

Hybrid Neural Network and Physics-Based Digital Twins for Condition-Based Maintenance

Andrew Lammas, Russell Brinkworth, Karl Sammut
Luke Voigt, Giselle Rampersad, and Nicholas Brealey

Flinders University Tonsley
Adelaide SA
AUSTRALIA

Defence Science and Technology Group Edinburgh
Adelaide SA
AUSTRALIA

andrew.lammas@flinders.edu.au, russell.brinkworth@flinders.edu.au, karl.sammut@flinders.edu.au,
luke.voigt@defence.gov.au, giselle.rampersad@flinders.edu.au, nicholas.brealey@defence.gov.au

ABSTRACT

This paper presents a case study of the challenges faced in developing a reliable, robust, and accurate digital twin system for an automotive shock absorber. Specifically, this digital twin system's role is to estimate the current gas pressure in the reservoir chamber and compare it with the expected pressure. It is quantitatively demonstrated that design choices of sensors and algorithms have a significant effect on the accuracy of the system which is not proportionate to the hardware costs of the digital twin system. The evaluated sensor suites cost is significant, with an overall cost ranging from A\$297 to A\$4,292, representing a 14-fold difference in costs. The study shows that the use of an expansive and costly sensor suite does not necessarily reflect proportionately in the accuracy of the system. The algorithms and sensors utilised in the digital twin architecture have a significant effect on the accuracy of the system with the RMSE ranging from 3.83 Bar to 0.85 Bar, a four-fold variation in accuracy. The digital twin approach showed significant benefit in accuracy highlighted by the most accurate sensor only approach achieving a RMSE of 2.27 compared to the 0.84 of the full digital twin approach. The lowest cost system which maximally utilised Bayesian methods and physical modelling generated the second most accurate estimate with a RMSE of 1.4 Bar, 165% of the most accurate system, which is still effective for the task, but at 7% of the cost. This demonstrates that by leveraging algorithmic development in a hybrid architecture, performance can be significantly improved and both dataset sizes and training times for the neural network components can be significantly reduced.

1.0 INTRODUCTION

Determining the health and performance of complex systems is challenging, however the use of real-time operational reference systems, or digital twins, can aid in this endeavour [1]. As such, digital twins are increasingly being employed for condition-based maintenance, performance-based asset management, fault detection and performance prediction applications [2]. Such innovative and revolutionary digital twin technology can be used for military fleet sustainment to ensure platform safety, providing better information on the condition of each asset, reducing maintenance costs, and increasing platform availability by enabling better maintenance planning and reducing the occurrence of unanticipated damage.

In our study we evaluated the design choices available in developing a digital twin and examined the consequences of these design choices on hardware costs, development time, and accuracy of the developed system. For our case study, we developed a digital twin capable of estimating, in real time, the static gas pressure in a reservoir chamber of a shock absorber. This estimate can then be used to determine the current

health of the shock absorber by comparing the current pressure with the expected pressure, with any differences potentially indicating the loss of gas volume. For the purposes of this paper, we will look at the estimated static gas pressure as a measure of the accuracy of the digital twin system.

2.0 CHALLENGES WITH DEVELOPING DIGITAL TWINS

Employing digital twins imposes upfront costs, such as software development and sensor integration, and therefore, reducing these costs increases system affordability and adoption. While software development costs may be significant, such costs can be amortised across many units. The same, however, does not hold for the sensors themselves. As such, we will demonstrate the development of a digital twin that balances the accuracy and effectiveness of the system with reduced sensor hardware costs.

2.1 Structured vs Unstructured

Much of the focus of digital twin development to-date has been on high-level modelling of whole systems (e.g., entire vehicles) or end-to-end manufacturing processes. At the component level of service-phase digital twins (i.e., for operational and sustainment rather than manufacturing), the emphasis has been on either physical modelling of the system or the architecture of neural networks that correlate and infer system states based on available measurements.

The terms structured and unstructured represent the extremes of potential architectures. A classical estimation (CE) approach is an example of a fully structured architecture and would employ explicitly developed models and estimators to perform the task of estimation. Conversely, a fully connected End-to-End (E2E) Deep Neural Network (DNN) is an example of a completely unstructured approach. The middle ground between these extremes includes systems that employ a hybrid approach [3]. The CE has the benefit of predictable performance of the system within a range of operational conditions, however, it has a disadvantage in that the system can only incorporate features explicitly encoded into the system. The DNN has the advantage that it can potentially encode any information embedded in the data streams from the sensors without explicit descriptions of these features, however, the quality of this encoding is dependent on the sufficiency of the training dataset. This defines the contrast between these two extremes; CE maximum design effort and predictability, with minimum data and computation requirements; E2E-DNN minimum design effort and predictability, and maximum data and computation requirements.

