MONASH University

GROUP OF EIGHT

Information Technology

Bayesian networks for decision making under uncertainty

How to combine data, evidence, opinion and guesstimates to make decisions

Professor Ann Nicholson

Deputy Dean

Faculty of Information Technology

Monash University

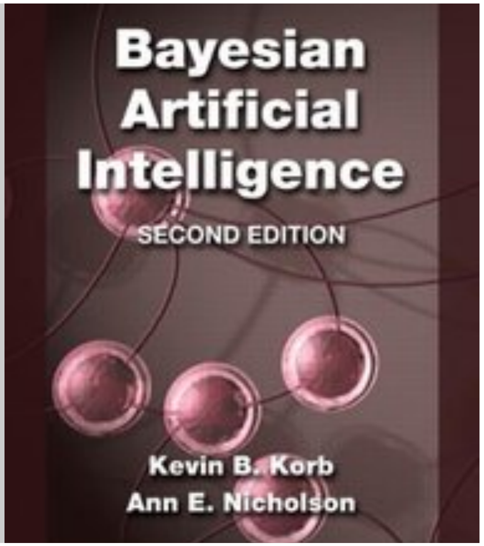

Sources of Uncertainty

Ignorance

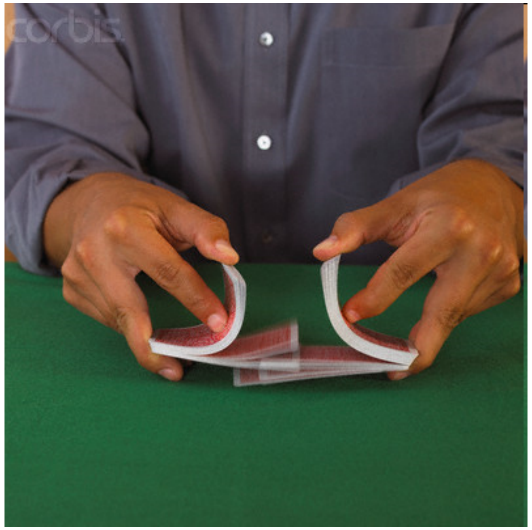

Physical randomness

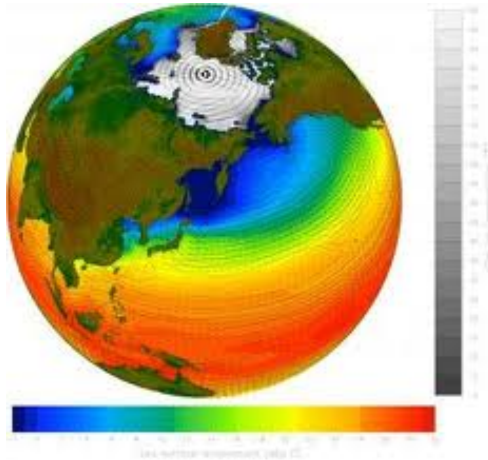

Complexity

- Solution: Probability Theory

The reasoning process

1. Start with a belief in a proposition

“The imported mango is infested with a pest”

“The patient has lung cancer”

“The applicant will pay back the loan”

“Sales of a product will increase”

“Beliefs”

- Very unlikely
- 1% chance
- Odds of 100 to 1
- **0.01 probability**

From

- Gut feeling
- Expert opinion
- Data

The reasoning process

1. Start with a belief in a proposition

"The imported mango is infested with a pest"

"The patient has lung cancer"

"The applicant will pay ba the loan"

"Sales of a product will increase"

2. New information becomes available

"It is from a country where the pest is endemic"

"The patient is a smoker and has a cough"

"The applicant defaulted on a previous loan"

"Some models are recalled for safety reasons"

3. Update your beliefs

But how?

Probability theory for representing uncertainty

- Propositions are either true or false

  “Patient has cancer”
- Assigns a numerical degree of belief between 0 and 1 to propositions

  P(“Patient has cancer”) = 0.001

  prior probability (unconditional)
- We can now represent the impact of evidence on belief

  P(“Patient has cancer | positive mammogram”) = 0. 8
  - *Conditional* probability
  - Or, *posterior* probability (by way of Bayes’ theorem)

The Bayesian approach

The Rev. Thomas Bayes
1702?-1761

- Represent uncertainty by probabilities
- Use **Bayes' theorem:**

  h = hypothesis

  e = evidence

Starting belief='prior'

$$P(h|e) \leftarrow \frac{P(e|h) \times P(h)}{P(e)}$$

New belief

Estimating Risk

Scenario: An athlete is tested for steroid use and the test comes back positive.

