nature and foreshadows modern discussions of smartphones and nature. The research paper is an inductive study of sources of enjoyment in figure skating for elite skaters published in 1989. It focuses on enjoyment–an important concept in sports at all levels. My tentative pre-seminar grand challenge is: redefine SportsHCI to enhance the wellness benefits of lifelong participation in sports.

## 4.14 Data visualization in sports

*Paolo Buono (University of Bari, IT, paolo.buono@uniba.it)*

**License** Creative Commons BY 4.0 International license
© Paolo Buono

Data visualization is an established field of HCI. One of the main goals is to help users to understand and analyze data through visualization techniques and interaction. Visualizations have the property of being processed in parallel by the human perceptive system. Good visualizations should allow the observer to quickly understand the data and make decisions. Classic approaches, such as Shneiderman's Mantra, cannot be implemented anymore due to the massive amount of data. AI might help, but in many cases, the role of the user is not clear. Sports data are multivariate and are produced massively, making them challenging to compute and visualize. We aim to investigate the feasibility and utility of data visualization in sports, analyzing first soccer data visualization and interactive interfaces for video analysis.

## 4.15 Towards the Design of Playful water activities

*Maria Fernanda Montoya Vega (Monash University – Melbourne, AU, maria.montoyavega@monash.edu)*

**License** Creative Commons BY 4.0 International license
© Maria Fernanda Montoya Vega

Water's pleasant nature and associated health benefits have captivated the interest of HCI researchers. Prior WaterHCI work mainly focused on advancing instrumental applications, such as improving swimming performance, and less on designing systems that support interacting with technology in water in more playful contexts. In this regard, I explore the somaesthetic design of playful interactive water experiences, specifically, floating, surfing and diving, aiming to enrich the experiential aspect of the water activity through technology. Employing somaesthetic design, I aim to develop different playful prototypes to understand the integration of the soma (human mind and body), the technology and the water, towards a synergy that enriches the water experience. To date, I have designed an extended reality system for a flotation tank experience, which suggests interesting opportunities for the technological enrichment of water experiences. Ultimately, I hope that our WaterHCI work supports people to be playful and benefit from the many advantages of being in the water.

## 4.16 Blending micro-health behaviors into everyday activities through "invitations" as SportsHCI

*Xipei Ren (Beijing Institute of Technology, CN, x.ren@bit.edu.cn)*

**License** Creative Commons BY 4.0 International license
© Xipei Ren

Sitting in front of computers has become a major part of our workaday routines, challenging us in maintaining active and healthy lifestyles. This challenge becomes even more salient and more relatable to many of us after COVID-19. Is it possible to design interface and interaction so that the health behaviour intervention is integrated into the target users' established lifestyles and daily routines? We argue that by leveraging existing social mechanisms, a SportsHCI design can enable health interventions in a natural and effective way. One simple example could be facilitating a social context where users feel being invited for micro health gains. To better illustrate this design mechanism, we use several design instances (Step-by-Step, Anti-Sedentary Robot, Co-Drink, Co-Coffee) to exemplify our proposal. We would also like to learn your advice to inspire us in figuring out the design, evaluation and scaling-up methods of this persuasive strategy.

## 4.17 Experiencing Cycling Indoors and Outdoors

*Andrii Matviienko (KTH Royal Institute of Technology – Stockholm, SE, matviienko.andrii@gmail.com)*

**License** Creative Commons BY 4.0 International license
© Andrii Matviienko

Cycling is a great way to remain physically active, maintain cardiovascular health, and improve physical shape. While cycling outdoors is de facto a standard that includes cycling tours, training, or commuting, cycling indoors has become more popular over the last decade. In my research, I explore technological improvements for indoor and outdoor cycling. I will present how we employed multimodal and extended reality user interfaces and new tandem-based cycling simulation methods to facilitate cycling safety and realism. The results of these works demonstrate how technology can improve the cycling experience and reduce motion sickness indoors and facilitate safety outdoors.

## 4.18 AI+MR for augmented sport training systems

*Fabio Zambetta (RMIT University – Melbourne, AU, fabio.zambetta@rmit.edu.au)*

**License** Creative Commons BY 4.0 International license
© Fabio Zambetta

SportsHCI can benefit from the integration of MR (Mixed Reality) interfaces with AI (Artificial Intelligence) algorithms, particularly CV (computer vision) and ML (Machine Learning algorithms). We discuss the blend of such technologies in use cases related to training high performing athletes in table tennis, albeit ideas and principles can be extended to other racket sports and, in fact, to a variety of other sports, as well. We are also very interested in the interplay of such technologies with humans in the loop (athletes and coaches), specifically wrt their acceptance of such technologies, their perception of user experience and the opportunities for such technologies to be assisting athletes and para athletes.

## 4.19 Towards SportsHCI to improve skill acquisition and performance in athletes

*Florian Daiber (German Research Center for Artificial Intelligence (DFKI) – Saarbrücken, DE, florian.daiber@dfki.de)*

**License** Creative Commons BY 4.0 International license
© Florian Daiber

In the last years, sports technology has become ubiquitous and there has been a large body of work in HCI and Ubicomp as well as commercial products including apps, wearables and smart sports environments. SportsHCI has the potential to analyze complex human movements and provide guidance when learning a new motor skill. In our research we investigate intelligent assistants to support skill acquisition and improve performance in athletes. To achieve this goal we investigate different feedback techniques, both on the athlete's body and in their training environment. This includes smart environments using intelligent sports equipment and projections as well as wearables and virtual reality headsets, which enable us to provide in-situ feedback to enhance motor learning. We are particularly interested in implicit, in-situ and real-time feedback for example by using EMS but also playful approaches in mixed realities.

## 4.20 The Role of HCI in Enhancing Inclusivity and Encouraging Engagement in Sports and Exercise

*Daniel Harrison (Northumbria University – Newcastle-upon-Tyne, UK, daniel.b.p.harrison@northumbria.ac.uk)*

**License** Creative Commons BY 4.0 International license
© Daniel Harrison

Much of the historic focus of physical activity-related HCI has either been related to: those starting out their journey; or, to performance, particularly in (semi-) professional athletes. Over time this has broadened to include a range of sports, fitness levels, and technologies. However, more work remains to be done to establish a more inclusive SportsHCI research agenda that caters for all athletes, as well as promotes enjoyment and continued engagement regardless of gender, race, background, ability, body size, or time-pressures. While there is a broader movement in HCI around designing inclusive and accessible technologies, this should be more explicit in SportsHCI. Here I will bring examples from the cycling communities to highlight this need. I argue that a Feminist HCI framework can help emphasise inclusivity and diversity from an intersectional standpoint, allowing SportsHCI to better address individual needs and preferences, challenge stereotypes, and create a more welcoming environment for all.

### 4.21 SportsHCI benefits from integration

*Florian 'Floyd' Mueller (Monash University – Melbourne, AU, floyd@exertiongameslab.org)*

**License** Creative Commons BY 4.0 International license
© Florian 'Floyd' Mueller

SportsHCI benefits from the integration of the human body and interactive technology. We demonstrate this through a series of research design works around integrated cycling experiences, integrated entertainment experiences, and integrated arts experiences. The results of these works suggest interesting ways forward for SportsHCI research, in particular how the design of integrated SportsHCI can highlight experiential aspects, facilitating playful exertion experiences. Ultimately, with our work, we want to enhance our knowledge around the design of integrated SportsHCI experiences to help people understand who they are, who they want to become, and how to get there.

### 4.22 Digital Twins and SportsHCI: Design Challenges and Benefits

*Regina Bernhaupt (Eindhoven University of Technology (TU/e), NL, r.bernhaupt@tue.nl)*

**License** Creative Commons BY 4.0 International license
© Regina Bernhaupt

Data – qualitative or quantitative – has become central for Sports HCI. My personal take on data is from the perspective of digital twins. The goal of digital twins is to have a full representation of an artefact, process, machine or even a digital replica of a human. The identified key challenge from my perspective is that time, evolvement over time and especially how to enable people with the future. Digital twins have the ability to simulate and predict in real time – making it imperative for people to interact with future states and alternative futures. We know that humans are not very good when it comes to planning, predicting the future and knowing what to come is sometimes disturbing for people and how they react. It will be important to consider the ethics when designing with digital twins, to ensure that sports systems will be a positive contributor to humans health and their planning of sport activities and does not become a burden or a scary prediction instrument that is foreshading possible diseases or injuries. The key topic emerging from this challenge is to understand overall on how to interact with data and representations that relate to the future: "future interactions".

## 5 Overview of Interactivity Session

The second day of the seminar was introduced by organizer Dennis Reidsma, starting off with an "interactivity session" that involved live interactive demonstrations of systems and technologies relating to sports HCI. Numerous participants in the seminar showcased their prototypes and technologies, strategically arranged throughout the main room and the nearby hall. This layout allowed other attendees to freely explore the setups, test the systems, and engage in interactive discussions. The main goal was to provide a hands-on demonstration of interactive technologies, hoping to inspire participants to conceptualize their innovative interactive systems for different sports contexts.

Laia Turmo Vidal demonstrated “SoniBand” [1], a wearable device designed for real-time sonification of movement angles, through a range of movement-generated sounds. Embedded in a patch of fabric in a bracelet, SoniBand can be worn in various body locations (e.g. arm, leg, neck). Soniband integrates a BITalino R-IoT embedding a 9-axis Inertial Motion Unit (IMU). The R-IoT transmits movement angle data wirelessly to a Raspberry Pi Zero, which can be controlled using a web browser e.g., in a smartphone. The device registers the minimum and maximum angle of the body part (calibration), and then it sonifies the movement angle. SoniBand includes different metaphorical sonifications, such as wind, water, or rusty gears.

**Figure 3** SoniBand.

Michael Jones [2] presented a measurement system for figure skating jump detection, consisting of a wearable IMU and a smartphone. The wearable IMU was affixed to the participant’s lower back, which sampled motion at 120 Hz and sent all readings wirelessly to the smartphone. The smartphone ran a jump detection algorithm tuned to specific properties of figure skating jumps. A jump was detected when the airtime of the participant’s jump was sufficiently large and when the rotational velocity around the participant’s longitudinal axis surpassed a given threshold. In this demo, Mike recorded the participants’ results on a whiteboard to encourage competition among them. Participants achieving a jump off the ice created a greater appreciation for the difficulty of spinning jumps performed on ice. Particularly, participants rarely completed more than one rotation while in the air. In contrast, skaters in international competitions routinely perform 2.5 or more rotations in the air.

**Figure 4** Wearable IMU for figure skating jumps.

Perttu Hämäläinen demonstrated an anticipatory visualization for VR dancing. A key visuomotor control challenge in dancing is following the choreography or movements demonstrated by someone else. This is hard to do in real-time due to the inertia of the body and the delays inherent in human visuomotor control. The only way to perfectly follow a choreography in real-time is to have some anticipatory knowledge about it, either through memorization or additional cues such as a visualization of the upcoming movements. Perttu presented a novel way to do this for complex full-body contemporary dance movements, going beyond the highly simplified visualization techniques of current dancing and rhythm games that only specify parts of the movement such as footstep positions in space. Participants engaged in this demo by using the VR headset and following the choreography visualised in the VR environment.

Robby van Delden presented “FireFly Island”, an interactive social VR world created for a dyad-based exploration through which social intimacy can grow. This demo was developed on the basis of Savio Menifer’s MSc thesis[2] with further input of Medra, Joris Weijdom, and Dirk Heylen. In this demo, a set of interactive triggers facilitate moments of non-sexual intimacy, such as an area with proximity-sensitive witch hats and selfie mirrors that physically brings people closer together. The demonstrator identified how relatedness also in sports and through the medium of VR/XR can provide ample opportunities to deliberately provide epochs of reflection to stimulate relatedness beyond or mixed with more fast-paced physical activity-based interactions.

**Figure 5** Firefly Island by Savio Menifer.

Lars Elbæk presented the “MeCaMInD” method card toolbox [3]. He explained how researchers, teachers and potentially the seminar attendees could use this toolbox to access a design methods collection that will equip them with tangible tools to integrate movement into their design practice, transforming traditional creative methodologies with an innovative, movement-centred approach. The Method Cards for Movement-based Interaction Design (MeCaMInD) project is an ambitious initiative funded by the EU through the Erasmus+ strategic partnership program. [3] The project’s objective is to introduce movement as an integral component in designing new movement practices, artefacts, and interaction designs for sports and physical activities. The intent is to drive the development of more sustainable movement technologies and sports concepts, enhancing health and well-being for people.

**Figure 6** Method cards for movement-based interaction design.

Carine Lallemand presented the Vitality database, a collection of student and researchers projects on the topic of vitality realized at the Eindhoven University of Technology from 2017-2023. It addresses the challenge in design education or design research communities

[2] https://essay.utwente.nl/89341/
[3] https://mecamind.eu

to document and share past projects in order to learn from them and potentially valorize them through dissemination activities. In the Vitality database, each project is tagged with metadata to allow for easy sorting and filtering by topic, method, and material used. The idea behind the demo was to envision how such a database could be used by the SportsHCI community as a whole (rather than by a single institution). First, the design exemplars can inspire future products or be used for educational purposes so that students working on a SportsHCI project have prior design work to review and build on. These design exemplars can also be combined by topics to create papers such as annotated portfolios or explorations of design spaces. Other tabs can be created in the database for recommended readings and publications.

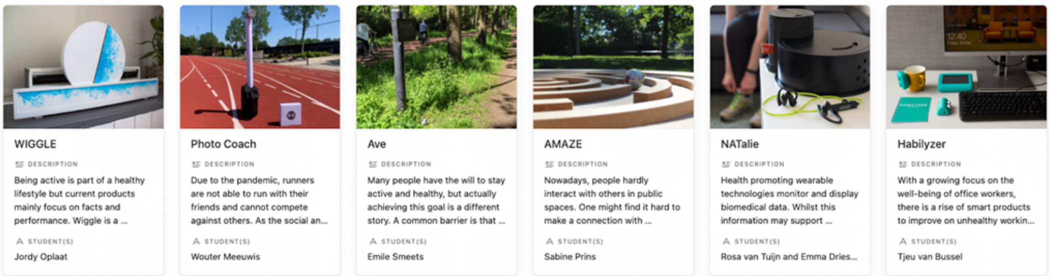


**Figure 7** A screenshot of student projects in the Vitality database.

Florian Daiber presented slackliner 2.0 [4], an interactive slackline training assistant which features head and skeleton tracking, and real-time feedback through life-size projection. Like in other sports, proper training leads to a faster increase of skill and lessens the risk of injuries. We chose a set of exercises from slackline literature and implemented an interactive trainer which guides the user through the exercises giving feedback if the exercises were executed correctly. The present demo showcases an interactive sports training system that provides in-situ feedback while following a well-guided learning procedure.

**Figure 8** Demonstration of Slackliner 2.0.

Vincent van Rheden presented the Digital Breathing Coach [5], which supports Locomotor-Respiratory Coupling (LRC), also known as Rhythmic Breathing. Locomotor-Respiratory Coupling (LRC), a breathing technique in which the breath is coupled to steps, can positively impact running experience and economy. Locomotor-Respiratory Coupling (LRC), typically is done in whole-integer ratios, counted in steps per inhalation and exhalation. For example, one could use a 2:2 ratio, implying that the runner performs two steps per inhalation and two steps per exhalation. BreathTool, an Android application picks up the runner's steps and guides the breath through sound feedback. Based on a proof of concept a full fledged Android application is designed and developed to support longitudinal studies.

**Figure 9** BreathTool: a Digital Breathing Coach.

Florian Daiber, finally, presented InfinitiWall [6], an integration of a Virtual Reality outdoor rock climbing experience coupled with a rock climbing treadmill which allows the virtual experience to extend infinitely in vertical space.

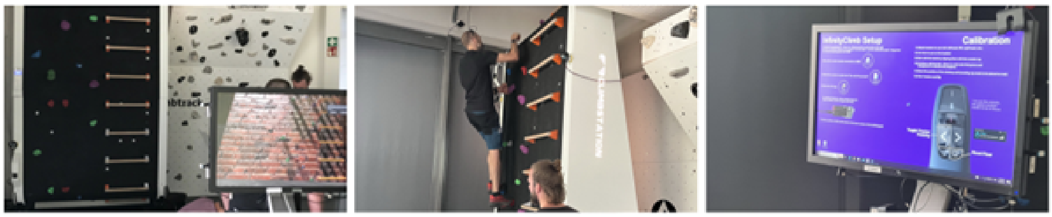

**Figure 10** InfinityWall: VR climbing on a vertical rock climbing treadmill.

## 6 Defining the Grand Challenges of Sports HCI

Through the Pecha Kucha presentations, a comprehensive list of “challenges” was collated by the time every participant had presented, providing a strong foundation for steering discussion during the remainder of the seminar toward topics that require further elaboration. Previous works have argued that HCI requires “grand challenges”, namely in that it provides a steering force to drive coordinated action in guiding future research, theory, design, and commercial development [9, 10, 8, 7]. Acknowledging this need, these challenges were further consolidated into a set of concrete “Grand Challenges for SportsHCI”, with the intention their articulation would give guidance to researchers wishing to contribute to the development of the theory by providing specific and actionable gaps in knowledge or capability to which they can contribute through future research.

To determine what constitutes a Grand Challenge in SportsHCI, some inclusion criteria were discussed:

- Is the challenge specific to SportsHCI or more salient in our field? If not, does it play out differently?
- Is the challenge important for the field and not easily solved?
- Is the challenge not addressed yet in current work?
- Can the challenge be solved within the next 10 years?

With the help of the Miro board (see Figure 11) and according to the abovementioned criteria, participants discussed a list of potential grand challenges. This included: reconciling performance and experience, the feeling of data (objective vs. subjective data), the temporal aspect in SportsHCI, sports data in a wider context (home, nature, city, work) from both a technology and an activity perspective, designing for political futures, athletic performance from the experiential perspective, the role of the audience, promoting physical literacy, addressing inequality by reaching people at a disadvantage, and developing a strategic vision for the field. Then, in breakout groups, all 22 participants were broken into four groups, with each group tasked with discussing one of the proposed grand challenges and taking notes on the Miro board. A second round was conducted with the reshuffling of the groups with a different challenge. After each round, a plenary presentation of highlights was given for each challenge addressed by a breakout group.

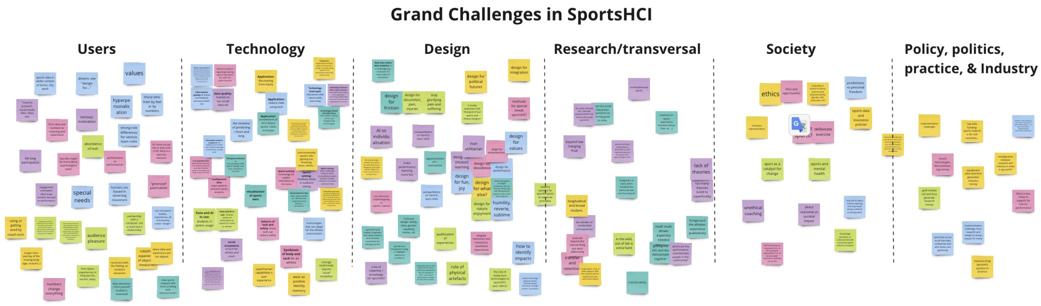


**Figure 11** Overview of initial challenges that were grouped and sorted into the preliminary Grand Challenges.

Following the completion of all participant presentations, a discussion took place in which the participants synthesized the concepts brought to light during the presentations, as well as the opportunities and challenges that were also highlighted.