There has been much hype regarding the potential capability of DNNs, but they place the burden of work on the data collection and processing steps, which depending on the required levels of accuracy and/or robustness, may be significant or even impractical.

2.2 Deployability in the Field

As mentioned in Section 2.1, the limitation of an unstructured system is the limited predictability of the performance of the system when presented with novel data [4]. This is of particular concern when operating in the field where the operational environment cannot be controlled, and the potential permutations of sensors states are predominantly unconstrained. Combining the previous two features means that the data and processing required to achieve a given performance and reliability are generally significantly higher, or possibly unfeasible, for digital twins operating in the field.

Due to this limitation, the utility of DNN systems is limited when it comes to mission critical systems where operating boundaries, and the probability and modes of failure, need to be quantified to the highest level of confidence possible. One approach to make DNNs more trusted is to make the internal parameters generated reflect real world values. There are various ways to achieve this, such as constraining connectivity to neighbouring nodes in the preceding layers, as in Convolutional Neural Networks (CNNs) [5] or forcing some nodes of the architecture to generate an explicit real-world value as in hybrid approaches.

3.0 DIGITAL TWIN APPROACHES

While DNNs are in principle capable of producing the desired performance and robustness, given sufficient training, the training process can be highly time consuming and demanding in terms of the data set. There is a balance to be found between the structured and unstructured approaches to leverage the ability of DNNs to extract unidentified features in a given data stream. This has been employed explicitly in the case of shock absorbers via various approaches either using splines [6] or piecewise linear functions [7] for averaged behaviour with residuals from this developed model estimated by a DNN or alternatively modelling an individual component using a DNN [8]. This allows DNNs to be applied to as small as possible part of the solution, to minimise the size of the training dataset necessary to achieve the required performance and robustness, while maximising the confidence in the behaviour of the DNN. A drawback in the previous examples is by using the data driven analytical modelling approach no insight to the functioning of individual components can be derived.

As mentioned previously, a structured approach would entail explicitly describing the relationships between input, output, and intermediate states either by explicitly modelling sub-components or utilising constrained polynomial as previously demonstrated [6], [7]. When modelling a system as a set of subsystems, equations are developed to quantify the exchange and conversion of quantities between these states. These quantities and states are generally dependent on the function of the system of interest. Since a shock absorber performs a physical task of exerting force on a moving component by converting the energy to heat, a classical approach would be to utilise a physics-based modelling approach and model the components that contribute to the functioning of this conversion process. This is performed by developing equations that describe the inter-relationship between individual damping components, under differing operation conditions, the process of conversion to heat and the propagation of this heat through the system to the surrounding environment.

We illustrate this concept in our case study where we developed a digital twin of a shock absorber for a land-based vehicle by focusing on the individual mechanical components of the shock absorber. The digital twin thus developed infers the current gas pressure in the reserve chamber; something that is vital for functionality, but difficult to measure directly in-situ. This pressure estimation was based on immediately available sensor information, and model-based information that can be derived from this data, parsed to a neural network to perform the inference. Although neural networks can theoretically learn (almost) any arbitrary pattern, the size of the network and the training data required can limit this in practice. One way to overcome these practical limitations is to pre-process the data, hence outsourcing some of the processing that the network would otherwise have to learn. This additional data can then be used as an augmented input to the network, thereby reducing the network size, increasing accuracy and/or decreasing training time.

4.0 EXPERIMENTAL SETUP

The development of the digital twin required a dataset to validate the empirical models and the estimation systems, as well as to train the neural network. The experimental setup consisted of actuating the shock absorber on a dynamometer and collecting data that could be directly or indirectly related to the operational state of the shock absorber. This could be reasonably expected to be analogous to an instrumented shock absorber that is deployable in the field.

4.1 Dynamometer Testing

The testing phase consisted of actuating the shock absorber on a dynamometer at three frequencies: 1, 2, and 4 Hz. These differing frequencies in combination with a fixed peak to peak displacement of 80 mm generates different levels of excitation i.e., velocity, and acceleration. These excitations are repeated for varying gas pressures in the reserve chamber: 0.5, 4, 7, 10, and 12 (nominal) Bar to emulate different stages of gas leakage.