You know that:

- One in 100 competitors are thought to take steroids
- The test isn't always accurate

  False positive rate: 10%

  False negative rate: 20%

Q. What is the probability the athlete is a drug cheat?

Bayes' Theorem for Estimating Risk

Suppose:

- $\boldsymbol{h}$ = "the athlete is taking steroids"
- $\boldsymbol{e}$ = "test result is positive"

And:

- $P(h) = 0.01$ (one in 100 people)
- $P(e|h) = 0.8$ (true positive rate)
- $P(e|\text{not } h) = 0.1$ (false positive rate)

h = hypothesis

e = evidence

$$\underbrace{P(h|e)}_{\text{New belief}} \leftarrow \frac{\underbrace{P(e|h)}_{\text{Likelihood}} \times \overbrace{P(h)}^{\text{Old belief}}}{P(e)}$$

What is $P(h|e)$? ≈ 0.075 (7.5%)

In general, people can't do Bayes Theorem (well) off hand!

And how do we scale up to $X_1, X_2, \ldots.. X_{100}, \ldots. \quad X_{1000}$??

Bayesian Networks

- Developed by graphical modeling & AI communities in 1980s for probabilistic reasoning under uncertainty
- Many synonyms
  - Bayes nets, Bayesian belief networks, directed acyclic graphs, probabilistic networks

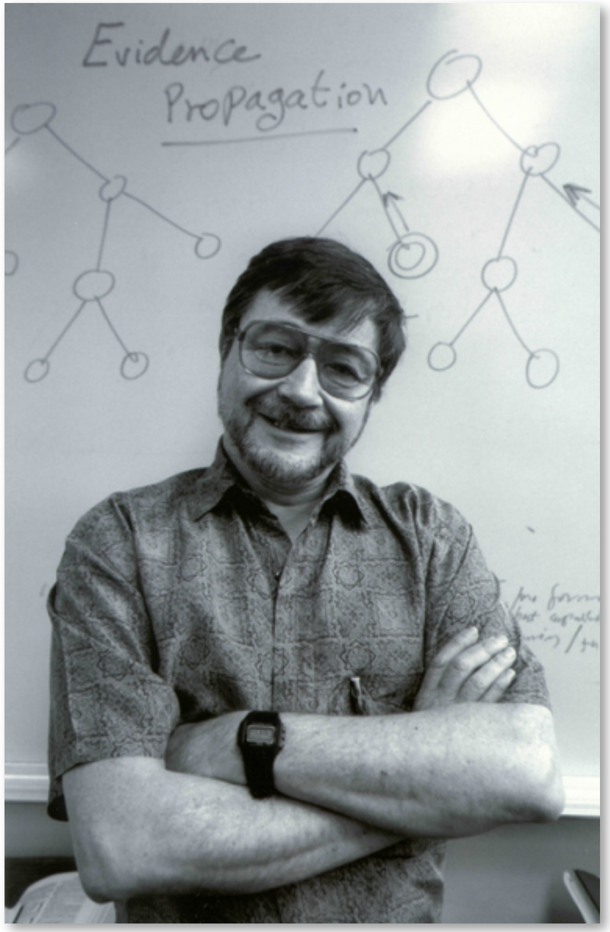

Judea Pearl
2012 Turing
Award

Probabilistic Graphical Models

For a target system…

… capture the process structure (not black box)