The challenges were extensively documented in a separate initial draft for a Grand Challenges paper; here we only provide initial summaries of the Grand Challenges that were separately discussed during the Seminar.

## 6.1 Time in SportsHCI

Time emerges as a critical aspect in sports experiences, encompassing various phases that demand consideration. These phases range from the immediate timespan of the sports activity itself to preparatory and recovery phases. Additionally, long-term time periods, like seasons and training for major competitions, like the Olympics, hold significance. Aging also influences sports, prompting changes across different life phases. The influence of aging on sports prompts adaptations across different life phases. Consequently, designing with time as a central consideration becomes indispensable, involving careful attention to time span and data sampling for efficient monitoring and feedback. Moreover, despite its challenges, predictive methods and visualizations based on temporal data are crucial for positively manipulating athletes' time span, mitigating skill deterioration over time. Furthermore, understanding the body's changes over time in SportsHCI design poses underlying challenges, mainly due to the scarcity of longitudinal data. Addressing this requires implementing longitudinal methods for subjective evaluation, and once human interactions are successfully modeled, the question remains of how to effectively engage with these models and interact with predictions and simulated outcomes.

## 6.2 Strategic vision for the field of SportsHCI

Much of our discussion (across both sessions) moved around the fact that sports themselves are often not inclusive, and by promoting a generic message of “we should be more inclusive” we are not providing an actionable vision for the future. Beyond inclusivity (which we re-categorised as “access and participation” to more accurately reflect the conversations we were having in the sessions, we also discussed sustainability and how similarly this was too broad of a term to be useful in providing a useful strategic vision – this is something we better need to define and more explicitly break down in order to be useful (do we mean sustainable in human-costs, in fiscal costs, in environmentalism, in terms of access to facilities?). We also discussed that, by their very nature, conversations around the strategic vision for SportsHCI are political and how the political landscape across different places and sports will have a significant influence on our work, meaning that our strategic visions would need to be aware of this and sufficiently broad. Within the context of inequalities, we discussed how SportsHCI tech can exacerbate (or even be the root cause of) inequalities, because of design decisions, target audiences, and costs.

## 6.3 Performance and Experience

We discussed other ways to “enjoy” sports beyond improving athletic performance. We realised an experience as joy is very narrow and using “appreciation” of the experience describes better how to enclose the sport experience and provide a holistic view. In that sense we as sports researchers can support several ways to appreciate sports. We also discussed how our field can learn from other fields that moved from performance to the experience, such as game research, in which at the start they were focused on developing for competition and then they created more experiential games based on arts, for example. Performance, measurement, judgement, competition, winning and losing are core elements of sport. However, the unmeasured subjective athlete experience drives both performance and participation. Conversely, performance strongly impacts the athlete experience. Performance and experience are closely interrelated in SportsHCI. The grand challenge is to incorporate elements of both performance and experience in HCI for sports. Doing so may increase performance and participation while improving the athlete’s performance.

## 6.4 Subjective and Objective Data

We discussed the need to understand how we can provide meaningful data independently of being subjective or objective, and also not only meaningful for the athlete but also for the coach. We discussed how there is a two-way feedback loop where: (1) objective data is “translated” to actionable, subjective data (i.e., joint data is turned into coach advice to an athlete, such as “raise our right elbow”); (2) subjective data is “verified” (fact checked) using objective data (e.g., misconceptions about dancers’ lightness when dancing actually depends on the colour of their shoes). A lot of research work is still based on quantification and measurement. While there is certainly room for that, there are plenty of opportunities to capture people’s experience of sport (how they feel about it). This is important in Sports HCI, because for many people the feeling of being “immersed” in sports is key.

## 6.5 Engineering SportsHCI

SportsHCI introduces a number of very specific engineering challenges that may not be so salient/prominent or even present in some other HCI fields. Compared to the other application fields of HCI, SportsHCI often puts more emphasis on accuracy, timely measurement, embodied experiences, and high-level integration with bodily exertion. This requires higher fidelity levels and robustness of the prototype. Yet, there lacks toolkit and design methods to help designers easily engineer their out-of-box ideas into testable/experienceable probes, which sometimes also limit the novelty of SportsHCI research and work efficiency. Partly because of this, the prototypes for SportsHCI are not that easy to avoid disturbing people in use due to their physical obtrusiveness, making it challenging to create an ultra-realistic setup for testing proposed design concepts. Additionally, unlike other computing technologies that were designed for constant usage, many SportsHCI techs are designed for special sports, which makes them not being used very frequently. How should we deal with these occasionally used technologies? And how can we make sure such technologies can be evolved to fit the changes of users' needs? A systemic thinking might be needed to study and design SportsHCIs ecologically. Some specific challenges include:

**Complex data.** Data consists of large and complex time series, that require processing to turn into meaningful interpretations. For HCI applications often the detections and interpretations need to happen on the spot so the interactive application can immediately respond to it. At the same time, discovering meaningful interpretations sometimes needs a) too much processing power and b) to work on the full data (including "future" data). The challenge is to find ways to develop (possibly based on training stages) the AI that can deliver real-time immediate meaningful interpretations for the HCI to respond to the user‘s activity.

**Movement artefacts.** Not all, but many measuring systems are prone to "movement artefacts" (EEG, eye tracking, ...). This means a user must sit pretty still when working with that measurement system. Unlike in other HCI fields, SportsHCI is by definition high on movement, so if we want to use those kinds of measurement systems, we need to find fundamentally new ways to deal with / get around those movement artefacts.

**Robust prototyping platforms.** Existing prototyping platforms are often not robust enough for use in the SportsHCI context; there is a need for other, more robust prototyping toolkits tailored to both the setting in terms of sensing and actuation paradigms, and to the context in terms of robustness.

**Sensing sports movement.** Compared to some other movement-based contexts in HCI, sports see a wide range of movement speed and power; measurement systems require high resolution at low values but also enough room at the high end to deal with explosive actions (cf the figure skaters who jump so hard that the IMU is maxed out, in the demo of Mike Jones during this seminar).

**Robust and challenging environmental circumstances.** Sports sensing and actuation happen in challenging settings that require more robust technology compared to many other mobile/wearable/ubiquitous works. Water, sand, ice, bumping into rocks and mountains, collisions between athletes, etcetera. This robustness problem is exacerbated by the need for long-term studies – ideally, systems could survive year-long deployments, which for prototypes in more lenient circumstances is already challenging, but in sports settings even more so.

**Safety, sports regulations, and wearables.** Wearables may run into regulations forbidding athletes from wearing anything (including jewelry, etc.) for safety reasons; the engineering of the thing needs to prevent any safety risk (and regulations then must adapt). Wearables interfere with, and are interfered with by, athletic movement. The physical obtrusiveness of wearables is a problem for athletes' movement. Solutions may involve off-body sensing, smaller components integrated across the body rather than bulky wearables, different weight distribution of wearable, flat electronics that don't extrude, and other directions. The athletes' movement is also a problem for the physical wearable: straps are not good enough, and we need new attachments that don't shift and don't hinder.

**Per-sport specific things to measure and model.** In sports we often deal with very specific types and categories of movement that do not automatically generalize to other sports. A spin in ice skating is (data-wise) not the same as a spin in skateboarding or gymnastics. So where quite a bit of work on movement-centric HCI looks for generalizable models, we need approaches and methodologies that help us, in any sports context, to quickly get at those sport-specific movements.

## 6.6 Other Challenges

Other challenges that were discussed but not yet elaborated during the seminar, include the challenge of involving movement in the design of new practices and artefacts; the challenge of doing SportsHCI designs and studies that do not focus on novelty but instead offers the opportunity to triangulate already-known results (that is, go from "having a bunch of cases" to "having generalizable, well-grounded knowledge"); and the challenge of finding good methodologies to do forms of meta-analysis in this field.

In summary, the discussions of grand challenges during the seminar, of which an initial selection has been summarized above, has led to a deep insight of the community of seminar participants into where the future directions of the field lie; these insights will be written up in a manuscript for a major conference.

**References**


1 Judith Ley-Flores, Laia Turmo Vidal, Nadia Berthouze, Aneesha Singh, Frédéric Bevilacqua, and Ana Tajadura-Jiménez. Soniband: Understanding the effects of metaphorical movement sonifications on body perception and physical activity. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021.

2 Michael Jones, Sarah Ridge, Mia Caminita, Kirk E Bassett, and Dustin Bruening. Automatic classification of take-off type in figure skating jumps using a wearable sensor. 2022.

3 Lars Elbæk, Rasmus Vestergaard Andersen, Robby W Van Delden, José María Font Fernández, René Engelhardt Hansen, Perttu Hämäläinen, Mats Johnsson, Lærke Schjødt Rasmussen, Søren Lekbo, Solip Park, et al. Method cards for movement-based design. 2023.

4 Christian Murlowski, Florian Daiber, Felix Kosmalla, and Antonio Krüger. Slackliner 2.0: Real-time training assistance through life-size feedback. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–4, 2019.

5 Vincent Van Rheden, Eric Harbour, Thomas Finkenzeller, Lisa Anneke Burr, Alexander Meschtscherjakov, and Manfred Tscheligi. Run, beep, breathe: exploring the effects on adherence and user experience of 5 breathing instruction sounds while running. In *Proceedings of the 16th International Audio Mostly Conference*, pages 16–23, 2021.

6 Felix Kosmalla, Florian Daiber, and Antonio Krüger. Infinitywall–vertical locomotion in virtual reality using a rock climbing treadmill. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–6, 2022.

7 Eleonora Mencarini, Amon Rapp, Ashley Colley, Florian Daiber, Michael D Jones, Felix Kosmalla, Stephan Lukosch, Jasmin Niess, Evangelos Niforatos, Paweł W Woźniak, et al. New trends in hci and sports. In *Adjunct Publication of the 24th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 1–5, 2022.

8 Eleonora Mencarini, Amon Rapp, Lia Tirabeni, and Massimo Zancanaro. Designing wearable systems for sports: a review of trends and opportunities in human–computer interaction. *IEEE Transactions on Human-Machine Systems*, 49(4):314–325, 2019.

9 Florian Mueller, Rohit A Khot, Alan D Chatham, Sebastiaan Pijnappel, Cagdas" Chad" Toprak, and Joe Marshall. Hci with sports. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pages 2509–2512. 2013.

10 Stina Nylander, Jakob Tholander, Florian Mueller, and Joe Marshall. Hci and sports. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*, pages 115–118. 2014.


## 7 Design for SportsHCI: Design Lenses, Role Play, and Impact Pathways

On Thursday morning, Carine Lallemand conducted a design activity based on the paper "10 Lenses to Design Sports-HCI" [1]. Participants were randomly divided into 5 groups of 4, each group being randomly assigned one of the lenses. Each group was invited to read the section of the paper related to the lens they had been assigned and to brainstorm potential design concepts representative of that lens. The concepts were presented to the entire group in the form of role-play, along with insights into the design for this lens. The concepts produced were intended to be used as starting points to think of impact pathways in SportsHCI research.

On Thursday afternoon, Dennis Reidsma gave a presentation on how to articulate the scientific and societal contributions of our SportsHCI work, for the purpose of papers and grant proposals, in a way that outsiders can follow the line of argumentation.

We started off with a presentation on the Impact tool of Erasmus+ Partnerships [4] which provides a schematic to articulate the difference between output, outcome, and impact in terms of how much it is under direct control or influence from the project activities (see Figure 12). The same tool also provides extensive help in building up a storyline, starting from the impact you want to make, all the way down to the input and activities that are needed for this, and then going the same path back to show how the input and activities make it likely that the impact will indeed occur.

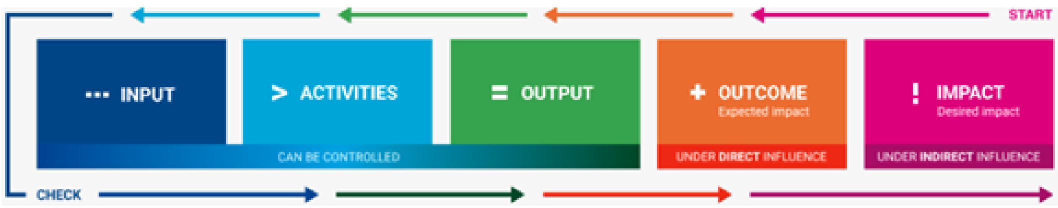


**Figure 12** Erasmus+ impact tool: how to build up the storyline about the impact that a project will have.

[4] https://www.erasmusplus.nl/en/impacttool-strategicpartnerships

Next, we discussed the impact pathways (see Figure 13) that are currently used to structure many National [5] and European [6] grant proposals. These build up a similar story as the Erasmus+ tool but are a bit more articulate about how impact, outcome, and necessary output are tied together by assumptions about underlying causes for problems and knowledge gaps that prevent us from resolving the underlying cause.

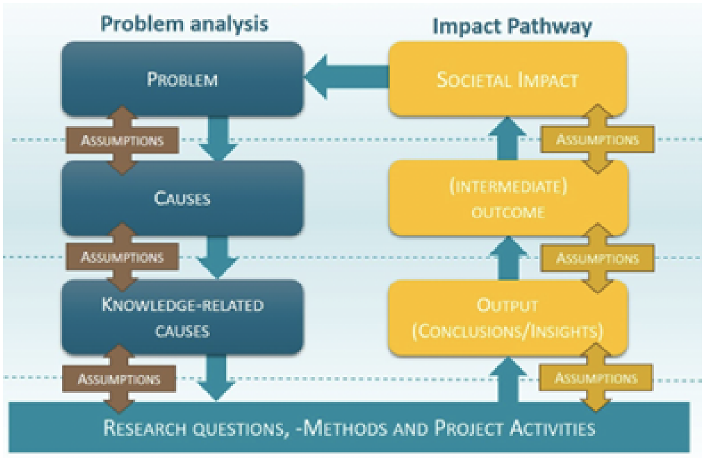


**Figure 13** Impact Pathways used for Dutch NWO grant proposals, with a central role of "assumptions".

Based on the concepts presented above, all groups worked on articulating an "impact pathway story" about one of their design concepts from the morning session, as reported in the remainder of this chapter.

The remainder of this section presents both the design concepts and their corresponding impact pathways.

### References

**1** Florian Mueller, Damon Young, et al. 10 lenses to design sports-hci. *Foundations and Trends® in Human–Computer Interaction*, 12(3):172–237, 2018.

## 7.1 SportsHCI lens: Reverie

*Laia Turmo Vidal, Carine Lallemand, Lars Elbæk, Daniel Harrison*

**License** Ⓒ Creative Commons BY 4.0 International license
© Laia Turmo Vidal, Carine Lallemand, Lars Elbæk, Daniel Harrison

After having read about the Reverie lens from the paper, the group summarized what they understood about the concept and the three subfacets (Void, Solitude, and Mini-Holidays). We highlighted in the text a few elements that could act as design triggers to create reverie (e.g., removing stimulation in the environment or adding extra stimulation). Then, each participant shared their own experiences of reverie, in a sport or non-sport context. Void and solitude were often hard to distinguish:

[5] E.g., for Dutch NWO proposals: `https://www.nwo.nl/en/impact-plan-approach`
[6] We discussed how EU grant proposals are typically expected to show expected impact in scientific, societal, as well as economic ways – although in the later exercise we did not address the latter at all.

- Walking in the forest, foraging = safe solitude
- Cycling to work with a toddler in a cargo bike = co-experience of solitude (alone together)
- Solo hot air balloon flight = feeling insignificant in the greater scheme of things
- Rhythm of breathing during a run
- Slackline, being in one's bubble despite the audience (note: for some, it is rather the opposite, where the audience is intimidating)

Personal experiences resonating with the mini-holidays sub-facet were:
- Using swings on children's playground as an adult.
- Listening to an audiobook while using an indoor bike trainer
- Knitting (automatized gestures), wood-turning, pottery

The group tried to identify higher-level concepts within these examples. The ideas of contemplation, patterns/rhythm, craft/material experience, and purposeless activity were emphasized. However, here again, interindividual differences in the appreciation of craft (fear of non-mastery) or purposeless activities emerged.
Besides a general audience, possible target groups of reverie technology were discussed as a means to reach a more concrete level in ideation:
- Trauma survivors, where vulnerability and fears of specific situations are a challenge
- Eating disorders, with opposing needs where refraining from compulsive physical activity is a goal
- Aging population, with the opportunity to reframe solitude as positive

The group identified some opportunities linked to reverie Tech, e.g., overcoming obstacles in exercising related to the security to embrace an activity (yoga class), the body awareness or the threshold to get started.

**Roleplay.** The group roleplayed two concepts:
- Speed down: Contrary to the optimization and HR quantification (high exertion), we tried to slow down movement speed to support the pathway of getting into reverie. That was exemplified by showing, and afterwards, everyone successfully tried the "Beat Saver" slowing down exercise.
- Proprioceptive sonified gym ball: A gym ball that would turn the attention of the user inward through sonified feedback related to the propriocepsis of moving the ball along one's own body.

**Figure 14** Roleplaying the two concepts related to the SportsHCI lens Reverie.

**Impact pathway.** Burnout and stress are on the rise, negatively impacting mental health. Research shows that physical activity is a key factor in preventing burnout and stress, and improving mental health in overall. Yet, people experiencing burnout or stress often face a set of barriers (e.g. feeling not good enough to engage in PA, feeling tired due to the burnout, feeling they have no time because they are stressed, etc). Hence, these people who would benefit significantly from PA are not doing so enough. Reverie technologies could offer a solution in that regard. Building on principles of positive solitude, feeling void and experiencing mini-holidays, reverie technologies promote light physical activity, which can be done with little equipment and alone, and that is intrinsically motivating. It can help reframe solitude as a positive feeling. The fact that it is quite opposed to the current narrative around sports is revolutionary and is an ideal entrance point for people to engage with movement. Yet, there is a need to understand exactly how these technologies can be designed to address different people/different contexts, what are the triggers of reverie and how to design for the right balance in order to reach positive reverie experiences. A second design challenge related to the use of technology in reverie experiences, as those are usually low-tech or even purposively done to disengage with the stress and fast pace of life attached to the use of technology.

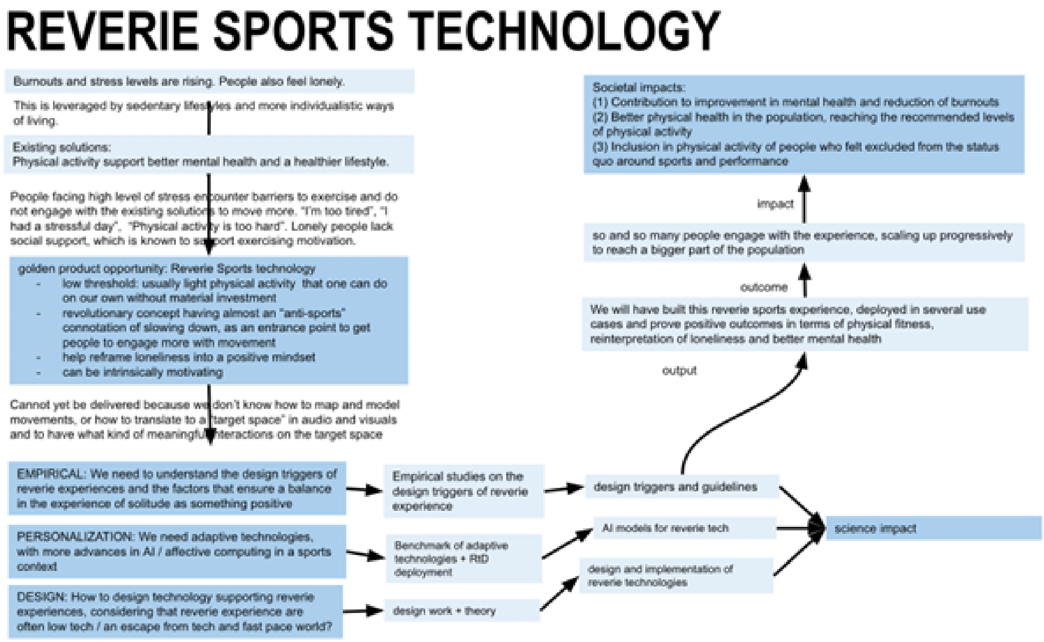


**Figure 15** Impact pathway sketch for "reverie sports technology".

The group brainstormed two alternative entry points/problem statements that could justify Reverie Sports technology.