For each of these scenarios, the test started when the shock absorber was within 2°C of ambient temperature and was continued until the maximum normal operating temperature, 80°C, was reached, or in the case of the 1 Hz scenario when the temperature reached equilibrium, at around 60°C. This facilitated the collection of data over a typical range of excitations, temperatures, and in the case of a gas leak failure an expected range of pressures.

Figure 1 shows the dynamometer system used to excite the shock absorber to generate the above scenarios in addition to a diagram that highlights the configuration of sensors used.




Figure 1: Dynamometer test setup and instrumentation diagram Note: the shock absorber in the image was not the device under test.

4.2 Data Collected

For each of the scenarios, presented in Section 4.1, multiple streams of data were captured from the sensor configuration, shown in Figure 1, and is presented in Table 1. This data collection started 1 minute prior to the start of each scenario and continued till the termination condition was met, as specified in Section 4.1.

Table 1: Sensors and Sensed Information.

Value	Rate (Hz)	Sensor	Notes
Position	100	suspension travel sensor	linear potentiometer
Acceleration	1000	2x 8G G-Force sensors	difference between accelerometers
Force	50	20 tonne loadcell	strain gauge
Shock Temperature	1	thermocouple	K-Type
Ambient Temperature	1	thermocouple	K-Type

5.0 ALGORITHM DEVELOPEMENT AND NETWORK TRAINING

To augment the measurements supplied to the DNN, physics-based modelling and Bayesian estimation were used to generate additional data from the stream of available sensor readings. Using this approach, three additional sets of information were generated:

- The velocity of the shock absorber using a Kalman Filter [9].
- The nominal behaviour of the shock absorber, using fluid dynamic principles [10]:
 - Force exerted.
 - Heat generated.

5.1 Velocity Estimation

Velocity is an essential parameter since the behaviour of a shock absorber is dominated by dissipative forces dependent on the velocity. Thus, an accurate measure of velocity would provide a useful metric of the shock absorber's performance and provide critical information to the physics-based model. Directly measuring the velocity of the shock absorber arm is, however, difficult. In contrast, position and acceleration sensors are relatively inexpensive, easy to install and readily available. An algorithm was selected to explicitly calculate the velocity, as the relationship between position, velocity, and acceleration is well known and can be solved optimally using a small set of linear algebraic equations, i.e., a Kalman Filter.

Essentially, the state vector, i.e., the values of interest, to be estimated are position, y_t , & velocity, \dot{y}_t , of the shock absorber extension, and accelerometer bias, b_a , of the accelerometer sensors. These are estimated by correlating the position sensor, z_t , and accelerometers, $a_{1,t-1}$ & $a_{2,t-1}$, with the kinematic position equation and the associated velocity equation (1).

$$\hat{y}_t = \hat{y}_{t-1} + T\hat{\dot{y}}_{t-1} + \frac{T^2}{2}((a_{1,t-1} - a_{2,t-1}) - \hat{b}_a) \quad (1)$$

$$\hat{\dot{y}}_t = \hat{\dot{y}}_{t-1} + T((a_{1,t-1} - a_{2,t-1}) - \hat{b}_a)$$