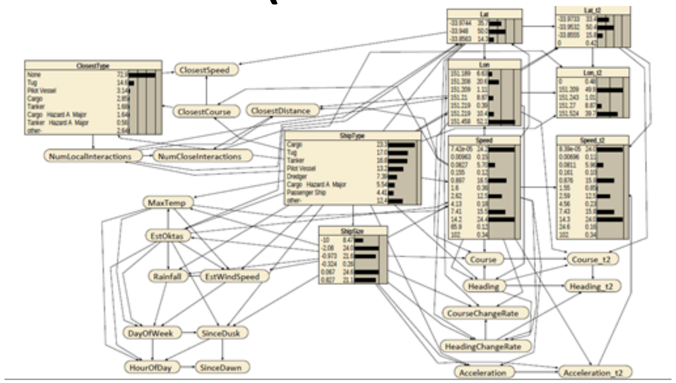

…and quantify the uncertainty

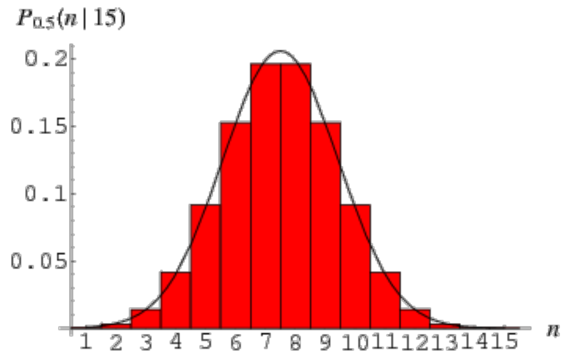

Why use models?

- Increases understanding
- Supports decision making
- Use new data and evaluation to improve over time

Native Fish Example

A local river with tree-lined **banks** is known to contain **native fish populations**, which need to be conserved. Parts of the **river** pass through **croplands**, and parts are susceptible to **drought conditions**. **Pesticides** are known to be used on the crops. **Rainfall** helps native fish populations by maintaining **water flow**, which increases **habitat suitability** as well as **connectivity** between different habitat areas. However rain can also wash pesticides that are dangerous to fish from the croplands into the river. There is concern that the **trees** and native fish will be affected by drought conditions and crop pesticides.