- The first one was simply to start from the observation that people do not reach the recommended levels of physical activity → Individual sports are easier and have a lower threshold to start (walking and running are the most popular options) → Despite this low threshold, people start running or walking but do not manage to sustain this practice due to several barriers → Reverie Technology as a solution.
- The second one took a complete different approach, with a problem statement outside of the realm of sports. There is an urgency to take action to combat climate change and adopt more sustainable lifestyles → Connection with nature can be a driver for individual behavior change → Reverie technology can be used to raise awareness for nature.

## 7.2 SportsHCI lens: Pleasure

*Regina Bernhaupt, Maria Fernanda Montoya Vega, Perttu Hämäläinen, Vincent van Rheden*

**License** Creative Commons BY 4.0 International license
© Regina Bernhaupt, Maria Fernanda Montoya Vega, Perttu Hämäläinen, Vincent van Rheden

First, the group spent time reading the description of the pleasure lens from the paper. Then, we discussed our understanding about the lenses and how we conceive the definition of pleasure in different ways than the ones provided by the paper. Next, we provide our first person experiences about the pleasure we encounter from physical activity, such as intense exertion through sports like cycling and crossfit, very similar to the discomfort subcategory in the paper. On the contrary, we also discussed the pleasure we get from less intense exertion but more expressive physical activity, such as dance, which is not described in the paper. We then discussed how this last one is less explored in our field and then we brainstormed in which scenarios we encounter pleasure from dance and music. Finally, we came out with the case of a work office environment in which workers are seated all day and are under great amounts of stress. Most of the time these workers don't have the chance to have active breaks, sometimes because they don't have access to facilities that allowed them to have a physical active break, or sometimes they don't have the time. Therefore it would be useful to provide and active break through dance, allowing them to get rid of the stress while performing bodily movements following the music. We also imagined a system, for example VR, that allows the worker to switch off the external work environmental and encourage them to move however they want avoiding feeling judged by others.

**Roleplay.** The group roleplayed a concept where a boring and stressful moment in the life of an office working suddenly gets transformed into an inspiring, enthusiasm-eliciting, physically active episode through the use of immersive VR.

**Figure 16** Roleplaying the concept of pleasurable and active work environments.

**Impact pathway.** Work environments are known for be stressful. Moreover, in work environments, jobs proliferate that involve sedentary behaviors such as computer-based work where workers have to remain seated most of the time. Although the promotion of physical activity to tackle stress and sendentarity is a worldwide concern, workplaces have not yet profited from physical activity during working hours, where adults spend an average of 7.9

hours per day. Barriers to perform physical activity while in work include low access to outdoor spaces, low access to specialized physical activity spaces, overload of work, poor engaging alternatives such as “active breaks”. However, there are opportunities to support more active work environments through the use of interactive technologies, while providing opportunities for workers to express themselves and engage in fun activities. Particularly, interactive technologies that encourage the pleasure from physical activity, such as the pleasure from exertion and the pleasure from being recognized by others when doing physical activity, present great potential for this. Therefore, there is a need to understand how to design more active work environments using interactive technologies such as virtual and augmented reality, what are the suitable physical activities that workers will consider more engaging during working hours, and when to promote the development of these physical activities.

To begin building this understanding, we propose a research through design methodology involving workers and managers of a company’s office space. Through this design process we aim to gain insights of the feasibility of interactive technologies to be adapted in a working space and the perceptions of workers towards them. The final design should help workers to reduce their stress through physical activity during working hours, influence their productivity and their willingness to develop more physically intense breaks. Ultimately, this proposal aims to impact the health of workers during their jobs, and inspire HCI researchers to design interactive systems to encourage physical activity in workplaces.

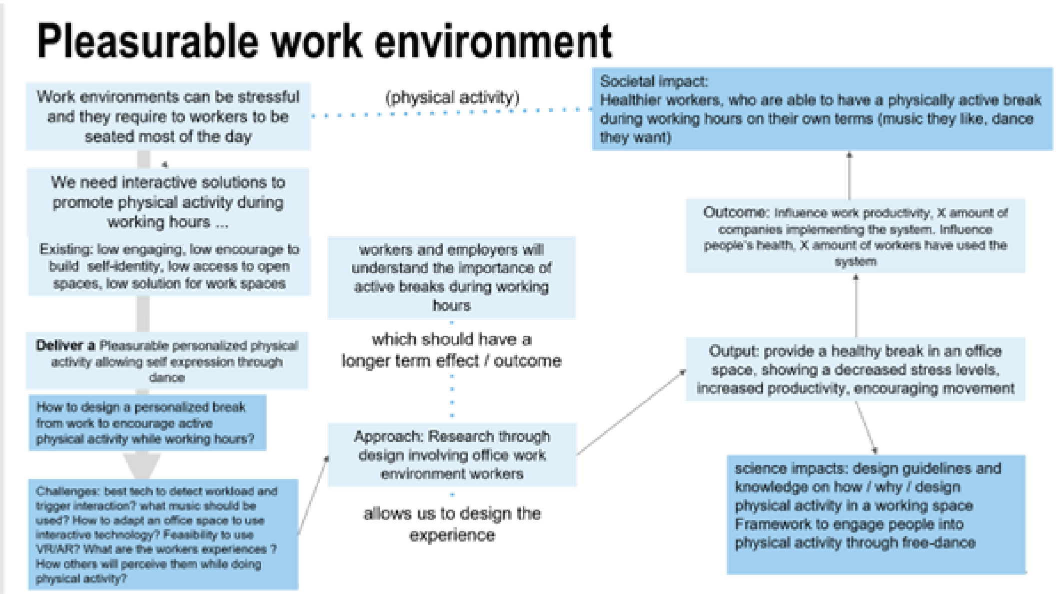


**Figure 17** Impact pathway sketch for “pleasurable work environment”.

## 7.3 SportsHCI lens: Beauty in movement – Multi-equality spaces in sports

*Dennis Reidsma, Andrii Matviienko, Xipei Ren, Paolo Buono*

License Creative Commons BY 4.0 International license
© Dennis Reidsma, Andrii Matviienko, Xipei Ren, Paolo Buono

First, we spent time reading 2.5 pages about the beauty lens from the paper. Second, we freely associated and reflected on the main aspects of beauty such as rhythm, expressed our opinions about them, and proposed possible extensions, such as a timely unfolding of a sports performance and the beauty of time in sports. Then, we did solitary ideation: keeping the lens and those extra associations in mind, what ideas come up to design for beauty? We spent five minutes writing down our ideas about how technology can possibly facilitate beauty in sport activities and ended up with six ideas. The descriptions of the lens and sub-facets and the concrete design suggestions in the paper provided an entry point to start thinking about design ideas once we identified a goal besides the "beauty". In addition, the associative concepts that we came up with, made the design possibilities feel more nuanced and detailed. Once we had a number of concepts, we presented them to each other and shortly discussed and extended them. Afterward, we briefly discussed which of these ideas can be clearly presented via a pantomime or a bodily representation. In the end, we decided to pick an idea that would fit the best for bodily representation and would be easy for an audience to understand.

**Associative concepts.** tension, development of tension in movement; expressive range of movement; internal appreciation of beauty and external appreciation of beauty; synchronization and harmony: alternating between making and breaking the harmony (especially in certain antagonistic sports, but also in dance or music); time speeding up and slowing down as expressive parameter; repetitive or patterned movement without sticking strictly to fixed beats; virtuosity and mastery; shared beauty in the interaction between people; buildup, tension, climax – unfolding time in the beauty of sports and movement.

**Ideas.** Xipei Ren. (1) Outdoor running; Augmenting the outdoor experience of wind, birds, and other beautiful environmental factors through feedback technologies; using this to adjust pace of runner. Andrii Matviienko. (2) EMS system that amplifies your movements so you can consciously experience the space of dynamic tension in movement. Once you are familiar with the feeling you will be more able to use that yourself in expressive movement, thanks to the earlier EMS-enhanced experience. (3) Symbiosis between athlete and fan: sensors on athlete allow fans to co-experience the dynamic tension in the virtuoso movements of the athlete. Feedback of the audience's appreciation to the athlete (embodied/haptic/tactile feedback?). (4) Appreciating the beauty of doing sports movements, for non-athletes: first-person enhanced multimodal replay of movements that were previously measured on an athlete. Dennis Reidsma. (5) Movement mapping for equal participation in sports. Given that some people struggle to express conventional beauty in movement (e.g., spasm, paralysis, etc.), measure the movements they can make within their own range/space of expressive movement; similarly measure the movement through the expressive range of other people; re-create the movement in a mapped audiovisual space that is the same for both people by having unique personal mappings; this allows co-equal participation in joint movement expressed in the mapped space. Both participants can together enjoy the unfolding, dynamic, joint movement while acting from within their own movement capacity. Paolo Buono. (6) Starting from the premise that without experience, it is hard to perceive (let alone appreciate) the beauty in certain sports (or music, or dance...). Measure and detect the patterns in expressive activity; then use multimodal representation to reveal and augment / highlight the patterns so it is easier for audience to see and appreciate the beauty of complex expressive movement activities.

**Role play.** We picked the equal participation idea and titled it "multi-equality spaces". The role play was as follows: Two people (picture right) represent the "athletes" or human interactors. They make the same movement, which could conceivably be part of a dance or sports movement. One of the athletes has a much wider range of flexibility and control of motion than the other, who has only a very limited range of hand and arm motion due to e.g. spasm (bottom right). Two other people in the role-play (left in the picture) represent the computer's interpretation of these movements in an equality space in which the movements of the two people become comparable. This is presumably done by calibrating the space of movement of each, and then mapping these spaces to equivalent equal movement expression. On that mapping, one could then conceivably build a sport or other movement activity in which the two athletes could equally participate on their own movement terms and characteristics while becoming comparable in "effectivities".

**Figure 18** Roleplaying the concept of multi-equality spaces.

**Impact pathway.** Sport is good for maintaining health and other beneficial outcomes but is not equally accessible for all. People do not have the same capabilities to move expressively and competently in the same manner and amount. Existing solutions typically involve people with different needs or capabilities being put apart in their own sports category, e.g., paralympic games, gender-specific competition, age clusters in competition, etc. Thus, athletes cannot meaningfully interact with all others in meaningful expressive sports movement, on their own terms. This is the core of our Golden Product Opportunity: making a system that allows people with different capabilities to compare/compete/interact with each other in sports and expressive movement. The system normalizes the movement space of people by measuring movement in a person's range of motion, and mapping it to a common multi-equality space, generating audiovisual and other multimodal output in that multi-equality space, after which interaction between people's movements happens in that equality space. Although the idea would allow for equalized interaction between differently-bodied people, this product cannot yet be delivered because (1) we don't know how to measure, map, and model a person's movements in their unique movement space, (2) we don't know how to translate the mapped movements to a "target space" in audio and visuals, and (3) we don't know how to have (what kind of) meaningful interactions within that target space.

To address these challenges we work on four distinct science problems.

**Machine Learning to model personal movement:** Before we can model and analyze personalized movement spaces, we address the lack of datasets with different, personal and idiosyncratic ranges of sport movements. We collect and model different types of movements. This will yield a fair data set with a range of idiosyncratic expressive movements.

**Algorithms to map movement to a multi-equality space:** Even when we can measure expressive movement within a personalized range of capabilities, there is a lack of novel models and approaches to map these movements to a joint multi-equality space in which the movements of differently capable people become comparable. We will address this by developing models that provide mappings and weight coefficients of different types of sports movements, which will also yield new data analytics methods and approaches for normalizing multi-equality spaces in sports.

**Multimodal environment for expressing movement in the multi-equality space:** Provided that, through the previous step, we can map users' movements to a shared and equal space, we do not yet know well enough what kind of multimodal output form should be generated from the mapped movements, such that it generates a meaningful embodied experience of dynamic expression for all types of participants, and such that it allows for meaningful back-and-forth interaction between different users. We will address this through movement-based and participatory design methods. This will lead to a multi-person, multimodal expressive environment for meaningful interaction, as well as a new theory about expressive visualization of movement for interaction.

**Evaluation of the impact on embodied participatory sensemaking in sports:** Once the platform is available, we must evaluate the effects of interacting in that environment on users' physical, emotional, and social experiences. Building on insights from the earlier codesign stage, we will carry out empirical studies with our system. This leads not only to insights about the appreciation and feasibility of our system, but will also contribute to more generalizable insights about challenges and opportunities in movement-based interaction between differently capable people.

The above-described work will have a scientific impact: The datasets, models, and algorithms can be used by other researchers to advance the state of the art in personalized movement modeling in various domains. The multimodal expressive output generation can be re-used and extended for a wide range of interaction domains including dance, music, sports, and social relationship building. Finally, the insights about movement-based interaction between differently capable people provide inspiration for new directions of inclusive movement-based design.

The above described will also yield, through its output, societal impact. Given the output of measurements, algorithms, models, and interactive systems, the outcome of our work is that we will have built this new co-equal sports and movement experience, and shown that it leads to a sense of being jointly engaged in coordinated, equalized, and harmonic sports movement. Through our codesign, evaluation, and dissemination activities, <unspecified number> of users will have experienced the possibilities of our system, and <unspecified number> of organizations will have encountered our new technology through participating in the research and dissemination (or even already incorporated the new technology in their organizational environment). This will lead to a growing impact on differently capable people who will fundamentally be more able to experience sports together in meaningful interaction, on their own terms, in equal interaction with all others – no matter body capability, gender, sportiness, skill level, or possible disabilities.

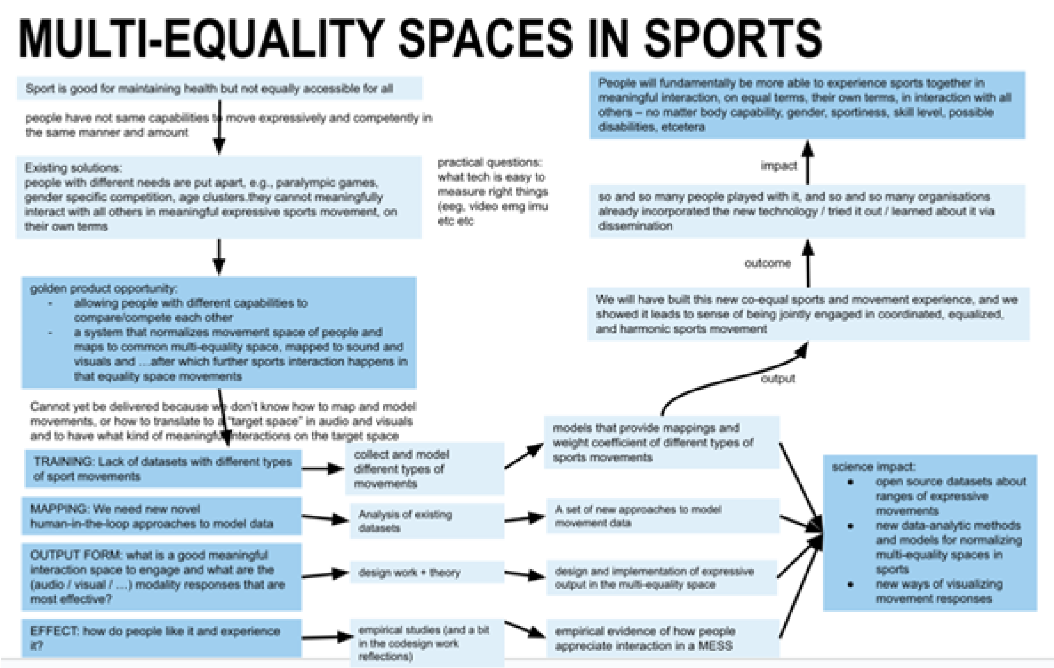


**Figure 19** Impact pathway sketch for "beautiful interaction in multi-equality spaces in sports".

## 7.4 SportsHCI lens: Pain in sports

*Michael Jones, Fabio Zambetta, Armağan Karahanoğlu, Don Samitha Elvitigala*

**License** 🅭 Creative Commons BY 4.0 International license
© Michael Jones, Fabio Zambetta, Armağan Karahanoğlu, Don Samitha Elvitigala

The group first reviewed the section on "pain" as a lens in the Sports HCI Lenses paper. We discussed the meaning of the word pain as constructed in that paper as well as in another paper Mike is working on. We focused on the concept that pain is a voluntary experience in sports.

- We reflected on our personal experiences in sport which included table tennis and running.
- We reflected on our experiences with parents and children in sport. This included some experiences with our own children and some experiences we'd seen involving other children.
- We generated several ideas about how to generate pain in several contexts.

**Tug of war concept.** We settled on "Tug of war" as the key inspiration, extending the game, as such:

- Two teams pulling the rope (could be 1 v 1)
- Players know that there will be pain, but don't know how we will distribute it.
- EMG sensors could be used to identify which muscle is working the most and direct a punch there.
- The pain is distributed equally across the groups.
- The order of pain stimuli is assigned randomly and is not known ahead of time. In addition, we envisage that the system implementing the gamified embodiment of pain will possess the following characteristics: (a) pain recalibration of pain (e.g., if after the first punch pain is 4/10, then the next may be stronger), (b) a recording system to relive the pain, as well as analysing where pain was inflicted more objectively.

**Figure 20** Roleplaying the Tug of War concept related to the SportsHCI lens Pain.

A key outcome of our system is triggering reflection on pain, what that pain means and whether one's sports goals are worth achieving. While participants may feel and internalise their own pain, they would not know how do that for others. We believe our system is geared towards someone who is already engaged in sport. Interesting additional questions included "How do you metabolise the loss or the win?".

**Impact pathway.** Pain is likely to occur in sports training and competitions. Child-athletes often give up when they encounter pain in sports, especially if their parents and/or coaches are not able to manage the adverse effects of pain, and rather push them to their physical or psychological limits. A holistic view of pain that involves perspectives from psychology, sports and technology is needed to prevent the emergence of such negative behaviour. In this project, we set out to normalise pain in sport, by delivering a holistic experience for embracing and reflecting on pain via a gamified embodied experience. In this experience, pain is delivered in a game with two teams through an exogenous process so that the pain cannot be directly related to the game and the pain is delivered equally to both the winning and losing teams. Participants and their parents reflect on the experience after the game ends. We do not yet understand pain in this context because we do not know: (1) the minimal amount of pain to administer in order to trigger reflection, and (2) the impact of pain in this setting. We will conduct studies to calibrate pain "dosage" per individual participant and to understand the experience of playing the game from both the child-athlete and parent (or coaches) perspectives. These studies will result in quantitative data about the physiological perception of pain in this context and qualitative data about the experience of pain in sports, for children, coaches and parents.

This project allows us to present pain in a balanced and constructive way to parents, child-athletes, and coaches. This presentation will help them better understand pain in a sports context. Our work may lead to improving environments for child-athletes so while they may experience pain in sports, this will not have an adverse effect on participation in sports. Furthermore, this work may contribute to insights into how to develop more healthy and less toxic coaching environments generally.