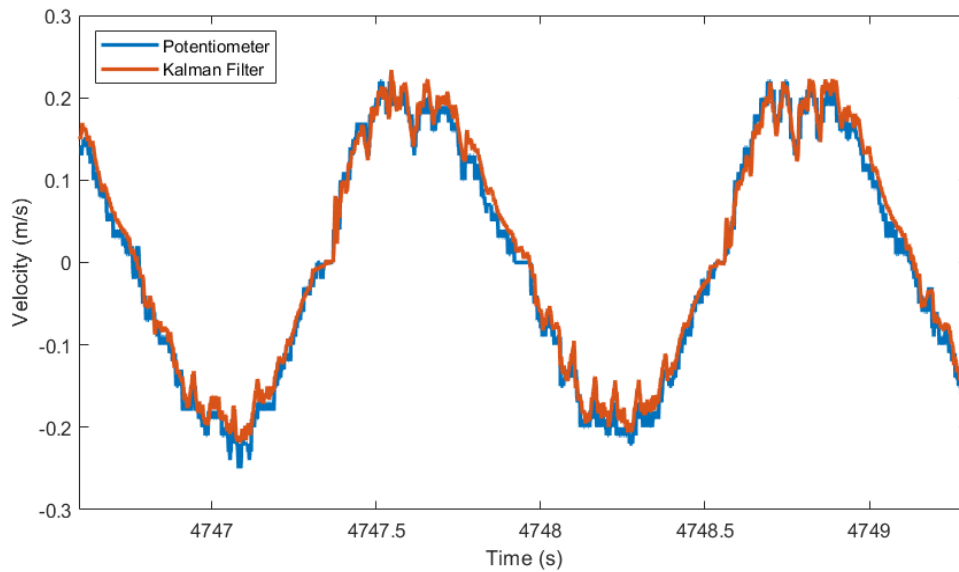


Figure 2: Estimated and measured, via time differencing the potentiometer measurement, velocity of the 1Hz test.

5.2 Thermal Modelling

A thermal model of the shock absorber was created to model the heat generated by, and conducted through, the shock absorber as the shock absorber was agitated. The model consists of two equations describing the relationship of heat transfer. This parametric model was also fitted to the shock absorber test measurements generated in its nominal state.

5.2.1 Thermal Transfer of Kinetic Generated Heat to Oil to Cylinder Body

The first equation of the thermal model (2) models the rate of change of the temperature of the damping fluid, T_{oil} , which is defined by the total heat transfer to the oil and moderated by the mass and thermal capacity of the oil, m_{oil} and c_{oil} . The heat transfer is determined by the kinetic energy converted to heat, $F_{damping}\dot{y}$, and the conduction to the shock cylinder based on the relative temperature, $T_{cyl} - T_{oil}$.

$$m_{oil}c_{oil} \frac{dT_{oil}}{dt} = h_{oil}A_{cyl}^{in}(T_{cyl} - T_{oil}) + F_{damping}\dot{y} \quad (2)$$

5.2.2 Thermal Transfer from Cylinder Body to Environment

The second equation of the thermal model (3) models the rate of change of the temperature of the metal cylinder of the shock absorber, T_{cyl} , which is defined by the total heat transfer to the metal and moderated by the mass and thermal capacity of the cylinder, m_{cyl} and c_{cyl} . The heat transfer is determined by the temperature difference of the cylinder relative to the damping fluid, $T_{cyl} - T_{oil}$, as well as convection and radiation to the surrounding environment, $T_{cyl} - T_{amb}$.

$$m_{cyl}c_{cyl} \frac{dT_{cyl}}{dt} = h_{oil}A_{cyl}^{in}(T_{oil} - T_{cyl}) + h_{amb}A_{cyl}^{out}(T_{amb} - T_{cyl}) + \varepsilon\sigma A_{cyl}^{out}(T_{amb}^4 - T_{cyl}^4) \quad (3)$$

5.2.3 Empirical Estimation

The model presented, (2)-(3), was fitted to the collected sensor data. The model was fitted by estimating the lumped parameters, i.e., $m_{oil}c_{oil} \Rightarrow A$, that minimised the RMSE of the generated cylinder temperature profile, T_{cyl} , in relation to the collected temperature data. The temperature data was the data collected using the cylinder thermocouple, for the 1, 2, and 4 Hz scenarios for the nominal operating condition of 12 Bar. The model inputs consisted of the sensor data collected from the loadcell, the ambient thermocouple, and the velocity as derived in Section 5.1. The output of the fitted model in relation the sensor data is presented in Figure 3.




Figure 3: Modelled and measured temperature profile of the 1Hz test at nominal gas pressure.

5.3 Force Modelling

A common structure of a shock absorber, shown in Figure 4, consists of a piston, a disk like shape partially blocking the flow of dampening fluid, attached to a connecting rod, in a cylinder filled with dampening fluid connected to a second chamber for holding a reserve oil pressurised by a gas chamber separated by a floating piston [11]. The piston, shown in Figure 4(a) in turn has holes, ports, that allow a limited rate of fluid through, some of which are in part covered by a shim stack, that deflects due the force exerted by fluid pressure created as the piston moves within the cylinder. To generate an estimate of the force expected to be generated at any point in time, a first principle parametric physical model of the shock absorber, Figure 4(b), was utilised to determine the force generated by the individual components within the shock absorber and their individual forces combined to generate an overall force [11]. The model presented here is in principle equivalent to the model resented in Reybrouck [12] with functional modifications to extend the applicability of the model. There are three types of force generating components in a shock absorber;