## 7.5 SportsHCI lens: Humility in sports

*Florian Daiber, Floyd 'Floyd' Mueller, Dees Postma, Robby van Delden*

**License** Creative Commons BY 4.0 International license
© Florian Daiber, Floyd 'Floyd' Mueller, Dees Postma, Robby van Delden

We discussed what the lens of humility was about and went through the three types, mostly directly in relation to some of the examples given in the paper. We reiterated that the big mountain in front of you brings a sense of humility, talked about the humility from competing

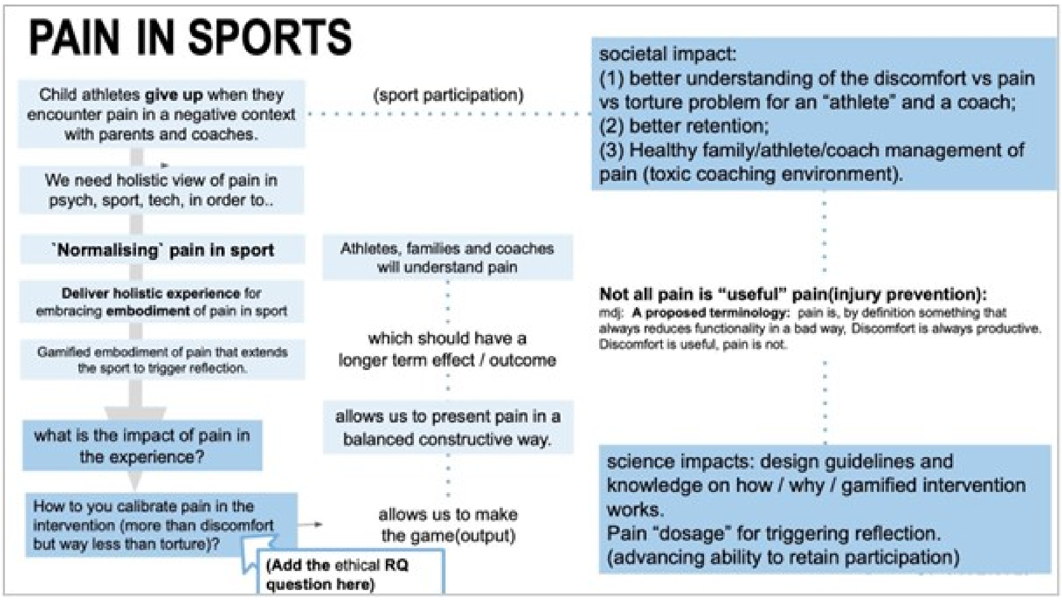


**Figure 21** Impact pathway sketch for "pain in sports".

against someone who is expert or better, only later more as a side note for completion we briefly mentioned but not really discussed trying again. Instead, we did discuss trash talk and bluffing in relation to humility. There is a difference of knowing your skills/cards (e.g. in poker you might exactly know what you have in the last round), versus not yet knowing your current capabilities of that day before going all out. The trash talks was also a picture in one of the slides shown after it was mentioned in the discussion, we related this also to the current sports event of the Tour de France. For instance, the coach sort of coaxing by saying "you can do it!" to Pogačar who was already completely beat. We later mentioned exemplars from media/every day use and own experiences. This led to the focus on experiencing actual parts of being a pro, which interestingly enough is also in the paper but more as a almost comparison rather than a more complete feeling of humility. Through the generation of the design space we saw a variation in first person experiences of pro sports which might make one humbled. After the discussion we wondered how other theoretical contributions or theory would play a role, or could help to shape the design space. Here, Newell's theory of constraints, phenomenology, and remediation/ post-phenemonology were mentioned but not discussed in too much depth.

Reflecting on the process, we see a recall and sharing of exemplar 1st/2nd/3rd person experiences, some focusing on the somaesthetic qualities (e.g., going a ridiculous speed in real or virtual life). There is a back and forth of ideas that would be interesting to experience humility, ideas that inspire a direct transformation of that idea into something related (e.g., from wavelight to seeing someone descend in the Tour de France), then talking these and associated ideas through we see other examples that already exist (e.g. cycling at higher speed due to wrong weight setting, experiencing downhill footage in connection to simulator) which we can build on but require and instigate a transformation of the last proposed own idea into something new (e.g., camera-based system currently not keeping up if you don't break downhill, rather than seeing the pace align the pace changing you with your "known" speed). This thus naturally transitions between inspiration, associated idea generation, to converging or actually taking steps sideways, see Figure below.

**Concept.** We identified that systems that allow you to compare your athletic performance to a professional athlete can facilitate humility. They are emerging (examples: wavelight, ...) and span a new design space. By allowing to compare yourself to a professional athlete, facilitating humility, we can help society appreciate the achievements, efforts and investments of professional athletes more, supporting a more empathetic culture.

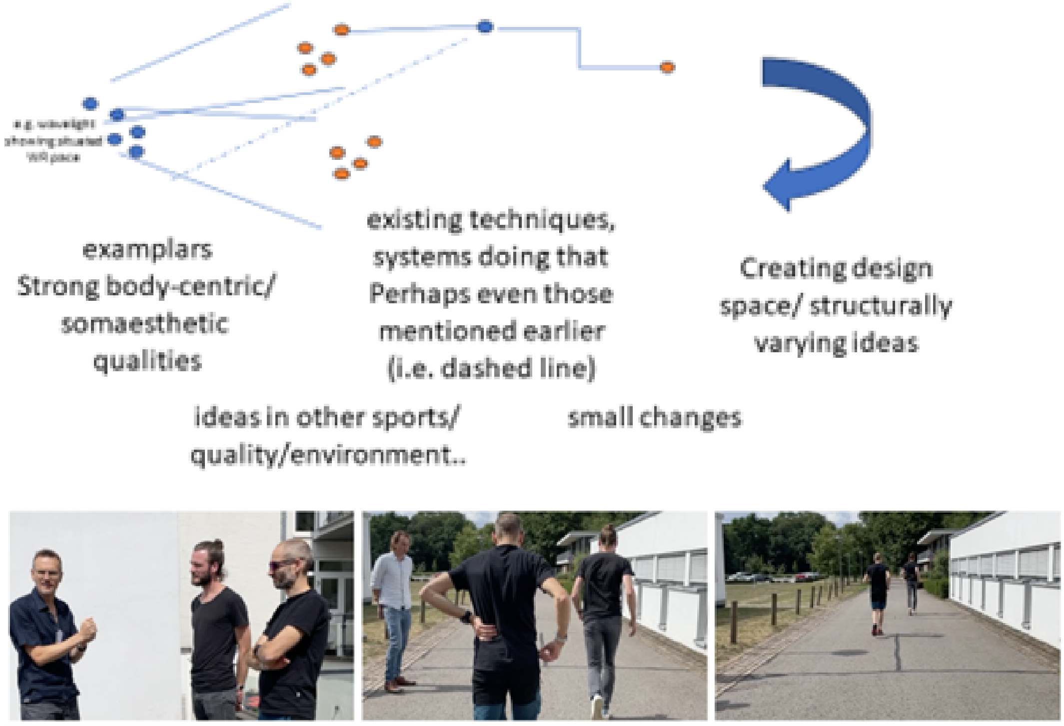


**Figure 22** Humility in SportsHCI: compare your performance to a pro.

**Impact pathway.** We mostly worked on furthering the design space. This also led to a remark about current Dutch grants on researching the value of elite sports for society/trickling down to amateur/grassroots sports. Building on the existing framework explained by Dennis, we see top right Sports is a human right and then counterclockwise:

- people are in bad shape (see papers on skills going down, lack of moving and interrelated issues, alternatively “not enough people experience the joy of moving often enough”)
- One powerful experience could be to see the pro athletes enjoyment/experience as a possible goal to work towards. It might be bringing a feeling of humility, triggering one to attempt it again and again.
- Technology is used to map or remediation between what the bodily experience/capabilities of the pro-athlete are versus what those bodily experiences/capabilities of the novice athlete are.
- These mappings are however initiated only sporadically for some sports and without a clear experiential framework and wider accessibility. So the scientific challenge becomes what can we remediate of the pro-athletes experiences, and how should we map these remediations?
- If we answer these in positive ways (where our related work/exemplars already show some opportunities) we offer a framework and set of experiences to more fully appreciate the pro athlete’s experience of pro-ness.

- This in turn we expect will lead to a wider uptake of such experiences, potentially instigating people's interest in sports, and fulfilling our fundamental right to experience sports and moving (cf. IOC's view and UN goals).

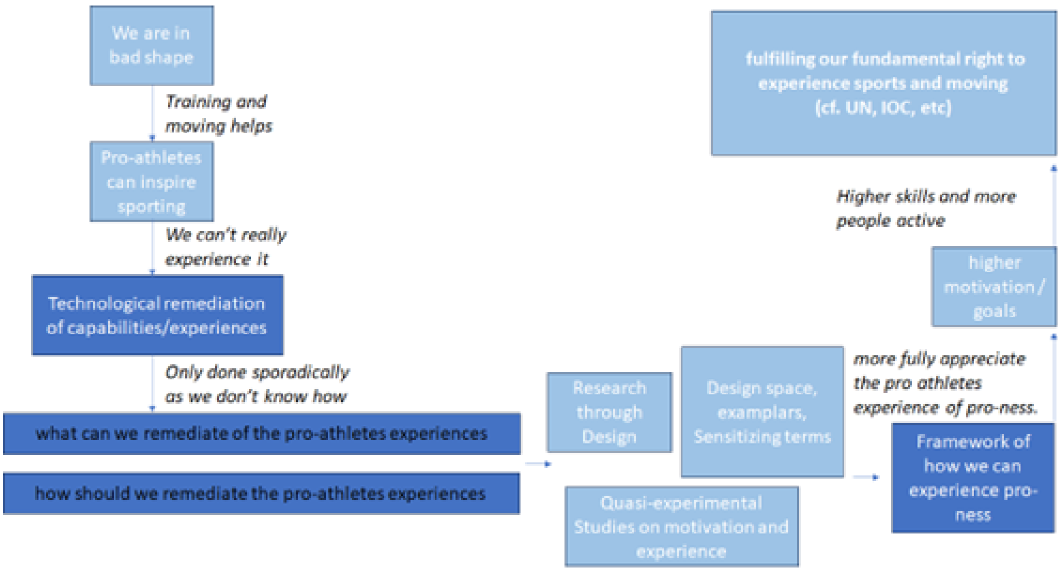


**Figure 23** Impact pathway sketch for "humility in sports".

## 8 Acknowledgements

We thank Dagstuhl for their extensive support and all the participants who contributed to this report as part of a collective effort. A particular thank you to Maria Fernanda 'Mafe' Montoya Vega for volunteering first to help and supporting the organizers in finalizing the report.

Dennis Reidsma: "As an organiser, I felt we could really focus on the content and participate in the work, thanks to the smooth facilitation by the Dagstuhl team from the start of the application all the way through the actual seminar."

Andrii Matviienko: "This is the best summer camp for nerds I've ever attended!"

Carine Lallemand: "Dagstuhl seminars are a truly unique format in academia. Magic happens when bringing together top researchers in a field and allowing them to truly focus and network during 5 days."

Laia Turmo Vidal: "I truly enjoyed the format of focused work in such a pleasant and nice environment! I will treasure the Dagstuhl experience and I am looking forward to great things coming out of it."

Lars Elbæk: "The Dagstuhl seminar will bring my cross-field of technology and interaction design in my primary field of sports science further in my future research. This was a highly experienced expert group that brings in-depth and a wide variety of knowledge into sports science will lightly influence future sports science. On the contrary, sports science will probably also affect future SportsHCI, I believe."

Mike Jones: "A perfect combination of academic discussions and informal activities. This was a great place to both imagine new research directions and build collaborative relationships with colleagues. It became clear to me that while competition is a key part of sports, other aspects of the sports experience such as enjoyment and discomfort are also a key part of the sports experience. By the end of the week, it was clear that experience and competition, enjoyment and discomfort have complicated relationships that need further study. It was valuable to wrestle with these concepts with a group of experts in sports and human-computer interaction."

Armağan Karahanoğlu: "Dagstuhl provided a fantastic natural sports opportunity to contemplate the future of SportsHCI. It was a great place to discuss the research opportunities to bridge sports science, HCI and human-centred design."

Vincent van Rheden: "Dagstuhl felt like a retreat for researchers, helping to deepen topics, connect with fellow researchers, create a feeling of community, being situated in a beautiful location away from the hectic of normal life. SportsHCI is where sports and interactive technology come together to create motivating, inspiring, and innovative work to help people move, grow, and enjoy sports and exertion activities."

## Participants


- Regina Bernhaupt
TU Eindhoven, NL
- Paolo Buono
University of Bari, IT
- Lisa Anneke Burr
Paris Lodron Universität Salzburg, AT
- Florian Daiber
DFKI – Saarbrücken, DE
- Lars Elbaek
University of Southern Denmark – Odense, DK
- Don Samitha Elvitigala
Monash University – Clayton, AU
- Perttu Hämäläinen
Aalto University, FI
- Daniel Harrison
University of Northumbria – Newcastle, GB
- Michael Jones
Brigham Young Univ. – Provo, US
- Armagan Karahanoglu
University of Twente, NL
- Carine Lallemand
University of Luxembourg, LU
- Andrii Matviienko
KTH Royal Institute of Technology – Stockholm, SE
- Maria Fernanda Montoya Vega
Monash University – Clayton, AU
- Florian 'Floyd' Mueller
Monash University – Clayton, AU
- Dees Postma
University of Twente, NL
- Dennis Reidsma
University of Twente – Enschede, NL
- Xipei Ren
Beijing Institute of Technology, CN
- Laia Turmo Vidal
Carlos III University of Madrid, ES
- Robby van Delden
University of Twente – Enschede, NL
- Elise van den Hoven
University of Technology – Sydney, AU
- Vincent van Rheden
Paris Lodron Universität Salzburg, AT
- Fabio Zambetta
RMIT University – Melbourne, AU


Report from Dagstuhl Seminar 23301

# Computational Proteomics

**Rebekah Gundry*[1], Lennart Martens*[2], and Magnus Palmblad*[3]**

1 **University of Nebraska – Omaha, US.** rebekah.gundry@unmc.edu
2 **Ghent University, BE.** lennart.martens@ugent.be
3 **Leiden University Medical Center, NL.** n.m.palmblad@lumc.nl

**Abstract**

This report documents the program and the outcomes of Dagstuhl Seminar 23301 "Computational Proteomics". This seminar was built around three topics: the increasingly widespread use, and continously increasing promise of advanced machine learning approaches in proteomics; the highly exciting, yet fiendishly complicated, field of single cell proteomics, and the development of novel computational methods to analyse the highly challenging data obtained from the glycoproteome. These three topics fuelled three parallel breakout sessions, which ran in parallel at any given time throughout the seminar. A fourth, cross-cutting breakout session was created during the seminar as well, which dealt with the standardisation efforts in proteomics data, and explored the possibilities to upgrade these standards to better cope with the increasing demands being put on the relevant data storage and dissemination formats. This report comprises an Executive Summary of the Dagstuhl Seminar, which describes the overall seminar structure together with the key take-away messages for each of the three topics. This is followed by the abstracts, comprising three introduction talks, one for each topic, which were intended to whet the participants' appetite for each topic, while also introducing an expert perspective on the current challenges and opportunities in that topic. Along with the topic talks, two ad-hoc talks were presented during the seminar as well, and their abstracts are provided next. Moreover, each breakout session also comes with its own abstract, which provides insights into its discussions and relevant outcomes throughout the seminar.

**Seminar** July 23–28, 2023 – https://www.dagstuhl.de/23301
**2012 ACM Subject Classification** Applied computing → Bioinformatics
**Keywords and phrases** bioinformatics, glycoproteomics, machine learning, mass spectrometry, proteomics, single cell proteomics
**Digital Object Identifier** 10.4230/DagRep.13.7.152

## 1 Executive Summary

*Lennart Martens (Ghent University, BE)*
*Rebekah Gundry (University of Nebraska – Omaha, US)*
*Magnus Palmblad (Leiden University Medical Center, NL)*

**License** Creative Commons BY 4.0 International license
© Lennart Martens, Rebekah Gundry, and Magnus Palmblad

The Dagstuhl Seminar 23301 "Computational Proteomics" was based around three key topics of rapid development in mass spectrometry-based proteomics, which were discussed in-depth in light of their challenges and opportunities. These three topics were: (i) the expanding and highly successful adoption of machine learning (ML) approaches in proteomics; (ii) the varied computational challenges posed by the very recent, but very rapidly evolving field of single cell proteomics; and (iii) the possible paths to adoption of advanced computational

* Editor / Organizer

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

approaches in the challenging field of glycoproteomics. Each of these topics was introduced by a short lecture, delivered by an expert in the field, and focused on two main goals: (i) to provide an informed opinion of the current state of the field, while highlighting its key challenges; and (ii) to thus entice the participants to contribute their own views on this topic, to help set the agenda for the discussions throughout the remainder of the seminar. Apart from these three invited talks, two ad-hoc talks also emerged during the seminar, and these concerned the specific topic of large scale spectral clustering, and the promise of using the Rust programming language in proteomics applications.

Based on the ideas collected after each of the introductory, topic-specific presentations, a list of discussion points was collated for each topic, and three parallel breakout sessions were then organised in the mornings and afternoons around these discussion points. A final, joint session in the evening of each day served to bring all participants from the different breakout groups together again, and summarized the key points of their respective discussions. Moreover, these joint sessions were also used to update the lists of discussion points for the three topics with any newly emerged points, and to reprioritise discussion points for the next day's breakout sessions.

The Machine Learning Working Group had a lot of topics to explore, mostly focusing on refining current approaches, by, for instance, introducing quality control and explainable AI, as well as setting out new applications for ML in the field. The latter included the possibility of a foundational model, the extended prediction of analyte behaviour in the instrumentation, and the possibility to analyse the resulting models to gain a better understanding of the physico-chemical properties at play in the analytics workflow. Finally, community-building efforts were discussed, and suggestions for improvements of existing initiatives (notably proteomicsML.org), as well for novel community engagements were made.

The Single Cell Proteomics Working Group discussed the applicability of current tools in the new discipline of single cell proteomics. Correspondingly, issues in capacity in current tools also came up, in light of the fast-growing data sizes for singe cell experiments, which are currently at hundreds of analytical runs, but likely soon expanding towards thousands of runs; well beyond the capabilities of present-day algorithms. Standardisation of this field, and its metadata, which is even more inportant given the sheer size and complexity of typical single cell proteomics data sets, was also considered in some detail. This entailed the standardisation of the workflows and algorithm parameters, which are currentl very diverse and specific, as well as the standards for data and metadata representation and dissemination.

The Glycomics and Glycoproteomics Working Group saw plentiful opportunities for the field to strenghten their bioinformatics, and put emphasis on adopting machine learning techniques in their field. However, they also saw open challenges regarding the collection and annotation of their data. Some time was also spent on identifying current weaknesses in their field, notably the quantification of glycopeptides, and possible avenues for addressing these.

At the end of the seminar, a critical assessment of the seminar was performed by all participants, highlighting the strengths and improvement points of the overall Seminar organisation, and a list of potential future Dagstuhl Seminar topics was drafted based on the participant's input. The assessment of the Seminar highlighted in particular the extremely fruitful nature of the open and engaging discussions, the unique and highly valuable nature of Schloss Dagstuhl and its unmatched seminars, and the ongoing gratitude of the computational proteomics community for the opportunity to convene in this singular setting. Concerning possible future topics, a plethora of enticing options were put forward, indicating that the field of computational proteomics remains in full expansion and that it continues to brim with both challenges and promise!

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Topic introduction: Machine learning in Proteomics: everything everywhere all at once

*Robbin Bouwmeester (Ghent University, BE)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Robbin Bouwmeester

It can be clear that machine learning has integrated with many (if not outright most) of proteomics data analysis, as it has been a key ingredient in today's proteomics informatics advances. At the same time, this widespread uptake of machine learning in the field can correctly be labelled as only a beginning. Indeed, there are many different new models available in the domain of machine learning, including large language models, tarnsfer learning, and graph convulational networks, alongside many new applications in proteomics itself, such as the prediction of chromatographic elution profiles, optimal experiment design solutions, and the provision of confidence intervals on predictions, all to name but a few. Moreover, the onset of machine learning models in the field has alos led to a need for comparing and evaluating the impact of these models. Of course, the presence of machine learning-interested researchers in proteomics has also led to calls for the establishment of a machine learning in proteomics, community, which in turn also led to education efforts in the form of ProteomicsML. As a result, there are several interesting opportunities regarding the further advancement of machine learning in proteomics, which can be organised into thee three subtopics: (i) new models and applications; (ii) comparing of models and their downstream impact; and (iii) community-building and support around machine learning in proteomics.

## 3.2 Topic introduction: Mass spectrometry-based single-cell proteomics: current challenges

*Laurent Gatto (University of Louvain, BE)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Laurent Gatto

Single cell proteomics is a relatively new field within proteomics, as the analysis of the truly minute quantities of proteins from single cells had been regarded as well outside the range of the mass spectrometers' limits of detection. However, instrumentation and methodological advances have allowed these limits to be breached, and have resulted in a rapidly growing interest in, and uptake of, single cell proteomics. This is not to say that analysing the proteome of a single cell is now by any means a routine endeavour. It remains a highly challenging and work-intensive process, and requires dedicated instrumentation and extremely careful sample preparation and handling. However, the challenges in single cell proteomics are most certainly not only practical. Indeed, considerable computational challenges accompany this young field, and these can be summarised in three categories as identification challenges, quantification challenges, and statistical data analysis challenges. Indeed, at all levels of the data processing workflow, single cell proteomics provides unique challenges for present-day computational algorithms and approaches, and thus plentiful opportunities for research into novel methods.

## 3.3 Topic introduction: Glycomics and glycoproteomics – Computational challenges and opportunities

*Sriram Neelamegham (University at Buffalo – SUNY, US)*

**License** Creative Commons BY 4.0 International license
© Sriram Neelamegham

The field of glycomics/glycoproteomics is a challenging one, due to the combined complexity of requiring analysis of both glycan structures as well as their carrier peptide sequences. Importantly, while a vast array of software tools are available for glycomics (the study of the glycans) on the one hand, and glycoproteomics (the study of the peptides carrying the glycans) on the other hand, many of these tools have been deprecated swiftly after publication, are not open source, or lack modularity. This unfortunately greatly limits the actual availability of computational solutions in the field. A classification of the various computational approaches can be made in four types: peptide-first, glycan-first, O-glycan specific, and glycan database focused algorithms. These approaches vary in their sensitivities, and limited overlap between the results of various tools is the rule rather than the exception, as also highlighted in a recent Human Proteome Organisation (HUPO) study on the topic. However, a very promising new avenue is provided by machine learning approaches, and it is very encouraging to see heightened interest from the field in supporting, and subsequently adopting, such approaches. However, some key challenges remain in the field, notably the lack of large, well-curated data sets. There are interesting opportunities available for exploration, however, especially if curated data sets of sufficient size were to become available, and these include the use of ion mobility information in the identification process, and the possibility to predict retention times from glycan (or even peptido-glycan) structures. In summary, some of the most pressing challenges in the field of glycomics/glycoproteomics are clearly computational in nature, and there is therefore ample opportunity for innovation in this field.

## 3.4 Ad-hoc talk: Unsupervised clustering of spectra

*Lukas Käll (KTH Royal Institute of Technology – Solna, SE)*

**License** Creative Commons BY 4.0 International license
© Lukas Käll

Tandem mass spectrometry experiments generate large amounts of spectra, generally with high redundancy between and within samples. An exciting way to process such data is to use unsupervised clustering to group spectra with similar appearance that tentatively were generated by the same analyte across samples. By associating the spectra to their corresponding MS1 abundances, the approach enables the quantification of analytes represented by spectra before their identification. The approach also allows the merging of the spectra of the clusters into consensus spectra, potentially increasing the accuracy of further identification. This approach has successfully been applied to proteomics data (`https://doi.org/10.1038/s41467-020-17037-3`), and we discussed if the approach could be further extended into the analysis of glycans and glycopeptides.

### 3.5 Ad hoc talk: A community-driven project – Rusteomics

*Dirk Winkelhardt (Ruhr-Universität Bochum, DE & ELIXIR Germany – Jülich, DE)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Dirk Winkelhardt

The Rust programming language has been making inroads in recent years, due, amongst others, to its speed of execution, its use of centralised libraries (crates) coupled to an excellent dependency manager (Cargo), and the ability to be cross-compliled to multiple platforms. Moreover, Rust is conceived to be particularly safe, as the memory ownership model protects from memory leaks at compile time. In proteomics as well, Rust has started to gain some initial traction, and this talk documents a Rust in proteomics hackathon that took place at the EuBIC Developers Meeting in early 2023. After a brief introduction to the programming language, a simple exercise was completed by the participants in which the peak intensity of the peak closest to a target mass-over-charge was determined. The hackathon participants then engaged in testing the I/O speed of rust, which was deemed to be extraordinarily high for reading text-formatted mass spectra (in MGF format) and sequence database files (in FASTA format). In addition, efforts were spent on exploring the generation of API documentation (via Sphinx), project structure, and the use of Github-Actions on Rust and Python code. The participants also worked on bridging Microsoft .Net libraires with Rust code, as most vendors write their proprietary libraries (which are essential when reading their proprietary, binary data formats) in .Net. Finally, the participants drafted future plans for Rust in proteomics, which involved setting up a host of core mass spectrometry libraries, and integration of the Rust-written Sage open modification search engine in this overall framework.

## 4 Working groups

### 4.1 Working Group Report: Machine Learning in Proteomics

*Robbin Bouwmeester (Ghent University, BE), Viktoria Dorfer (University of Applied Sciences Upper Austria, AT), Laurent Gatto (University of Louvain, BE), Arzu Tugce Guler (Leiden, NL), Tiannan Guo (Westlake University – Hangzhou, CN), Michael Hoopmann (Institute for Systems Biology – Seattle, US), Lukas Käll (KTH Royal Institute of Technology – Solna, SE), Ville Koskinen (Matrix Science Ltd. – London, GB), Lennart Martens (Ghent University, BE), Magnus Palmblad (Leiden University Medical Center, NL), Tobias Schmidt (MSAID – Garching, DE), Veit Schwämmle (University of Southern Denmark – Odense, DK), Mathias Wilhelm (TU München – Freising, DE), and Dirk Winkelhardt (Ruhr-Universität Bochum, DE & ELIXIR Germany – Jülich, DE)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Robbin Bouwmeester, Viktoria Dorfer, Laurent Gatto, Arzu Tugce Guler, Tiannan Guo, Michael Hoopmann, Lukas Käll, Ville Koskinen, Lennart Martens, Magnus Palmblad, Tobias Schmidt, Veit Schwämmle, Mathias Wilhelm, and Dirk Winkelhardt

The Working Group on Machine Learning (ML) in Proteomics convened every day, and even split off into two parallel groups for certain sessions to pursue multiple topics of discussion. This abstract reflects these dicussions and their ouctomes in chronological order throughout the Seminar.

On the first day, the Working Group began by looking at quality control (QC) and ML, and first of all, the ambiguity in the topic title was unpacked. Indeed, one could look at the use of ML in QC, or conversely, at the QC of ML. On the topic of using ML in quality control (QC), a presentation of data from several labs and instruments using paired features that may be suitable for ML and prediction of the source of experimental problems. Highlighting the discrepancy between observed and expected (ML predicted) features is key for successful QC, whether it's done using internal standards or, for example, relative entropy. On the topic of the QC of ML, one of the key needs is for users to know when they are using a model (too far) out of context. We discussed some very simple ways on how this could be presented, such as a (box) plot of model performance on the current data set(s) in the context of a boxplot of the expected performance of that model. Tools that provide such feedback are much more likely to get used in practice, and should confer higher confidence in the applicability of the model to that data set. From this, a digression into explainable AI was made, where feature relevance could be assessed by looking at training gradients, or looking at discrimination power rather than correlation between predicted and measured features.

The second day began with a discussion centered around the concept of Foundation Models, which are versatile models trained on broad data for various downstream tasks. The concept and its applications were introduced, which included for instance gene expression prediction. Participants questioned the potential for a proteomics-focused Foundation Model, given the challenges of data annotation. The group outlined a multi-step plan involving data collection, processing, and integration into a Foundation Model. Challenges, such as fragmentary and poorly annotated data, and variations between labs and instruments, were acknowledged. The group suggested that separate Foundation Models for sequences, spectra, and other data types could be linked together for a combined output. The conversation shifted to a proposed first spectra-only Foundation Model project. Several data repositories were suggested for sourcing spectral data, with the intent to cover as much sequence diversity as possible. Converting data to tensors for GPU computation was discussed, along with specific spectral encodings. The session resulted in a concrete plan to start by applying for GPU resources and a number of other action points for the coming months.

On the afternoon of the second day, the Working Group focused on two distinct topics. The first of these centered on the machine-learning based prediction of the behaviour of peptide and small molecule analytes on the instruments. We aimed to identify explainable properties, the models to use, and any opportunities for deep learning. Trivial, or already ML-predicted properties like mass, isotopic distributions, peak shapes, diffusion rates, protein localization, and isoelectric point were mentioned, upon which the focus shifted to properties benefitting most from further investment of machine learning techniques, such as ion suppression, ion mobility/collision cross-section (CCS), pH, and chromatographic column temperature. Revisiting early work with molecular dynamics and the application of association mining were raised. The impact of post-translational modifications (PTMs) on retention times and collision energy optimization for improved fragmentation were considered. Applying AI methods for protein structural dynamics, and for exploring enzymatic digestion and protein fractionation were also discussed. Action items include a review manuscript focusing on successful deep learning models like pDeep, Prosit and DeepLC. Throughout, this discussion highlighted machine learning's continuing potential for understanding and predicting physicochemical properties in MS-based proteomics.

The second topic of the afternoon of the second day discussed a project that aims to develop an end-to-end model to predict protein quantities directly from raw data, potentially linking it to a phenotype or disease state. This is a deeply challenging endeavor, but there are already some preliminary data available for reprocessing. One project, soon to be published, offers around 20k DIA runs of serum samples with quantified proteins. Nevertheless, questions remain regarding the amount of available data. The model architecture is also a crucial consideration. While convolutional neural networks (CNNs) are one option, it is unclear whether the local structure is sufficient to become informative hyper-features. An alternative might involve using a transformer network that processes the output of a scan or window. The model's input is a tensor or matrix, possibly with mass-over-charge and intensity dimensions. However, there are concerns that the matrix might be too high-dimensional for conventional deep learning. The DIA tensor (DIAt) tool has been proposed to convert DIA (SWATH) data into a tensor, effectively creating large "images". Another possible concept to explore involves graph convolutional neural networks. With this approach, peaks are nodes, and correlations or co-occurrences between peaks form the edges. However, this approach could be problematic as calculating correlations between features might go against the idea of an end-to-end model and may complicate feature finding or peak picking. A possible starting point for the project could be focusing on fourteen proteins in serum or blood, to begin with. Incorporating ideas from the field of ion networks might also be beneficial, as this essentially formes a network of ions that could be used as a starting point.

The third day saw the Working Group first explore ways to interrogate deep learning models for their knowledge about physicochemical rules through data-driven exploration. Retention time and fragmentation spectrum prediction models seem to be most suitable to extract rules that are "semi-orthogonal" to what they have been trained with. That could include details about binding energies, structure and the impact of motifs. The approach would be to generate simple hypotheses and then explore their potential impact on the predictions. This approach can be automated into building a Great Hypothesis Tester built on regular expressions for the peptide input.

The afternoon of the third day considered the ProteomicsML.org resource, which was first demonstrated, followed by general discussions on possible additions and modifications. Suggestions include adding recommended reading for ML and proteomics beginners, linking to existing tutorials, and covering relevant frameworks. Currently, the platform uses various ML frameworks and mainly focuses on behavioral predictions. Considerations include simplifying ML frameworks, introducing structural methods like AlphaFold, and clarifying the purpose of ProteomicsML. Outreach should emphasize seeking contributors, offering lessons and tutorials, and addressing platform maintenance.

On the fourth day, finally, the discussions began by considering community aspects of ML model reproducibility through better access to data and code. Platforms such as ProteomicsML, DLOmix and Zenodo were suggested for sharing and comparing models. Ongoing efforts are being coordinated with BioHackathon Europe, and their BioModelsML project. Another point was that model evaluation requires better comparison metrics, possibly focusing on relevant biological information such as coverage of relevant pathways, rather than simply counting peptide-to-spectrum matches. Other topics covered variational autoencoders, probabilistic modeling, and missing value imputation, with suggestions to use noise injection and quantitative data, building on the histone-based "proteomic ruler".

## 4.2 Working Group Report: Single Cell Proteomics

*Laurent Gatto (University of Louvain, BE), Bernard Delanghe (Thermo Fisher GmbH – Bremen, DE), Melanie Föll (Universitätsklinikum Freiburg, DE), Lukas Käll (KTH Royal Institute of Technology – Solna, SE), Ville Koskinen (Matrix Science Ltd. – London, GB), Lennart Martens (Ghent University, BE), Sriram Neelamegham (University at Buffalo – SUNY, US), Tobias Schmidt (MSAID – Garching, DE), Veit Schwämmle (University of Southern Denmark – Odense, DK), Mathias Wilhelm (TU München – Freising, DE), and Gamze Nur Yapici (Koc University – Istanbul, TR)*

**License** Creative Commons BY 4.0 International license
© Laurent Gatto, Bernard Delanghe, Melanie Föll, Lukas Käll, Ville Koskinen, Lennart Martens, Sriram Neelamegham, Tobias Schmidt, Veit Schwämmle, Mathias Wilhelm, and Gamze Nur Yapici

This abstract documents the discussions of the Working Group on Single Cell Proteomics, which took place throughout the Seminar, and whose subtopics are presented here in chronological order.

On the first day, the Working Group discussed that, whenever a novel technology is proposed, or an existing one is pushed beyond what was considered possible, such as in mass spectrometry-based single-cell proteomics (SCP), one can wonder whether existing software are still applicable. A notable comparison between fragmentation mass spectrometry scans from bulk and single-cell data highlighted possible differences, and thus wondered if existing search engines can be used as they are. Our working hypothesis is that the differences between SCP and bulk fragmentation mass spectra are mainly due to the lower intensity of the precursors, leading to some fragments getting lost, and that current tools can in fact still be used. We plan to test our hypothesis using data readily available from the members of our Working Group, comparing identifications with and without re-scoring of features, exploration of these features, and incorporating precursor intensity into the identification models.

On the second day, the discussion initially focused on SDRF metadata for proteomics. We first discussed which adaptations and challenges arise when using SDRF to annotate single cell proteomics experiments. We furthermore discussed how to generally improve the SDRF format e.g. more automatic collection of metadata and using it to facilitate writing materials & methods sections. Both would likely improve usage of SDRF by the community.

In the afternoon of the second day, the Working Group concentrated on recommendations for single-cell DIA data analysis. Single-cell TMT techniques are already "solved", because FDR control and batch effects are the same as non-single cell data. We further narrowed down to identification. With DIA, the two main tools are DIA-NN and Spectronaut. Users treat both as "black boxes", and much more transparency is needed on configuration options as well as the rationale for enabling or disabling various processing steps. We recommend saving the log file and exporting the config options at minimum, and tool developers are strongly recommended to provide sensible defaults and hide less-commonly used advanced options. It is also common to analyse a bulk sample together with single-cell raw files and use options like match-between-runs (MBR) to boost peptide and protein IDs. We will test the validity of this procedure by acquiring human single-cell data and analysing it together with 1) a bulk human sample (say 100 cells) and 2) an E. coli sample. The expectation is that the latter has no effect on identification rates. The opposite outcome would demonstrate the procedure is not valid.

The third day saw continued discussion on practical pipeline issues. Scaling of the data processing capability already is a current, and will be an especially important future challenge. Processing 100 raw files is still doable, but no tools currently handle thousands of samples.

Data size poses storage issues as well as data transfer problems, such as that one cannot easily upload a terabyte of data to the cloud. Moreover, many institutions forbid cloud processing of sensitive proteomics data, which differs from transcriptomics and genomics data processing. File formats like mzML are space inefficient, so can a better standard format be selected or developed? Or is it better to use just vendor raw files? The field also needs to clarify the research question at the end of a single-cell study. It is infeasible to interpret a spreadsheet with a thousand columns, so is it sufficient for software to only report protein differences between cells? To answer these questions, we need a concrete example, such as the 100 cell data set already discussed.

On the fourth day, the discussion touched on several different angles on how and whether transcriptomics of single cells could be combined with single-cell proteomics. Due to the extremely low sample amounts, it is difficult to do both on the same cell. Mosaic integration could allow some shared dimensions. Another discussion focused on the importance of the different modalities, and that we should expect to see the same effect. This, and the substantial differences in dimensions (a thousand for SCP, but ten thousand for scRNA-Seq) would also make their integration difficult or even questionable. Even within a single modality, focusing on proteomics, different features will provide different information. Carefully selected surface markers used in flow cytometry, for instance, will not necessarily be recapitulated, at least not as clearly, by hundreds or thousands of proteins; the same large, well differentiated clusters/types are expected to be recovered, but with much more variability. We finally discussed the notion of stable and unstable balance – protein abundances won't change easily if many copies are present in the cell, whereas proteins with only few copies can much more easily suffer changes in intensity. Moreover, (post-translational) modifications could also easily trigger quantification or cellular balances.

## 4.3 Working Group Report: Glycosylation and Glycoproteomics

*Rebekah Gundry (University of Nebraska – Omaha, US), Kiyoko Aoki-Kinoshita (Soka University – Tokyo, JP), Robbin Bouwmeester (Ghent University, BE), Robert Chalkley (University of California – San Francisco, US), Bernard Delanghe (Thermo Fisher GmbH – Bremen, DE), Viktoria Dorfer (University of Applied Sciences Upper Austria, AT), Melanie Föll (Universitätsklinikum Freiburg, DE), Arzu Tugce Guler (Leiden, NL), Catherine Hayes (Swiss Institute of Bioinformatics – Geneva, CH), Michael Hoopmann (Institute for Systems Biology – Seattle, US), Lukas Käll (KTH Royal Institute of Technology – Solna, SE), Ville Koskinen (Matrix Science Ltd. – London, GB), Karina Martinez (George Washington University – Washington, DC, US), Sriram Neelamegham (University at Buffalo – SUNY, US), Magnus Palmblad (Leiden University Medical Center, NL), Erdmann Rapp (MPI – Magdeburg, DE), Tobias Schmidt (MSAID – Garching, DE), Mathias Wilhelm (TU München – Freising, DE), Bernd Wollscheid (ETH Zürich, CH), and Gamze Nur Yapici (Koc University – Istanbul, TR)*

**License** Ⓒ Creative Commons BY 4.0 International license
© Rebekah Gundry, Kiyoko Aoki-Kinoshita, Robbin Bouwmeester, Robert Chalkley, Bernard Delanghe, Viktoria Dorfer, Melanie Föll, Arzu Tugce Guler, Catherine Hayes, Michael Hoopmann, Lukas Käll, Ville Koskinen, Karina Martinez, Sriram Neelamegham, Magnus Palmblad, Erdmann Rapp, Tobias Schmidt, Mathias Wilhelm, Bernd Wollscheid, and Gamze Nur Yapici

The Glycosylation and Glycoproteomics Working Group convened in several sessions throughout the Seminar, and discussed a variety of topics, which are listed in chronological order in this abstract.