- **Port** – orifices that limit the flow of the dampening fluid, there are two versions, low speed (leak) port and high speed (main) port. There is a pair of each of these ports, one for the compression direction and the other for rebound.
- **Shim stack valve** – a set of annular springs that are stacked to control the flow of fluid through the main port. Similarly, there is a pair of shim stacks, one for the compression direction and the other for rebound.
- **Gas pressure** – the force generated in extension due to gas pressure in the reserve chamber and differential cross-sectional area either side of the piston due to the presence of a connecting rod on one side.

Each of these force components have coefficients that weight the individual terms and have different values during compression or rebound motion. A summary of these forces is presented in the following subsections.




Figure 4: The shock absorber physical representation.

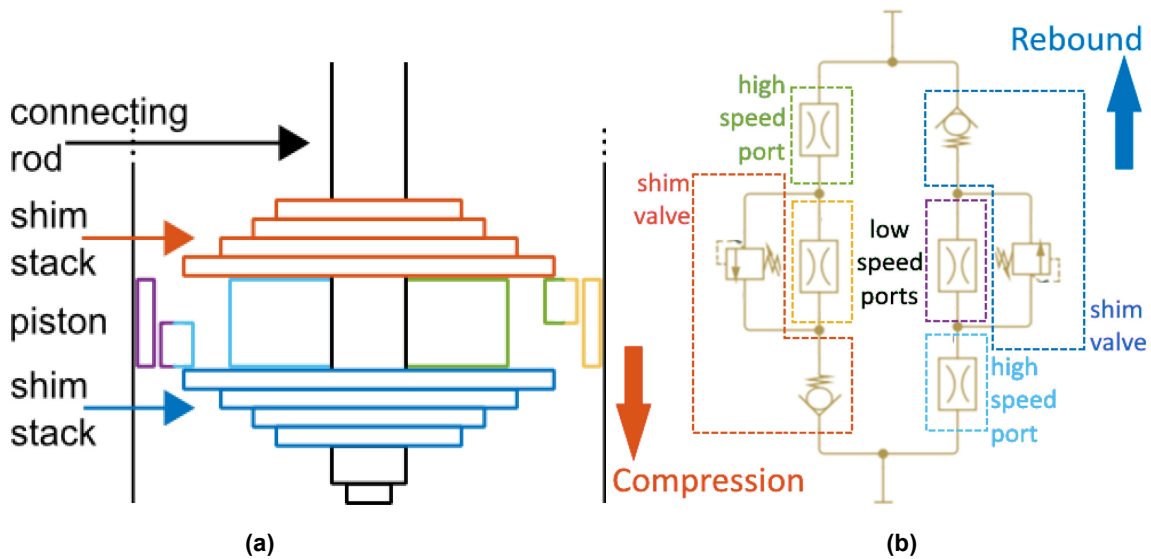


Figure 5: The dampening piston physical and schematic representation.

5.3.1 Leak Port Force

Leak port force (4) is the force generated by the leak port fluid restriction in the piston as the fluid attempts to bypass the pressure valve formed by the shim stack, depending on port design force can have differing relationship to velocity, in this case the force is linearly sensitive to the velocity of the fluid [13], equivalent to the velocity of the shock absorber, \dot{y} . Additionally, due to the relatively small velocities involved, this force is also sensitive to the acceleration of the fluid, \ddot{y} , due to hysteresis.

$$F_{leak} = K_{leak}v\dot{y} - K_{hys}\ddot{y} \quad (4)$$

5.3.2 Main Port Force

Main port force (5) is the force generated by the fluid restriction in the shock absorber piston that is regulated by the shim stack. Due to the fluid dynamics and the geometry of the port [13], the force is sensitive to the velocity of the fluid, \dot{y} , to the power of 1.75, and the fourth root of fluid viscosity.

$$F_{main} = K_{main}v^{0.25}\dot{y}^{1.75} \quad (5)$$

5.3.3 Shim Valve Force

Shim valve force (5) is the force due to the shim stack valve. This force manifests itself as flow valve regulated by a preloaded spring.