On the first day of the seminar, the Working Group discussed the idea that we do not yet know how or if the presence of glycosylation affects fragmentation of the peptide backbone when acquiring MS/MS data on intact glycopeptides. Does the glycan affect observation and/or intensities of fragment ions? The question is important to answer as, if the presence of glycosylation does affect fragmentation, it stands to reason that fragmentation prediction algorithms and fragmentation interpretation algorithms would need to consider this to ensure accurate interpretation of the data. Major questions to be anwsered include: 1) Does the presence of a glycan (site occupied vs site not occupied) affect fragmentation? 2) Does the fragmentation of the peptide differ if the glycan composition on the site is varied? 3) Does the fragmentation of the peptide differ if the glycan structure on the site is varied? 4) Does the fragmentation of the glycopeptide differ if the glycan is attached to a glycopeptide that is the result of full tryptic digestion vs. one with ragged (non-tryptic) terminus? 5) Does the presence of N-linked glycosylation have the same or different effect on peptide fragmentation as the presence of O-linked glycosylation?

The second day saw the Working Group discussing current challenges with glycopeptide and released glycan analyses as it relates to best practices for interpreting and reporting data. Major concepts included discussion of ambiguity in glycopeptide and glycan assignments, the use and misuse of oxonium ions in glycopeptide spectra interpretation, the impact of search space on false discovery rate (FDR) calculations and accuracy of assignments, and the call to action for software developers to incorporate strategies to highlight assignments that may be ambiguous and to include GlyTouCan accession numbers, the reference annotation language for glycan compositions, topologies, and structures. Our discussion led to the development of an outline for a perspective or white paper focused on best practices in data analysis and reporting, which will serve as a catalyst to promote scientists to follow MIRAGE guidelines. It will include real-world examples to promote scientists at any level of experience in glycoproteomics or glycomics to become aware of key issues which may otherwise not be obvious to those with limited experience in this discipline.

In the afternoon of the second day, the discussion turned to a variety of topics related to how we can make more and better use of the glycopeptide and glycomic data that we generate. We discussed available tools for interpretation and integration of data, including a tour of tools on GlyConnect website. On the subject of glycopeptide quantification, the opinion of the group was that this has so far remained an unmet goal, depsite any claims to the contrary. We decided to add this as point #5 in the best practices manuscript outlined in the morning session of the second day – to outline best practices for anyone publishing in a journal that is not Nature. Several studies have been published that compare "quantitative" comparisons of glycoforms of glycopeptides to suggest that some glycan classes are "more abundant" on a site than another. Given the inherent differences in ionization potential for glycans which have different compositions, this is problematic. We reviewed published data that demonstrates the difference in peak abundance for different glycoforms. We will use these data as examples in our best practices manuscript. We need to get more funding to support glycoproteomics and glycomics. While we have a lot of success stories regarding analytical capabilities, applications are less known. There are certainly success stories to demonstrate how our approaches have impacted physiology and disease, including fundamentals of physiological processes and clinical examples. However, not all examples are highly visible. We discussed ideas for creating resources as a community to help educate others of our value. This could be videos, documents, websites that summarize success stories that would resonate with institutional leaders, administrators, benefactors, funding agencies.

The third day focused initially on interactions between proteomics and glycomics, with a team of experts in glycomics and proteomics discussing ideas on how to transfer knowledge from the proteomics field for application in glycomics workflows. The first part of the discussion focussed on current challenges in the analysis of glycopeptides and released glycans, particularly with respect to: 1. Structure assignment, i.e. the challenges in the identification of glycan composition in glycoproteomics studies, and also with respect to topology and structure determination that are common to both the fields of glycomics and glycoproteomics. These challenges stem from, amongst others, the isomeric nature of glycans. 2. Variable experimental workflows: current best practices in glycomics studies and glycoproteomics wetlab studies were discussed with respect to different experimental workflows used by multiple groups. The use of multiple experimental methods, such as derivatization strategies and fragmentation modes results in lack of consensus in the field. 3. Limitations of current software solutions: It was highlighted that glycosciences is a niche area, with relatively few researchers deeply involved. Knowledge was shared about observations on how MS intensity (or relative intensity) may be used for glycan topology assignments. The conclusion of the first part was that the glycan search space is not so large and it may be possible to develop spectral libraries or clustering approaches to find reliable topology solutions. Additionally, already established proteomics tools may be adopted using transfer learning approaches to aid developments in glycoproteomics. Among the methodologies discussed were i) MaRaCluster as it is agnostic to MS/MS spectra and may be used to develop consensus spectra. While there is no perfect method to develop consensus spectra for glycomics, a variety of approaches were also discussed ii) GLEAMS as reference spectra may be populated in this framework, perhaps using Siamese/pairwise inputs so that the system may be trained to find commonalities and distinctions among glycan classes. While there were many advantages to this approach, concerns were raised regarding the ability of proteomics trained datasets to learn glycomics experimental results, and the fact that the reference spectra may vary with collision energy and that the actual experimental results would have to be tightly clustered to a reference dataset. iii) Application of transfer learning, possibly using graph convoluted neutral networks, to predict MS/MS intensity. Challenges were discussed here with respect to the representation of branched glycan structures as graphs in such modeling. The final portion discussed the need to either write an independent perspective article related to this subject, or simply to use this discussion as discussion material in the paper that is ‘forthcoming’ from the first day’s session. This will be important to raise awareness and obtain additional funding to support the project.

The afternoon of the thid day brought a discussion on the common concepts we consider when trying to connect our glycomics and glycoproteomics data to biology, including a review of available resources to support these efforts. We also reviewed a public repository of bioinformatic tools. We discussed the state of the art of tools for glycopeptide and glycomic analysis, including limitations and what would be important to consider for future versions. Overall, challenges to selecting which tools make sense to focus on include that there is high heterogeneity in the field regarding which approaches to use, so not immediately clear that there is a single solution to focus on as it would risk putting effort into a tool that would not be broadly used in the field.

The fourth day, finally, continued the discussion of available tools and their limitations. We had new discussion of how we can better make use of machine learning and neural network tools to advance the field. An example emerged that centered around how we could use machine learning to help the interpretation of glycosidase reactions. Briefly, we often treat samples with glycosidases to help determine structural details. These reactions cause mass and RT shifts in the glycans/glycopeptides. This is probably something where ML could help us to automate and interpret the results. We ended by planning details regarding the two manuscripts we anticipate submitting by end of 2023. The first will be a research article that uses existing data to answer the question of whether the glycan/peptide affects peptide fragmentation. The other is a guidelines/tutorial type paper likely to be submitted to MCP.

## Participants

- Kiyoko Aoki-Kinoshita
Soka University – Tokyo, JP
- Robbin Bouwmeester
Ghent University, BE
- Robert Chalkley
University of California – San Francisco, US
- Bernard Delanghe
Thermo Fisher GmbH – Bremen, DE
- Viktoria Dorfer
University of Applied Sciences Upper Austria, AT
- Melanie Föll
Universitätsklinikum Freiburg, DE
- Laurent Gatto
University of Louvain, BE
- Arzu Tugce Guler
Leiden, NL
- Rebekah Gundry
University of Nebraska – Omaha, US
- Tiannan Guo
Westlake University – Hangzhou, CN
- Catherine Hayes
Swiss Institute of Bioinformatics – Geneva, CH
- Michael Hoopmann
Institute for Systems Biology – Seattle, US
- Lukas Käll
KTH Royal Institute of Technology – Solna, SE
- Ville Koskinen
Matrix Science Ltd. – London, GB
- Lennart Martens
Ghent University, BE
- Karina Martinez
George Washington University – Washington, DC, US
- Sriram Neelamegham
University at Buffalo – SUNY, US
- Magnus Palmblad
Leiden University Medical Center, NL
- Erdmann Rapp
MPI – Magdeburg, DE
- Tobias Schmidt
MSAID – Garching, DE
- Veit Schwämmle
University of Southern Denmark – Odense, DK
- Mathias Wilhelm
TU München – Freising, DE
- Dirk Winkelhardt
Ruhr-Universität Bochum, DE & ELIXIR Germany – Jülich, DE
- Bernd Wollscheid
ETH Zürich, CH
- Gamze Nur Yapici
Koc University – Istanbul, TR

Report from Dagstuhl Seminar 23302

# Software Architecture and Machine Learning

**Grace A. Lewis*[1], Henry Muccini*[2], Ipek Ozkaya[3], Karthik Vaidhyanathan†[4], Roland Weiss*[5], and Liming Zhu*[6]**

1 **Carnegie Mellon Software Engineering Institute – Pittsburgh, US.** glewis@sei.cmu.edu
2 **University of L'Aquila, IT.** henry.muccini@univaq.it
3 **Carnegie Mellon Software Engineering Institute – Pittsburgh, US.** ozkaya@sei.cmu.edu
4 **IIIT Hyderabad, IN.** karthik.vaidhyanathan@iiit.ac.in
5 **ABB – Mannheim, DE.** roland.weiss@gmail.com
6 **Data61, CSIRO – Sydney, AU.** liming.zhu@data61.csiro.au

**Abstract**

This report documents the program and outcomes of Dagstuhl Seminar 23302, "Software Architecture and Machine Learning". We summarize the goals and format of the seminar, results from the breakout groups, key definitions relevant to machine learning-enabled systems that were discussed, and the research roadmap that emerged from the discussions during the seminar. The report also includes the abstracts of the talks presented at the seminar and summaries of open discussions.

**Seminar** July 23–28, 2023 – https://www.dagstuhl.de/23302
**2012 ACM Subject Classification** Software and its engineering → Software architectures; Computing methodologies → Machine learning; Software and its engineering → Extra-functional properties; Computing methodologies → Artificial intelligence; Software and its engineering
**Keywords and phrases** Architecting ML-enabled Systems, ML for Software Architecture, Software Architecture for ML, Machine Learning, Software Architecture, Software Engineering
**Digital Object Identifier** 10.4230/DagRep.13.7.166

## 1 Executive Summary

*Grace A. Lewis (Carnegie Mellon Software Engineering Institute – Pittsburgh, US)*
*Henry Muccini (University of L'Aquila, IT)*
*Ipek Ozkaya (Carnegie Mellon Software Engineering Institute – Pittsburgh, US)*
*Karthik Vaidhyanathan (IIIT – Hyderabad, IN)*
*Roland Weiss (ABB – Mannheim, DE)*
*Liming Zhu (Data61, CSIRO – Sydney, AU)*

**License** Creative Commons BY 4.0 International license
© Grace A. Lewis, Henry Muccini, Ipek Ozkaya, Karthik Vaidhyanathan, Roland Weiss, and Liming Zhu

The pervasive and distributed nature of many of today's software systems requires making complex design decisions to guarantee important system qualities such as performance, reliability, safety and security. The practices within the field of software architecture guide the design and development of software systems from its high-level blueprint down to their implementation and operations. While the fundamentals of software architecture practices

* Editor / Organizer
† Editorial Assistant / Collector

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

do not change, their execution evolves to address architecting with, and for, new system paradigms and emerging technologies. Incorporating machine learning (ML) elements into systems is advancing rapidly with the availability of more compute power, specialized infrastructure, and better and more efficient algorithms to process increasing amounts of data. This evolution has led many domains to leverage ML for automation, data analytics, decision support, and advanced user interfaces, among others. Experience published and shared by both software engineering research and practitioner communities shows that organizations struggle to move ML models and pilot projects into production. Reported software engineering challenges fall mostly outside data science expertise, which is focused on the development of ML models and not necessarily their production-readiness, and include areas such as testing, requirements management, software architecture, and configuration management. Among these, software architecture has a special role as it serves as an abstraction between requirements and implementation and drives the structure and behavior of systems.

Some of the reported challenges in developing ML-enabled systems, such as interface stability, data storage and access, and data transformation, are already addressed by existing software architecture techniques and practices. However, traditional software development and deployment practices are faced with and pose significant challenges when developing ML-enabled systems, which are systems that integrate ML components. These ML components include model training and updating components; model serving and inference components; infrastructure components to support data collection, processing, and servicing; and operations infrastructure, such as MLOps pipelines to support automated build and deployment. As these new ML-related components are introduced into systems, different architectural concerns and architecting challenges take higher priority, which include working with non-determinism, understanding and designing for new classes of dependencies, and co-architecting the system as well as the ML model development pipelines, among others. In addition, integration of ML components into software systems places greater emphasis on qualities that are not as common in non-ML systems, such as monitorability as a first-class citizen, designing for extensibility, and data-centricity.

This Dagstuhl Seminar culminated as a response to recognizing that software architecture research and practitioner communities have an opportunity, if not an obligation, to fill key software architecture principles and practices gaps that are particularly critical when incorporating ML components into software systems. The seminar focused on two key themes:
1. software architecture principles and practices for ML-enabled systems (SA4ML) and
2. application of ML techniques for improved architecting of software systems (ML4SA).

A key goal of the seminar was to enable technical exchange among otherwise scattered research and practitioner communities, such as software engineering, software architecture, self-adaptive systems, and machine learning around software architecture and ML (SA&ML). This Dagstuhl Seminar presented an opportunity to get these communities together to develop a common vocabulary and a coherent research agenda with a better understanding of problems faced in the industry. The goals of this seminar were to establish a common understanding of key concepts that are central to architecting ML-enabled systems, elicit challenges of meeting key quality attributes of ML-enabled systems, and build a research roadmap for future work to address these challenges. As such, this seminar marks an important milestone in accelerating research in SA4ML and ML4SA work by fostering better communication around key concepts and between diverse stakeholders. Hereafter, we will use the term SA&ML to include both SA4ML and ML4SA research directions.

## Seminar Format

The seminar on software architecture and machine learning was structured to foster an interactive and productive idea exchange environment for all participants. Participants were initially asked to prepare a single-slide introduction about themselves and their work in software architecture and machine learning, which helped to establish a foundational understanding among attendees. Prior to the workshop, organizers developed an initial version of a key concepts map and a research challenges map using a grounded theory approach, which was represented visually on miro boards. During the seminar, participants delivered 10-minute lightning talks about their SA&ML work, which were instrumental in refining the research and concept maps with keywords and challenges mentioned by the speakers. At the end of the first day, organizers summarized the discussions and identified key emerging topics and research areas, leading to an initial set of key quality attributes critical to ML-enabled system development. The second day continued with discussions and a working group formation session, where participants prioritized five critical quality attributes for ML-enabled systems: Evolvability, Uncertainty and Observability, Trust and Trustworthiness, and Data Centricity. Four working groups were formed around these quality attributes, with appointed editors documenting the discussions. Each group was provided with a draft template to guide their discussions for consistency, and organizers rotated among groups to offer diverse insights. Plenary meetings on the second and third days allowed groups to share progress and discuss challenges, complemented by end-of-day meetings among organizers to consolidate learning and monitor progress. The seminar concluded on the fourth day with presentations from each group, open for feedback from other participants. This collaborative review process led to the integration of feedback into the final discussions, resulting in a well-defined set of challenges and research directions, marking the seminar's successful completion. In addition, this format ensured a thorough exploration of the seminar's themes and fostered active participation and collaborative learning among attendees.

## 2 Table of Contents

## 3 Overview of Talks

The talks presented during the seminar included two industry keynotes. *Alex Serban from Siemens Healthineers* talked about *Software Engineering for Machine Learning: Past, Present and Future Conjectures*. The keynote explored the intricacies of machine learning, with a focus on the critical need for emphasizing robustness and trustworthiness in ML systems. It highlighted the challenges associated with integrating machine learning components into broader software architectures and underscored the importance of addressing uncertainty within these systems. The presentation delved into the development of engineering practices tailored to machine learning, emphasizing the need for adopting these practices to enhance agility, software quality, and team effectiveness. The keynote concluded by reflecting on the current and future trends in machine learning, particularly the role of language as a universal interface and the growing complexity of models, stressing the responsibility of software architects in crafting trustworthy and robust systems amidst emerging regulatory challenges.

*Thomas Kropf from Robert Bosch* in his keynote titled *Industrial AI – Real-World Applications, Challenges and Solution Approaches*, focused on the application and challenges of Industrial AI at Bosch, highlighting the journey from the establishment of the Bosch Center for AI to integrating AI across various products. It emphasized the distinct nature of Industrial AI, where quality, scaling, and algorithmic robustness are crucial, particularly in safety-critical applications, illustrated by examples such as Manufacturing Analytics and Automated Optical Inspection (AOI) systems. The keynote also covered significant challenges such as enterprise-level scalability, quality assurance amidst model drift, rapid AI evolution, computational constraints in embedded systems, and the necessity for extensive tool support for data management. The potential disruptive impact of Foundational Models on AI practices was also acknowledged. The keynote further described how these challenges influence software architectural choices, differentiating between cloud and low-power embedded AI solutions, and provided insights into Bosch's concrete strategies and architectures to address these issues, demonstrating the critical role of AI in industrial innovation and software architecture.

The examples of the state of the practice presented through these keynotes helped frame the remaining talks and discussions. What follows are the abstracts from the attendee talks.

### 3.1 Fifty Shades of Uncertainty

*Nelly Bencomo (Durham University, GB)*

**License** Creative Commons BY 4.0 International license
© Nelly Bencomo

There is growing uncertainty about the environment of software systems. Therefore, how the system should behave under different contexts cannot be fully predicted at design time. It is considerations such as these that have led to the development of self-adaptive systems (SAS), which can dynamically and autonomously reconfigure their behavior to respond to changing external conditions.

The scope of the talk is in the area of Requirements Engineering (RE) and the development of techniques to quantify uncertainty to improve decision-making. The explicit treatment of uncertainty by the running system improves its judgment to make decisions supported by evaluating evidence found during runtime, possibly including the human-in-the-loop. I will also discuss how quantification of uncertainty can be used to improve requirements elicitation (using simulations, for example).

The talk will cover different approaches to quantifying uncertainty and its role in Human-Machine Teaming.

## 3.2 Software Architecture Modeling of Machine Learning Systems: Unsolved Challenge or Old Wine in New Bottles?

*Justus Bogner (Universität Stuttgart, DE)*

**License** Creative Commons BY 4.0 International license
© Justus Bogner

More and more software systems incorporate techniques from machine learning (ML) to support decision-making or automate information processing. Such ML systems still incorporate many traditional components but nonetheless require specialized practices in certain areas. From a software architecture perspective, one question is if existing practices and frameworks are suitable for the effective modeling and documentation of ML systems. While there are many methods for architecture documentation based on concepts like viewpoints, views, and modeling notations, it is still unclear how ML professionals can best apply these practices to model the architecture of ML systems. The data dependency and uncertainty of ML components, ML stakeholder diversity, plus new quality concerns like explainability or model observability might require new types of views and diagrams. While a few specialized approaches have been proposed, other software architecture publications use vastly different notations and diagrams to model ML systems.

In this short talk, I want to lay the foundation for a discussion about this topic. I will briefly talk about challenges in this space and present a few existing approaches. My goal is to discuss if these challenges are different from architecture modeling for traditional systems and if new approaches are needed. Ideally, the discussion might lead to research collaborations in this space to a) better understand challenges and b) create or adapt methods to overcome them.

## 3.3 AI: From Offline & Centralized to Online & Federated

*Jan Bosch (Chalmers University of Technology – Göteborg, SE) and Helena Holmström Olsson (Malmö University, SE)*

**License** Creative Commons BY 4.0 International license
© Jan Bosch and Helena Holmström Olsson

Digitalization is concerned with software, data, and AI as enabling technologies. It changes the company and business ecosystem in which it operates from transactional to continuous, i.e., XOps. This requires fundamental changes in how we work with technology from offline and centralized to online and federated. The talk includes examples from Software Center (www.software-center.se) companies, including automated experimentation, federated learning, reinforcement learning, semi-supervised learning, and related topics.

### 3.4 Predicting Software Performance with Divide-and-Learn

*Tao Chen (University of Birmingham, GB)*

**License** Creative Commons BY 4.0 International license
© Tao Chen
**Joint work of** Tao Chen, Jingzhi Gong
**Main reference** Jingzhi Gong, Tao Chen: "Predicting Software Performance with Divide-and-Learn", in Proc. of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2023, San Francisco, CA, USA, December 3-9, 2023, pp. 858–870, ACM, 2023.
**URL** https://doi.org//10.1145/3611643.3616334

Predicting the performance of highly configurable software systems is the foundation for performance testing and quality assurance. To that end, recent work has been relying on machine/deep learning to model software performance. However, a crucial yet unaddressed challenge is how to cater for the sparsity inherited from the configuration landscape: the influence of configuration options (features) and the distribution of data samples are highly sparse. This talk is based on the paper in which we propose an approach based on the concept of "divide-and-learn", dubbed DaL. The basic idea is that to handle sample sparsity, we divide the samples from the configuration landscape into distant divisions, for each of which we build a regularized Deep Neural Network as the local model to deal with the feature sparsity. A newly given configuration would then be assigned to the right model of division for the final prediction.

Experiment results show that using strong domain knowledge to specialize AI performs significantly better in configuration performance learning.

### 3.5 Analyzing Greenability of Software Architectures for AI Systems: The GAISSA project

*Xavier Franch (UPC Barcelona Tech, ES)*

**License** Creative Commons BY 4.0 International license
© Xavier Franch

This presentation introduces the Spanish GAISSA project ("Towards Green AI-based Software Systems: An Architecture-centric Approach") run by the GESSI research group at UPC. The main hypothesis of GAISSA is that the impact of architectural decisions on environmental sustainability shall be understood, defined, reported, and managed in order to model, develop, and deploy green(er) AI-based systems. After presenting the objectives and vision, the presentation details the work done so far, which is basically a series of empirical studies on the effect of model and system architectures on environmental sustainability, exploring trade-offs with other attributes such as accuracy. Emerging challenges are the consolidation of the results of such studies, and how to boost the attention on sustainability in the community. In this respect, a result of the project is the so-called Energy Label, which categorizes the efficiency level (from A to E) of a trained model, defined in terms of several attributes. The concept has been proven using the HuggingFace repository. As future work, the talk mentions the link between requirements and architectures.

## 3.6 Why Organizations Fail to Implement AI

*Benjamin Klöpper (Capgemini – Stuttgart, DE)*

License Creative Commons BY 4.0 International license
© Benjamin Klöpper

Many organizations, especially enterprises, use a data-lake-based infrastructure to prepare their data for machine learning modeling. Data-lake-based systems ingest raw data from a wide variety of data sources such as databases, legacy information systems, automation systems, and more. The processed data is intended to serve the needs of different users and use cases. A data-lake-based system consists of a complex technology stack that implements distributed data processing and data storage. Typical transformation steps include ingestion, cleansing, enrichment, transformation, and finally, serving.

Typically, a centralized team of data experts runs the infrastructure and implements the data processing. The thinking is centered around different data pipelines.

The monolithic architecture of the data lake and the centralized organization result in a siloed way of working that puts the data team in a difficult position, often leading to unsatisfactory machine learning results and limiting the data culture of the data team. The data team suffers from

- Operational data that is not fit for analytics, and the team lacks a deep understanding of the data
- Frequent and unexpected downstream changes that require the team to patch data pipelines in response
- Time spent searching for and identifying the right data for use cases
- Lack of understanding of business requirements

The resulting monolithic architecture and centralized team setup ignore common software engineering best practices, such as reasonably sized teams of 5-8 people with clear ownership, lack of mechanisms to offload cognitive load, minimizing hand-offs through loose coupling, and clear interfaces between components and teams.

One architectural approach that tries to facilitate the software engineering best practices are data products introduced by Dehgani. They ingest data only from a few source systems or other data products, encapsulate the data transformation and data storage, provide an easy-to-use business-oriented API, and contain a human-readable description and code for observability. Data products group elements of data transformation so that they can be owned by a software team and consumed as a service by other teams.

Essential to the adoption of a data product is a data platform that makes it easy for data product developers to follow the governance policies of the surrounding organization. A key goal of the data platform is developer experience (DevEx). There are several points to learn from a data-product-centric architecture:

- Design software and organizational architecture for ML-enabled systems together.
- Defined components or products should be independently deployable and a team of 5-8 developers should be sufficient to develop them.
- There needs to be some kind of platform to reduce the cognitive load.
- The platform should be built in an evolutionary fashion starting from the thinnest viable platform (TVP).

## 3.7 SPIRA Challenges and Lessons Learned on Architecting an Intelligent System for Respiratory Insufficiency Detection – Notes on Hands-on Education on SA4AI

*Fabio Kon (University of Sao Paulo, BR)*

**License** Creative Commons BY 4.0 International license
© Fabio Kon

Our group at the University of São Paulo has 23 years of experience teaching an advanced software development course called "Agile Software Development Lab." It is extremely popular with students and has been very successful in teaching advanced software development concepts with a 'learn by doing', project-based approach with real external customers. However, many recent projects have dealt with machine learning and AI-based systems, which sometimes significantly increase the complexity of projects. We are then currently remodeling the course to incorporate pedagogical aspects related to AI-based systems development.

In this short talk, I provide an experience report about the development of a complex system for respiratory insufficiency detection that has been developed at the University of São Paulo in a collaboration between data scientists, physicians, linguists, and computer science professors and students. We discuss the challenges and lessons learned in architecting and building the system.

## 3.8 Software Architecture for Machine Learning Systems: Challenges, Practices, and Opportunities

*Grace A. Lewis (Carnegie Mellon Software Engineering Institute – Pittsburgh, US)*

**License** Creative Commons BY 4.0 International license
© Grace A. Lewis

Developing software systems that contain machine learning (ML) components requires an end-to-end perspective that considers the unique life cycle of these components – from data acquisition to model training to model deployment and evolution. While there is an understanding that ML components, in the end, are software components, there are some characteristics of ML components that bring challenges to software architecture and design activities, such as data-dependent behavior, the need to detect and respond to drift over time, and timely capture of logs, metrics, user input and labeled data to inform retraining. In this presentation, I talk about some of these challenges and propose a set of practices to address these challenges, such as co-architecting, the recognition of monitorability as a driving quality attribute, and system-level architecture patterns and tactics. The presentation concludes with a list of opportunities to make progress in software architecture for ML systems, which includes adapting and growing the software architecture body of knowledge for application to ML systems, bringing the software architecture body of knowledge and discipline to model development activities, and positioning software architecture as a unifying activity for system stakeholders (model developers, software engineers, and operations).

## 3.9 BeT (Behavior-enabled IoT)

*Henry Muccini (University of L'Aquila, IT)*

**License** Creative Commons BY 4.0 International license
© Henry Muccini
**Main reference** Henry Muccini, Barbara Russo, Eugenio Zimeo: "The BET project: Behavior-enabled IoT", CoRR, Vol. abs/2307.13186, 2023.
**URL** https://doi.org//10.48550/ARXIV.2307.13186

BeT (Behavior-enabled IoT)[1] is a project that aims to provide a reference architecture, conceptual framework, and related techniques to design behavior-enabled IoT systems and applications. The presentation introduces the concept of the Internet of Behaviors (IoB) and analyzes the human-system bi-causal quality connection effects. BeT analyzes the emerging problem of understanding the mutual influence between QoS (Quality of Service) and QoE (Quality of Experience) in IoB systems, where dynamic changes in the system and surroundings can give rise to nontrivial unforeseen cause-effect mutual relations between system and human behaviors.

## 3.10 Software Architecting in the Era of AI and AI-Augmented Development Tools

*Ipek Ozkaya (Carnegie Mellon Software Engineering Institute – Pittsburgh, US)*

**License** Creative Commons BY 4.0 International license
© Ipek Ozkaya

The boundary between SA4ML and ML4SA is diffusing as more general-purpose AI models are incorporated into systems. As more capable development tools enter the developer's ecosystem the role and responsibilities of the architect will shift [1]. As generative AI and foundation models increase in capabilities, they will help reveal the criticality of architecture knowledge and how it must guide development. As capabilities of large language models (LLM) evolve, however, there is a potential increasing gap between what tools supported by LLM can accomplish realizing narrowly scoped implementation tasks versus how they can do so while being cognizant of software architectural concerns [2]. This talk will emphasize some of the open questions and challenges that software architecture researchers need to address as LLM-supported development tools evolve rapidly. These include, but are not limited to, the following:

- What are the specific design and architectural concerns that need to be addressed when using generative AI?
- Can generative AI tools be used to improve the design and architecture of systems?
- Can these tools provide new features and capabilities that can be used to support the architectural design process?
- Could generative AI tools be used to generate design patterns and tactics, which could then be used to guide the development process?
- Can these tools accelerate the generation of alternative designs and their comparison?
- Can they be used to provide feedback on designs, such as identifying potential risks and issues?

### References

1 Ipek Ozkaya: Application of Large Language Models to Software Engineering Tasks: Opportunities, Risks, and Implications. IEEE Softw. 40(3): 4-8 (2023)
2 Ipek Ozkaya. Can Architecture Knowledge Guide Software Development With Generative AI? IEEE Softw. 40(5): 4-8 (2023)

## 3.11 Architecting Systems to Integrate Machine Learning

*Lena Pons (Carnegie Mellon Software Engineering Institute – Pittsburgh, US)*

**License** 🅭 Creative Commons BY 4.0 International license
© Lena Pons

ML-enabled systems introduce additional concerns to systems, which introduce new complexity into the software architecture of such systems. Design considerations for integrating an ML component into a larger software system may be unfamiliar to software architects. Software engineers and data scientists need to communicate information that crosses domains, which may result in some considerations not being exposed during architecture and design conversations. We present a set of driving quality attributes for machine learning systems and how they align with conventional software engineering practice versus attributes where the integration of an ML component requires new patterns and tactics. We present some questions to assist architects in eliciting better information about design concerns.

## 3.12 Machine Learning and Self-adaptation

*Bradley Schmerl (Carnegie Mellon University – Pittsburgh, US)*

**License** 🅭 Creative Commons BY 4.0 International license
© Bradley Schmerl

Architecture-based self-adaptation is a method for adding a control loop on top of systems to make changes and repairs at run time, with the software architecture as the basis for making decisions. But does this work for systems that have machine learning components in the architecture? We report on work that we are doing to investigate how to apply self-adaptation concepts to manage machine learning components. We report on the set of challenges related to the timeliness of observations, the propagation of uncertainty, and the lack of understood costs, benefits, and impacts of adaptation tactics, but also that the current concepts seem to apply. One of the key concepts in ML is explainability. We report on two aspects of explainability: 1) how we use planning models to support contrastive explanations, and 2) how standard techniques in ML, like principal component analysis, clustering, and decision tree learning, can be used to help explain choices that need to be made by architects, designers, etc. by focusing them on critical qualities/concerns and interactions among them.

## 3.13 A Vision and Challenges about Intelligent and Trustworthy IoT Systems

*Romina Spalazzese (Malmö University, SE)*

**License** 🅭 Creative Commons BY 4.0 International license
© Romina Spalazzese

Today's systems are more and more software/hardware intensive, very large, heterogeneous, data-driven, dynamic, evolving, intelligent, and involve humans. These characteristics make their engineering and architecture challenging tasks. Additionally, different applications have different desired quality characteristics concerning, e.g., response time, power consumption, interoperability, privacy, and trust, which influence their design. Examples of such modern complex systems are ML-enabled Internet of Things (IoT) systems that have to deal with many challenges at the same time, both from a human and technical perspective.

An overarching question in our ongoing research project is “How should Intelligent and Trustworthy IoT systems be designed?” and some of the more specific questions are “How and when should AI (in particular ML) be used to realize such systems?” as well as “How could and when should edge computing (including hybrid edge-cloud processing) be used to realize such systems?”. Towards identifying architectural approaches, patterns, and styles, as well as analysis models, metrics, and techniques for ML-enabled systems, I will discuss a vision and identify challenges related to an ML-enabled collaboration paradigm called IoT-Together.

## 3.14 SA, ML and Patterns

*Anastas Stoyanovsky (Amazon – Pittsburgh, US)*

**License** Creative Commons BY 4.0 International license
© Anastas Stoyanovsky

The introduction of machine learning has not changed anything fundamental about architectural thinking or methods. What is new is the confluence of particular quality attributes, which effectively reduces the discussion to a matter of practical execution. Several high-level patterns, both architectural and organizational, are discussed as foundations for further inquiry. Industry experience points to challenges common to both industry and academia, leading to a call to action around education.

## 3.15 Software Architecture Meets Machine Learning: A Tale of Convergence

*Karthik Vaidhyanathan (IIIT – Hyderabad, IN)*

**License** Creative Commons BY 4.0 International license
© Karthik Vaidhyanathan
**Main reference** Henry Muccini, Karthik Vaidhyanathan: “Software Architecture for ML-based Systems: What Exists and What Lies Ahead”, in Proc. of the 1st IEEE/ACM Workshop on AI Engineering – Software Engineering for AI, WAIN@ICSE 2021, Madrid, Spain, May 30-31, 2021, pp. 121–128, IEEE, 2021.
**URL** https://doi.org//10.1109/WAIN52551.2021.00026

Over the years, software systems have become increasingly complex due to their ever-increasing pervasive nature, resulting in various architecting challenges to ensure better performance, reliability, etc. On the other hand, machine learning (ML) has advanced rapidly with advancements in infrastructure, data availability, etc. However, the growing adoption of ML has given rise to challenges associated with development practices, deployments, ensuring data quality, etc., in addition to the challenges of a traditional software system. These challenges have resulted in the convergence of SA and ML, resulting in two broad research areas: i) ML4SA and ii) SA4ML.

This talk presents an overview of this convergence. The presentation then elaborates on the two research lines and the challenges the community must address going forward. The first part (ML4SA) focuses on using ML techniques to enable software systems to autonomously adapt and improve their architecture at run time to guarantee a better quality of service at run time and then using generative AI for architectural knowledge management for design-time aid for architects. The second part (SA4ML) elaborates on the challenges in architecting ML-enabled systems and further provides some insights on our recent work related to an agile methodology for ML-enabled systems [1] and self-adaptive architecture for ML-enabled systems [2].

**References**

1 K. Vaidhyanathan, A. Chandran, H. Muccini and R. Roy, "Agile4MLS—Leveraging Agile Practices for Developing Machine Learning-Enabled Systems: An Industrial Experience," in IEEE Software, vol. 39, no. 6, pp. 43-50, Nov.-Dec. 2022, doi: 10.1109/MS.2022.3195432.

2 K.Vaidhyanathan "Data-Driven Self-Adaptive Architecting Using Machine Learning," Ph.D. dissertation, Gran Sasso Science Institute, Italy, 2021, online at: http://hdl.handle.net/20.500.12571/15976

## 3.16 Generative AI at Fraunhofer and Research Roadmaps from the Software Architecture Community

*Ingo Weber (TU München – Garching, DE)*

**License** 🅭 Creative Commons BY 4.0 International license
© Ingo Weber

Fraunhofer is the world's largest organization focused on application-oriented research. Several of its institutes are at the forefront of research in generative AI, but the recent advances make it a topic relevant for all industry sectors and hence potentially all Fraunhofer customers and institutes. Accordingly, Fraunhofer is broadening and deepening its AI endeavours. One early step was the introduction of fhGPT, an AI chatbot based on Azure OpenAI GPT models. Furthermore, from the ICSA-lite conference 2022, a book of high relevance for this Dagstuhl seminar emerged, which is currently in print: Patrizio Pelliccione, Rick Kazman, Ingo Weber, Anna Liu, editors. Software Architecture – Research Roadmaps from the Community. Springer, 2023.

## 3.17 Building, Engineering & Operating Systems for Critical Infrastructure

*Roland Weiss (ABB – Mannheim, DE)*

**License** 🅭 Creative Commons BY 4.0 International license
© Roland Weiss

Control systems for process automation plants run business and mission critical infrastructure. They span various verticals, from chemical to power plants, from hydrogen generation to waste incineration. These applications require flawless execution 24 by 7, as malfunctions in the system can endanger people, the environment, or business targets. On the other hand, engineering such systems as well as operating them could benefit tremendously from the introduction of machine learning and artificial intelligence in general. In my talk, I explore the application of ML and AI in such systems.

In the engineering phase, generative AI has proven very promising in generating automation code (in IEC 61131 structured text). Based on a specification entered via a prompt interface, the code inside function blocks could be generated. Likewise, process graphics were translated into rules from older systems. Then, trained models validated the results and highlighted areas that deviated from the expected outcome. When running a process plant, operators have to monitor the system and, in the case of deviations or emergencies, interact with the system to either bring it to a safe state or to the desired stable state. Today, this requires very detailed domain know-how and extensive training. We are now exploring supporting the operators with ML/AI-based assistance systems. For example, in case of deteriorating product quality in a chemical facility, the operator can pull up similar occurrences from the past and get recommendations from the assistance systems on how to handle the situation.

The key challenges we are facing when introducing AI/ML components into control systems are the following:

- The behavior of the system must be explainable. In regulated and safety-critical environments, the system providers need to be able to provide detailed RCAs in case of incidents.
- Training data is hard to get for various reasons. The collection of the data can't interfere with the process itself. Also, the incidents are typically rare, and operators aim to avoid them in the first place.
- Last but not least, these systems have extremely long lifetimes, thus, introducing ML means bringing them into legacy systems with traditional development processes.

We are looking forward to tackling these research challenges with Academia to contribute to the collection of best practices for building systems for critical infrastructure in order to make the world a safer and more sustainable place.

## 3.18 Challenges of Integrating ML Models in Safety-Relevant Architectures

*Marc Zeller (Siemens – München, DE)*

License Creative Commons BY 4.0 International license
© Marc Zeller

Traditional automation technologies alone are not sufficient to enable the fully automated operation of trains. However, Artificial Intelligence (AI) and Machine Learning (ML) offer great potential to realize the mandatory novel functions to replace the tasks of a human train driver, such as obstacle detection on the tracks. The problem, which still remains unresolved, is to find a practical way to link AI/ML techniques with the requirements and approval processes that are applied in the railway domain. The safe.trAIn project aims to lay the foundation for the safe use of AI/ML to achieve the driverless operation of rail vehicles and thus addresses this key technological challenge hindering the adoption of unmanned rail transport. The project goals are to develop guidelines and methods for the reliable engineering and safety assurance of ML in the railway domain. Therefore, the project investigates methods to reliably design ML models and to prove the trustworthiness of AI-based architectures, taking robustness, uncertainty, and transparency aspects of the AI/ML model into account.