$$F_{shim} = F_{preload} + K_{spring}\dot{y} \quad (6)$$

5.3.4 Gas Pressure Force

The gas pressure force is the force used to pressurise the damping fluid while allowing the oil to occupy more space due to changes in the fluid density or displacement due to shock absorber extension. This force is composed of two components, the average gas pressure (7) and the instantaneous gas pressure (8) due to changes in extension, y , and adiabatic gas law and the adiabatic index of nitrogen of 1.4.

$$F_{gas,static} = \frac{mR_iA_{rod}T}{V_{gas,static}} \quad (7)$$

$$F_{gas,dyn} = \frac{F_{gas,static}V_{gas,static}^{1.4}}{(V_{gas,static} + A_{rod}y)^{1.4}} \quad (8)$$

5.3.5 Overall Force

The forces generated by the leak port, main port, shim valve, and gas pressure combine to generate the overall force generated by the shock absorber (9). The leak port and shim valve act in parallel as the shim valve bypasses the leak valve, these forces are combined using an empirical parallel combination of forces [6] and this combined force is added to the main port and gas pressure forces due to them acting in series.

$$F_{damping} = \frac{F_{leak}F_{shim}}{K_{tr}\sqrt{F_{leak}^{K_{tr}} + F_{shim}^{K_{tr}}}} + F_{main} - F_{gas,dyn} \quad (9)$$

5.3.6 Empirical Model Fitting

This parametric force model was then optimised, in combination with the already fitted thermal model, Section 5.2, to fit the sensor data recorded during the tests involving the shock absorber in its nominal state, i.e., 12 Bar. The relationship between the sensed and modelled force is shown in Figure 6 with the solid colour, and the colour bar, representing the modelled data and the dashed lines representing the average force measured at each velocity for the four temperature ranges defined by the centroid temperature shown in the legend.




Figure 6: Modelled and measured force of the 1 Hz test at nominal gas pressure.

5.4 Neural Network

The purpose of the neural network is to take various combinations of the sensor data, estimated velocity, and modelled thermal and force data, to generate an estimate of the average gas pressure based on immediate data, i.e. no memory, recursion or averaging, which could then be used to determine the operational condition of the shock absorber, as shown in Figure 7.

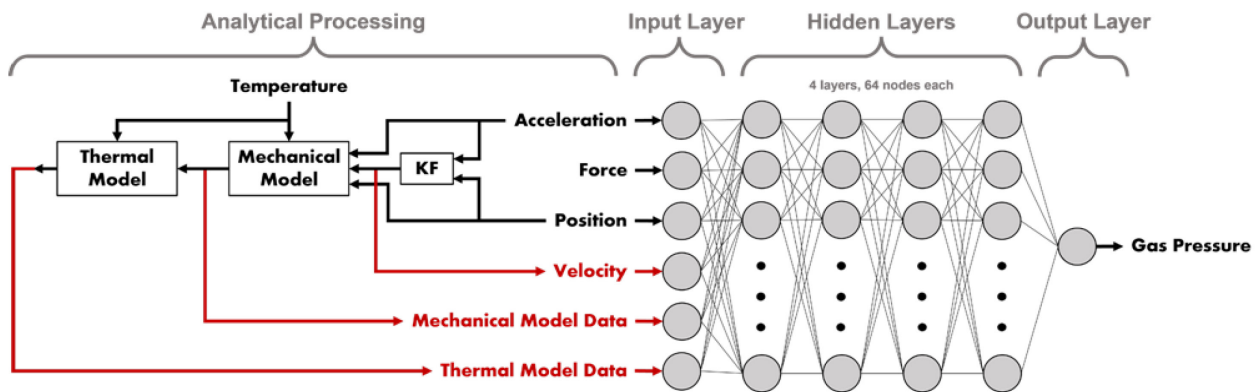


Figure 7: The digital twin architecture highlighting the network architecture and the connectivity of the modelling and estimation algorithms providing additional streams of information to the neural network.

5.4.1 Network Architecture

The architecture of the DNN was selected to be a feedforward neural network composed of 4 hidden layers with 64 nodes each. The method to determine the optimal architecture as well as more detailed analysis of the statistical behaviour of the network is outside of the scope of this publication. The overall architecture applied in this paper was chosen as it was the most straight forward approach in that any model generated data would be applied as additional inputs to the DNN. This architecture allowed the DNN to be treated conventional network simplifying the network development and training process. Additionally, simplifying

the explicit architecture allowed the analysis to focus on the effect of the derived data on the accuracy of the solution. The output layer consisted of one node corresponding to the gas pressure estimate while the input layer comprised N nodes corresponding to the number of inputs for the given permutation tested.