## 3.19 Software Architecture for Foundation Model-Based Systems

*Liming Zhu (Data61, CSIRO – Sydney, AU)*

License Creative Commons BY 4.0 International license
© Liming Zhu

With the successful implementation of Large Language Models (LLMs) in chatbots like ChatGPT, there is growing attention on foundation models (FMs), which are anticipated to serve as core components in the development of future AI systems. Yet, systematic exploration into the design of foundation model-based systems, particularly concerning risk management, trust, and trustworthiness, remains limited. In this talk, I propose the challenges and initial approaches in both architecting FM-based systems and how FMs have an impact on software engineering. I point to some initial directions, such as architecting as a process of understanding (rather than designing/building), designing guardrails (rather than quality attributes), and radical observability.

## 4 Working groups

The seminar provided an extensive and insightful exploration into the four key quality attributes that were prioritized by the attendees as critically important while architecting ML-enabled systems: *Data Centricity, Evolvability, Observability and Uncertainty, Trust and Trustworthiness*. The four working group (WG) discussions focusing on each quality attribute concern provided input for formulating the research agenda and key challenges to address. In this section, we describe the scope of each working group, and in the Seminar Overall Results section, we summarize the initial research roadmap and open challenges.

### 4.1 WG1: Architecting for Data Centricity

*Jan Bosch (Chalmers University of Technology – Göteborg, SE), Benjamin Klöpper (Capgemini – Stuttgart, DE), Ipek Ozkaya (Carnegie Mellon Software Engineering Institute – Pittsburgh, US), Lena Pons (Carnegie Mellon Software Engineering Institute – Pittsburgh, US), and Christoph Schröer (Universität Oldenburg, DE)*

**License** Creative Commons BY 4.0 International license
© Jan Bosch, Benjamin Klöpper, Ipek Ozkaya, Lena Pons, and Christoph Schröer

The main purpose of the WG1 group on "Data Centricity" was to explore and address the pivotal role of architectural elements that extract, transform, load, store, and share data in the architectures of ML and analytics systems. The group focused on examining the entire data life cycle – from acquisition to transformation to runtime consumption – and understanding how data is integrated and managed by an ML-enabled system throughout this data life cycle.

The discussions emphasized that effective data management throughout the data life cycle is crucial for the success of ML-enabled systems. This includes recognizing the challenges associated with the complexities of data architectures. The WG1 group also identified the need for a distinct architectural view that focuses specifically on data, given its critical importance in ML-enabled systems. They acknowledged that while current software architecture practices and architectures of non-ML systems address data processing and storage, there is a need for a more focused approach to identify strategies for addressing data-centricity to cater to the unique requirements and challenges posed by ML systems.

In summary, the WG1 group's discussions emphasize the centrality of data in ML systems and the need for specialized and clearly communicated architectural approaches to manage and utilize data in these systems effectively.

## 4.2 WG2: Evolvability

*Justus Bogner (Universität Stuttgart, DE), Jan Bosch (Chalmers University of Technology – Göteborg, SE), Helena Holmström Olsson (Malmö University, SE), Henry Muccini (University of L'Aquila, IT), Raghu Reddy (IIIT – Hyderabad, IN), Anastas Stoyanovsky (Amazon – Pittsburgh, US), Ingo Weber (TU München – Garching, DE), and Liming Zhu (Data61, CSIRO – Sydney, AU)*

**License** Creative Commons BY 4.0 International license
© Justus Bogner, Jan Bosch, Helena Holmström Olsson, Henry Muccini, Raghu Reddy, Anastas Stoyanovsky, Ingo Weber, and Liming Zhu

The primary objective of the WG2 group on "Evolvability" was to address the unique challenges in developing ML-enabled systems that can evolve efficiently in response to changing requirements and environmental factors. The group's primary focus was on understanding the dynamics of ML components within these systems, particularly how they interact with non-ML components and Machine Learning Operations (MLOps) infrastructures and practices.

A key finding of WG2 was the critical role of MLOps in ensuring the evolvability of ML components within ML-enabled systems. MLOps, supporting the entire life cycle of ML components from development to deployment and maintenance, emerged as a pivotal factor in managing and automating the continuous evolution of these systems. The group identified that the integration of MLOps within existing DevOps practices is not only essential but also challenging, given the unique characteristics of ML components.

Another significant aspect highlighted by WG2 was the concept of technical debt in ML systems. ML-enabled systems bring forth new classes of design and architectural challenges related to degradation in data or models as technical debt, which can impede the maintenance and evolution of ML-enabled systems. Addressing these forms of technical debt is crucial for sustaining the long-term viability and evolvability of these systems.

WG2 also emphasized the importance of the skill sets and roles of architects designing and maintaining ML-enabled systems. Architects need to possess a comprehensive understanding of aspects such as data engineering, model engineering, and ethical considerations to effectively manage the evolution of these complex systems.

In essence, the WG2 group underscores the complexities involved in creating evolvable ML-enabled systems and highlights the importance of MLOps, technical debt management, and specialized architectural knowledge in addressing these challenges.

## 4.3 WG3: Observability and Uncertainty

*Nelly Bencomo (Durham University, GB), Xavier Franch (UPC Barcelona Tech, ES), Fabio Kon (University of Sao Paulo, BR), Ipek Ozkaya (Carnegie Mellon Software Engineering Institute – Pittsburgh, US), Marie Platenius-Mohr (ABB – Ladenburg, DE), Bradley Schmerl (Carnegie Mellon University – Pittsburgh, US), Roland Weiss (ABB – Mannheim, DE), and Karthik Vaidhyanathan (IIIT – Hyderabad, IN)*

**License** Creative Commons BY 4.0 International license
© Nelly Bencomo, Xavier Franch, Fabio Kon, Ipek Ozkaya, Marie Platenius-Mohr, Bradley Schmerl, Roland Weiss, and Karthik Vaidhyanathan

The working group on "Observability and Uncertainty" primarily focused on how uncertainty and observability were intertwined and how they impact the design, implementation, and operation of ML-enabled systems.

A key finding of the group was the often inherent non-deterministic and statistical nature of ML techniques, which introduces a fundamental level of uncertainty in these systems. This uncertainty is not just a challenge but a necessary aspect to consider in the architecture of ML-enabled systems. The group discussed two main ways to handle uncertainty: (1) reducing or eliminating uncertainty through strategies similar to using redundancy in reliability engineering and (2) managing uncertainty, which includes strategies for dealing with incomplete information and uncertain components.

The group's discussions underscored the close relationship between uncertainty and observability. Observability – the ability to infer the internal state of a system based on its output – is crucial for understanding and managing uncertainty in ML systems. The degree of uncertainty dictates the requirements for observability, impacting various other aspects of software development, including trustworthiness and evolvability.

In summary, WG3's outcomes emphasize the importance of treating uncertainty as a primary concern in the design and implementation of ML-enabled systems. In addition, the discussions highlight the need for effective observability mechanisms to identify and manage the inherent uncertainties of ML-enabled systems to ensure their reliability and effectiveness.

## 4.4 WG4: Architecting for Trust and Trustworthiness

*Tao Chen (University of Birmingham, GB), Thomas Kropf (Robert Bosch GmbH – Renningen, DE), Grace A. Lewis (Carnegie Mellon Software Engineering Institute – Pittsburgh, US), Henry Muccini (University of L'Aquila, IT), Alex Serban (Siemens Healthineers, Erlangen, DE & University Transilvania of Brasov, RO), Romina Spalazzese (Malmö University, SE), and Marc Zeller (Siemens – München, DE)*

**License** Creative Commons BY 4.0 International license
© Tao Chen, Thomas Kropf, Grace A. Lewis, Henry Muccini, Alex Serban, Romina Spalazzese, and Marc Zeller

The main purpose of WG4 was to explore the key aspects and challenges in designing trustworthy ML systems, such as in safety-critical and autonomous contexts.

The outputs of the discussions in this group emphasize the distinction between trust and trustworthiness and various aspects of these concepts in ML-enabled systems, including state-of-the-art trustworthy AI, the role of trust as an element of overall responsible AI practices and their relevance to ML-enabled system development, and the engineering of trustworthy ML-enabled systems. It also examines the importance of providing evidence of trustworthiness tailored to different stakeholders for calibrated trust, such as through certification labels and empirical studies, and discusses enabling the design of trustworthy ML-enabled systems through architectural models and toolkits.

A significant portion of the discussions were dedicated to addressing the challenges in evaluating the trustworthiness of ML-enabled systems and the role of human-in-the-loop in ensuring trustworthiness. The discussions in this group focused on the trust and trustworthiness of ML-enabled systems and had a close relationship with considerations that apply to AI systems in general, such as responsible AI practices. The discussions, in particular, emphasized the need for explainable AI as a key component in building trust in ML-enabled systems.

In summary, WG4's observations emphasize the critical need for defining trust and trustworthiness clearly in ML-enabled systems, exploring the complexities of achieving the several qualities essential in building trust and trustworthiness and developing and extending frameworks and strategies to guide the development of trustworthy applications with ML components.

# 5 Open problems

The SA&ML research roadmap, which is one of the key outcomes of this Dagstuhl Seminar, is heavily influenced by the high-priority quality attributes that were discussed in depth in the working groups and their relationships. Figure 1 provides an overview of the relationships between the quality attributes and the interplay that exists between them, with Data Centricity being a cross-cutting quality attribute. It also shows Explainability and Adaptability as two additional key quality attributes that contribute to these relationships. The following subsections provide a summary of how each attribute was scoped by the working groups, open problems in each area, and a list of high-priority research areas that need to be explored in the near term to advance the practice of SA&ML.

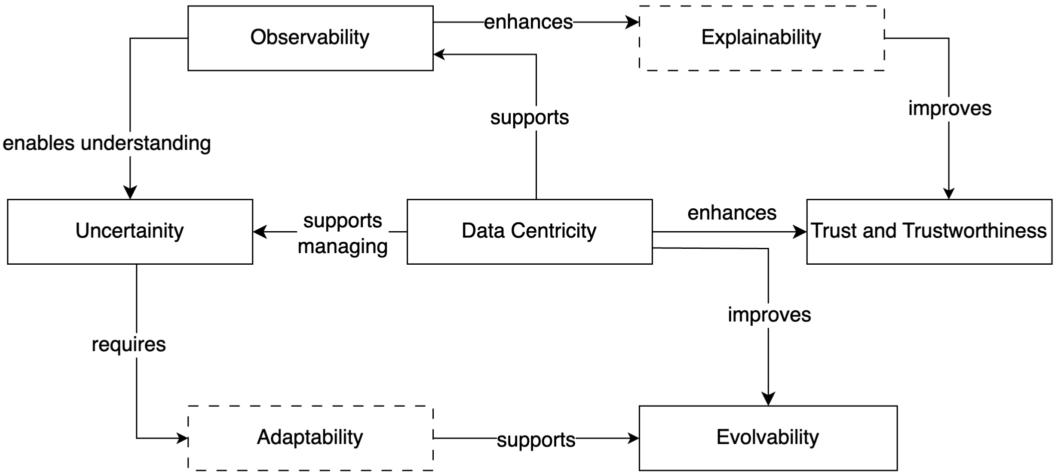


**Figure 1** Relationships between the different quality attributes.

## 5.1 Data Centricity

Data-centricity emphasizes the significance of data quality and management for ML-enabled systems. High-quality, well-managed data is crucial for the development of reliable ML models and for maintaining the trustworthiness of the system.

Data-centricity also impacts evolvability and adaptability, as systems that can adapt to changes in data or requirements are more sustainable over time.

### 5.1.1 Challenges and Open Problems

The seminar looked beyond basic data integrity and security issues, delving into more complex and ML-relevant issues such as ensuring data relevance over time, managing large and diverse datasets, and maintaining compliance with evolving data privacy laws and standards. The future directions suggested include developing data views that reflect architecturally-significant concerns, understanding architectural implications of online vs offline learning, and data decoupling strategies.

## 5.2 Uncertainty and Observability

**Uncertainty** is a fundamental aspect of ML-enabled systems due to the often non-deterministic and statistical nature of ML models. Uncertainty can arise from the quality of training data, the model design, interactions between ML and other components, and the environment in which the system operates.

**Observability** involves the ability to monitor the internal states and outputs of the system to enable understanding and make informed decisions about its functioning, especially in the face of uncertainty. Observability helps system developers and operators understand and manage uncertainty.

Observability can also contribute to enhancing explainability, as more observable systems provide more information that can be used to explain inference results to users and for engineers to diagnose systems. Observability also contributes to evolvability as runtime data can be used to determine, for example, when a model needs to be retrained in response to drift.

### 5.2.1 Challenges and Open Problems

The major challenge identified by the working group participants lies in quantifying and managing uncertainty, which is inherent in ML-enabled systems due to their often non-deterministic nature. Another challenge is developing methods for effective observability that can aid in explaining the behavior of ML systems. Future research is required to enhance methods for quantifying uncertainty, properly propagating uncertainty throughout the life cycle and system, developing advanced observability techniques, and exploring the impact of uncertainty on trustworthiness and evolvability. There is also an emphasis on creating strategies for adaptive actions in response to observed properties, including uncertainties.

## 5.3 Evolvability (and Adaptability)

**Evolvability** was defined by the seminar participants as a system's ability to adapt to changes in expectations (i.e., requirements), while adaptability refers to the ability of the system to handle (or be modified to deal with) changes in its environment.. Evolvability, therefore, is mostly a design-time concern, whereas **adaptability** is both a design-time and runtime concern. These attributes are essential for designing ML-enabled systems, as these systems must continuously evolve and adapt to remain effective. Evolvability and adaptability are influenced by observability and uncertainty, and together, they contribute towards maintainability; by addressing these attributes proactively and intentionally, systems can be designed to evolve and adapt more effectively. Evolvability and adaptability are also closely linked to data centricity, as the system's ability to evolve and adapt is dependent on the quality and relevance of the data it uses and how its evolution and adaptation are influenced by changes in data.

### 5.3.1 Challenges and Open Problems

Participants unanimously agreed that the main challenge in evolvability and adaptability lies in integrating continuous development and deployment practices for ML components, such as MLOps, with software development and deployment of non-ML components. Addressing the accruing technical debt and managing quality attributes of ML components are also highlighted as significant challenges. Further, participants suggest exploring effective integration of MLOps practices, identifying new evolvability and adaptability-specific technical debt items in ML-enabled systems, and conceptualizing and realizing architectural designs that support continuous change. Developing guidelines for managing the interaction between ML and non-ML components is also seen as a key future direction where more case studies and research need to be developed.

## 5.4 Trust and Trustworthiness

**Trust** is the subjective confidence stakeholders place in a system's quality. It is influenced by the system's trustworthiness.

**Trustworthiness** involves the system's objective ability to consistently perform according to quality expectations and ethical standards. It is bolstered by explainability, as stakeholders are more likely to trust a system they can understand. Trust and trustworthiness are impacted by data centricity, as the quality and management of data underpin the reliability and robustness of ML-enabled systems. Related, if the system provides better explainability, it fosters trust in the system. A system that is trustworthy (through evidence of achieving quality attributes such as security, reliability, and fairness) merits the trust placed in it by users.

### 5.4.1 Challenges and Open Problems

The participants elaborated on the challenges of balancing system transparency and explainability while complying with ethical standards and legal regulations in different domains. Proposed future research directions included the development of comprehensive frameworks for building trustworthy AI systems, tailoring trustworthiness evidence to different stakeholders for calibrated trust, exploring human-in-the-loop approaches to balance automation and human oversight, and creating dynamic processes for evolving trust and trustworthiness in line with technological and societal shifts.

## 5.5 High Priority Research Areas to Advance SA&ML

Several open research areas in SA&ML were identified as a result of the seminar discussions. High-priority research areas requiring near-term focus for advancing the state of the practice include the following:

- Architectural Design for Data-Centricity: Identifying and developing architectural approaches that effectively address the central role of data in ML-enabled systems, including data acquisition, processing, and management, as well as designing systems that can adapt to changes in data characteristics, need attention from researchers as well as practitioners in documenting their lessons learned. Existing architecture patterns, tactics, and quality attributes related to data architecting need to also be shared more consistently, and gaps need to be filled where applicable.
- Evolvability and adaptability of MLOps Architectures: Creating architectural designs that support the evolvability and adaptability of ML-enabled systems, particularly focusing on leveraging MLOps infrastructure, is not common practice. Challenge areas include addressing the lifecycle management of ML models and ensuring systems can evolve and adapt to new and changing requirements and environments.

- Uncertainty as a First-Class Concern: Integrating the management of uncertainty into the architectural design process for ML-enabled systems first requires identifying sources of uncertainty which can effectively be managed with architectural approaches. This involves developing methods for quantifying and mitigating uncertainty that can be leveraged by architecture practices and constructs.
- Observability in ML-enabled Systems: Existing approaches in system observability need to be enhanced to better observe and manage the behavior of ML components. This includes developing metrics and tools for monitoring and understanding ML components and ML-enabled system states.
- Trust, Trustworthiness, and Ethical Considerations: Building trustworthiness into the architecture of ML-enabled systems, including designing for ethical considerations, compliance with regulations, and ensuring transparency and explainability of AI decisions, is a growing need. Identifying what aspects of trustworthiness and ethical considerations can be handled architecturally and clearly communicating the remaining gaps is critical as research in this area progresses.
- Human-in-the-Loop AI Decision Making: Architecting systems that effectively incorporate human oversight and interaction, particularly in critical decision-making processes to ensure balanced autonomy and control in AI systems, will be an area of research that influences many quality attribute concerns, as well as how responsibilities are allocated to architectural elements, potentially resulting in new architecture patterns and tactics.

## 5.6 Follow-up Work

At the seminar, participants recognized that there is a lack of understanding and common definitions for many concepts related to ML-enabled systems and existing architecture practices and fundamentals. Furthermore, in addition to the set of key quality attributes discussed in detail and shared in this report, other quality attributes that require attention from the software architecture community were identified. Hence, a more detailed description of the concepts identified, extended discussions around the quality attributes discussed, as well as a more comprehensive explanation of the research roadmap, are planned for publication in appropriate venues.

## Participants

- Nelly Bencomo
Durham University, GB
- Justus Bogner
Universität Stuttgart, DE
- Jan Bosch
Chalmers University of Technology – Göteborg, SE
- Tao Chen
University of Birmingham, GB
- Xavier Franch
UPC Barcelona Tech, ES
- Helena Holmström Olsson
Malmö University, SE
- Benjamin Klöpper
Capgemini – Stuttgart, DE
- Fabio Kon
University of Sao Paulo, BR
- Thomas Kropf
Robert Bosch GmbH – Renningen, DE
- Grace A. Lewis
Carnegie Mellon University – Pittsburgh, US
- Henry Muccini
University of L'Aquila, IT
- Ipek Ozkaya
Carnegie Mellon University – Pittsburgh, US
- Marie Platenius-Mohr
ABB – Ladenburg, DE
- Lena Pons
Carnegie Mellon University – Pittsburgh, US
- Raghu Reddy
IIIT – Hyderabad, IN
- Bradley Schmerl
Carnegie Mellon University – Pittsburgh, US
- Christoph Schröer
Universität Oldenburg, DE
- Alex Serban
Siemens Healthineers, Erlangen, DE & University Transilvania of Brasov, RO
- Romina Spalazzese
Malmö University, SE
- Anastas Stoyanovsky
Amazon – Pittsburgh, US
- Karthik Vaidhyanathan
IIIT – Hyderabad, IN
- Ingo Weber
TU München – Garching, DE
- Roland Weiss
ABB – Mannheim, DE
- Marc Zeller
Siemens – München, DE
- Liming Zhu
Data61, CSIRO – Sydney, AU