5.4.2 Network Training

The neural network was presented with data collected from all tested sets of actuation frequency and gas pressure with the predefined labels based on the gas pressure measured at the beginning of the test scenario. The data was filtered and down sampled such that all the data arrived at the same rate of 50Hz. This dataset was randomly partitioned into training and testing sets, sized 75% and 25%, respectively. Each permutation of the neural network was then trained on the subset of data streams as defined by the sensor, modelling, estimation combination.

6.0 RESULTS

The training was completed as specified in Section 5.4.2 and the network RMSE then evaluated using the testing set for each of the 13 permutations, the results are presented in Table 2. It can immediately be seen that in general the more data inputs that are provided, be it sensor or derived, the more accurate the estimate is, but a more detailed look highlights which data streams have the most significant impact on the accuracy of the estimate.

When a loadcell was combined with positional measurements (as in models 3, 7, 9, 11 and, 13) this combination generated the most meaningful impact on accuracy. This was most notably seen in model 3 where it is more accurate than any other combination using only two of the available sensor inputs, and with only a marginally higher error than using all three kinetic sensors. The only difference in available data between models 7 and 9 was that in model 9, the estimated velocity was available from the Kalman Filter, while in model 7 the information necessary to infer velocity was available but had to be calculated by the network. Providing the velocity directly as an input to the network resulted in a 19% improvement in RMSE. Additionally, providing the estimated force from the mechanical model resulted in a 24% improvement from the baseline, a further 7% improvement over model 9. Finally, with the addition of the estimates from the thermal model, an improvement of 63% over baseline, or a 50% improvement without the thermal model was achieved.

The most significant finding from this experiment was that model 12, despite not utilising the loadcell, generated an instantaneous pressure estimate with an RSME of 1.40 bar at 50 Hz. When a moving average filter (of 1s) was applied to the estimate of pressure, the RMSE decreased to 0.67 bar. This showed that while the loadcell was the single most informative sensor for the network it could be removed without critically affecting the resulting system estimations.

The advantage of avoiding a loadcell in the sensor suite is shown in Table 3. The total cost of the sensors used in model 12 was only 7% of the cost of the alternatives that utilised the loadcell. Although model 12 was less accurate (65% higher RMSE) than model 13, it was sufficient to determine the gas pressure in the reserve chamber at a level that easily discriminated the pressure instances used, with the cost savings more than making up for the loss of precision.

Table 2: Tested permutations of available data to determine the accuracy of the neural network utilising combinations of sensor inputs.

Model ID	Model Name	Force	Position	Acceleration	Velocity	Mechanical Model Data	Thermal Model Data	RMSE (bar)
1	-							3.34
2	-							3.53
3	Force and Position DNN							2.51
4	-							3.52
5	-							3.18
6	-							3.18
7	Force, Position and Acceleration DNN							2.27
8	-							3.18
9	Kalman-Velocity Hybrid DNN							1.84
10	-							3.08
11	Mechanical Hybrid DNN							1.72
12	Thermomechanical Hybrid DNN (No Force)							1.40
13	Thermomechanical Hybrid DNN							0.85

Table 3: Shortlist of input permutations of the available sensors and derived data showing the significant cost of a loadcell and the ability of model derived estimates to compensate for the lack of a load cell.

Model ID	Force	Position	Acceleration	Velocity	Mechanical Model Data	Thermal Model Data	RMSE (bar)	No. of Sensors	Sensor Cost (\$ AUD)
3							2.51	2	\$4,270
7							2.27	3	\$4,281
9							1.84	3	\$4,281
11							1.72	4	\$4,292
12							1.40	3	\$297
13							0.85	4	\$4,292

This study has shown not only the effectiveness of a neural network based digital twin system but that the inclusion of selective pre-processing algorithms can substantially improve accuracy. Furthermore, through utilisation of system modelling, hardware costs can be reduced by avoiding otherwise essential sensors. These results also validate the general approach taken, that by factorising a system into smaller components, or features, a hybrid method can be used where the most appropriate algorithm is selected for each of the individual components. Thus, classical models should be used where solutions can be easily formulated and implemented while neural networks should be used for less well-defined tasks. This results in decreased development and training time for the network as well as a more accurate and traceable digital twin.

7.0 CONCLUSION

The most significant finding from this experiment was that it is possible to indirectly estimate parameter measurements from other sensor measurements without incurring the additional expense and complexity of fully measuring all parameters. As shown, despite not utilising a loadcell, model 12 generated an instantaneous pressure estimate with an RSME of 1.40 bar at 50 Hz. When a moving average filter (of 1s) was applied to the estimate of pressure, the RMSE decreased to 0.67 bar. This demonstrated that while the loadcell was the single most informative sensors for the network it could be removed without critically affecting the resulting system estimations. Finally, the case study presented here also serves as an analogue to larger and more complicated systems or platforms, as the factorisation approach presented here can be implemented using a variety of metrics including modularity, functional complexity, development time, and system risk assessment.

Our research suggests that when twinning a particular item or component, the most effective approach, in terms of accuracy, efficiency and interrogability, is to first factorise the object into elements that are as small as possible and optimising the approach for each element individually. These elements should then be recombined; with redundant, low contributing or high-cost elements selectively removed to produce a more efficient, but still functional digital twin.

8.0 ACKNOWLEDGEMENTS

The authors acknowledge the funding support provided by the Australian Defence Science and Technology Group for this research project. We also acknowledge the contributions by Supashock in providing the testing facilities necessary for conducting the research presented in this paper as well as the manufacturing resources for constructing a concept demonstrator to validate this work in the field.

9.0 REFERENCES

- [1] F. Tao, H. Zhang, A. Liu, and A.Y. Nee, "Digital twin in industry: State-of-the-art," *IEEE Transactions on industrial informatics*, vol. 15, no. 4, pp. 2405-2415, 2018.
- [2] S. Centomo, N. Dall'Ora, and F. Fummi, "The Design of a Digital-Twin for Predictive Maintenance," in *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2020, vol. 1: IEEE, pp. 1781-1788.
- [3] A. Kloss, S. Schaal, and J. Bohg, "Combining learned and analytical models for predicting action effects from sensory data," *The International Journal of Robotics Research*, vol. 41, no. 8, pp. 778-797, 2022.
- [4] M. Cheng, S. Nazarian, and P. Bogdan, "There is hope after all: Quantifying opinion and trustworthiness in neural networks," *Frontiers in Artificial Intelligence*, vol. 3, p. 54, 2020.
- [5] W. Samek, G. Montavon, A. Vedaldi, L.K. Hansen, and K.-R. Müller, *Explainable AI: interpreting, explaining and visualizing deep learning*. Springer Nature, 2019.
- [6] V. Pracny, M. Meywerk, and A. Lion, "Full vehicle simulation using thermomechanically coupled hybrid neural network shock absorber model," *Vehicle System Dynamics*, vol. 46, no. 3, pp. 229-238, 2008.

- [7] V. Barethiye, G. Pohit, and A. Mitra, “A combined nonlinear and hysteresis model of shock absorber for quarter car simulation on the basis of experimental data,” (in English), *Engineering Science and Technology-an International Journal-Jestech*, vol. 20, no. 6, pp. 1610-1622, Dec 2017, doi: 10.1016/j.jestch.2017.12.003.
- [8] J.L. Xu, J.L. Chu, and H.W. Ma, “Hybrid modeling and verification of disk-stacked shock absorber valve,” (in English), *Advances in Mechanical Engineering*, vol. 10, no. 2, p. 1687814018756398, Feb 6 2018, doi: Artn 1687814018756398
- [9] R.E. Kalman, “A new approach to linear filtering and prediction problems,” 1960.
- [10] C.Q. Howard, N. Sergiienko, and G. Gallasch, “Monitoring the age of vehicle shock absorbers,” in *International Conference on Science and Innovation for Land Power*, 2018, vol. 2018.
- [11] J.C. Dixon, *The shock absorber handbook*. John Wiley & Sons, 2008.
- [12] K. Reybrouck, “A Non Linear Parametric Model of an Automotive Shock Absorber,” *SAE Transactions*, vol. 103, pp. 1170-1177, 1994. [Online]. Available: <http://www.jstor.org/stable/44611831>. D. C. Rennels and H. M. Hudson, *Pipe flow: A practical and comprehensive guide*. John Wiley & Sons, 2012.
- [13] D.C. Rennels and H.M. Hudson, *Pipe flow: A practical and comprehensive guide*. John Wiley & Sons, 2012.