See http://bayesianintelligence.com/publications/TR2010_3_NativeFish.pdf

Bayesian Networks - Definition

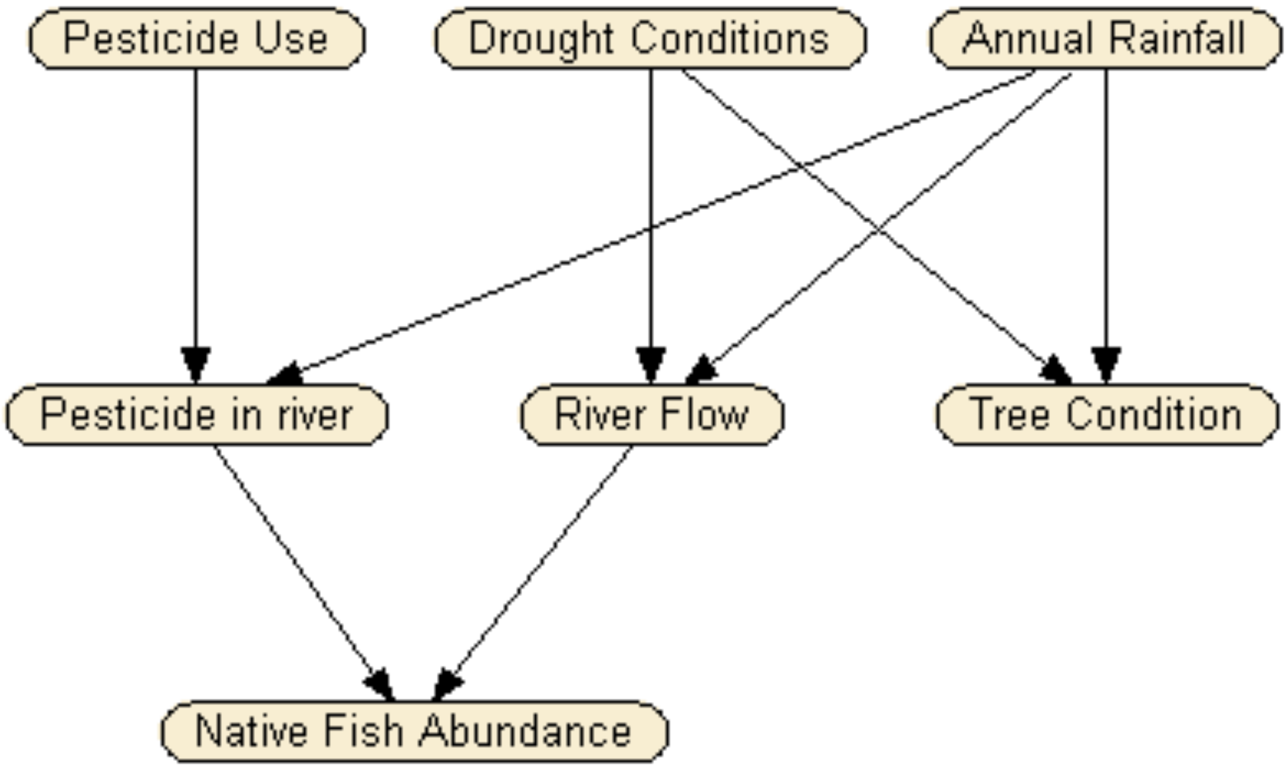

A **graph** in which the following holds:

1. A set of random variables = **nodes** in network
2. A set of directed **arcs** connects pairs of **nodes**

Structure represents the causal process.
(Anything missing?)

Bayesian Networks - Definition

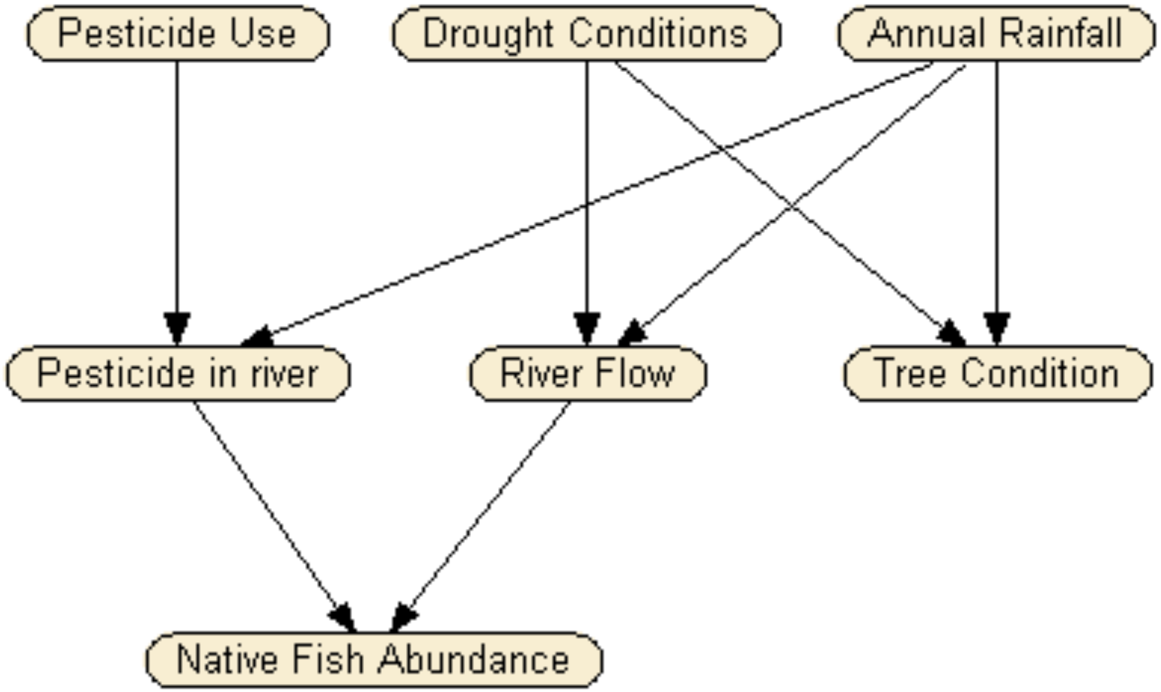

Pesticide in River	River Flow	P(FishAbundance | PesticideInRiver, RiverFlow) High	Medium	Low
High	Good	0.2	0.4	0.4
High	Poor	0.01	0.1	0.89
Low	Good	0.8	0.15	0.05
Low	Poor	0.05	0.15	0.8

WORST CASE

BEST CASE

Each row sums to 1

A **graph** in which the following holds:

1. A set of random variables = **nodes** in network
2. A set of directed **arcs** connects pairs of **nodes**
3. Each **node** has a **conditional probability table (CPT)** that quantifies the effects the ***parent*** nodes have on the ***child*** node
4. It is a directed acyclic graph (DAG), i.e. no directed cycles

Before you know anything
(no evidence)

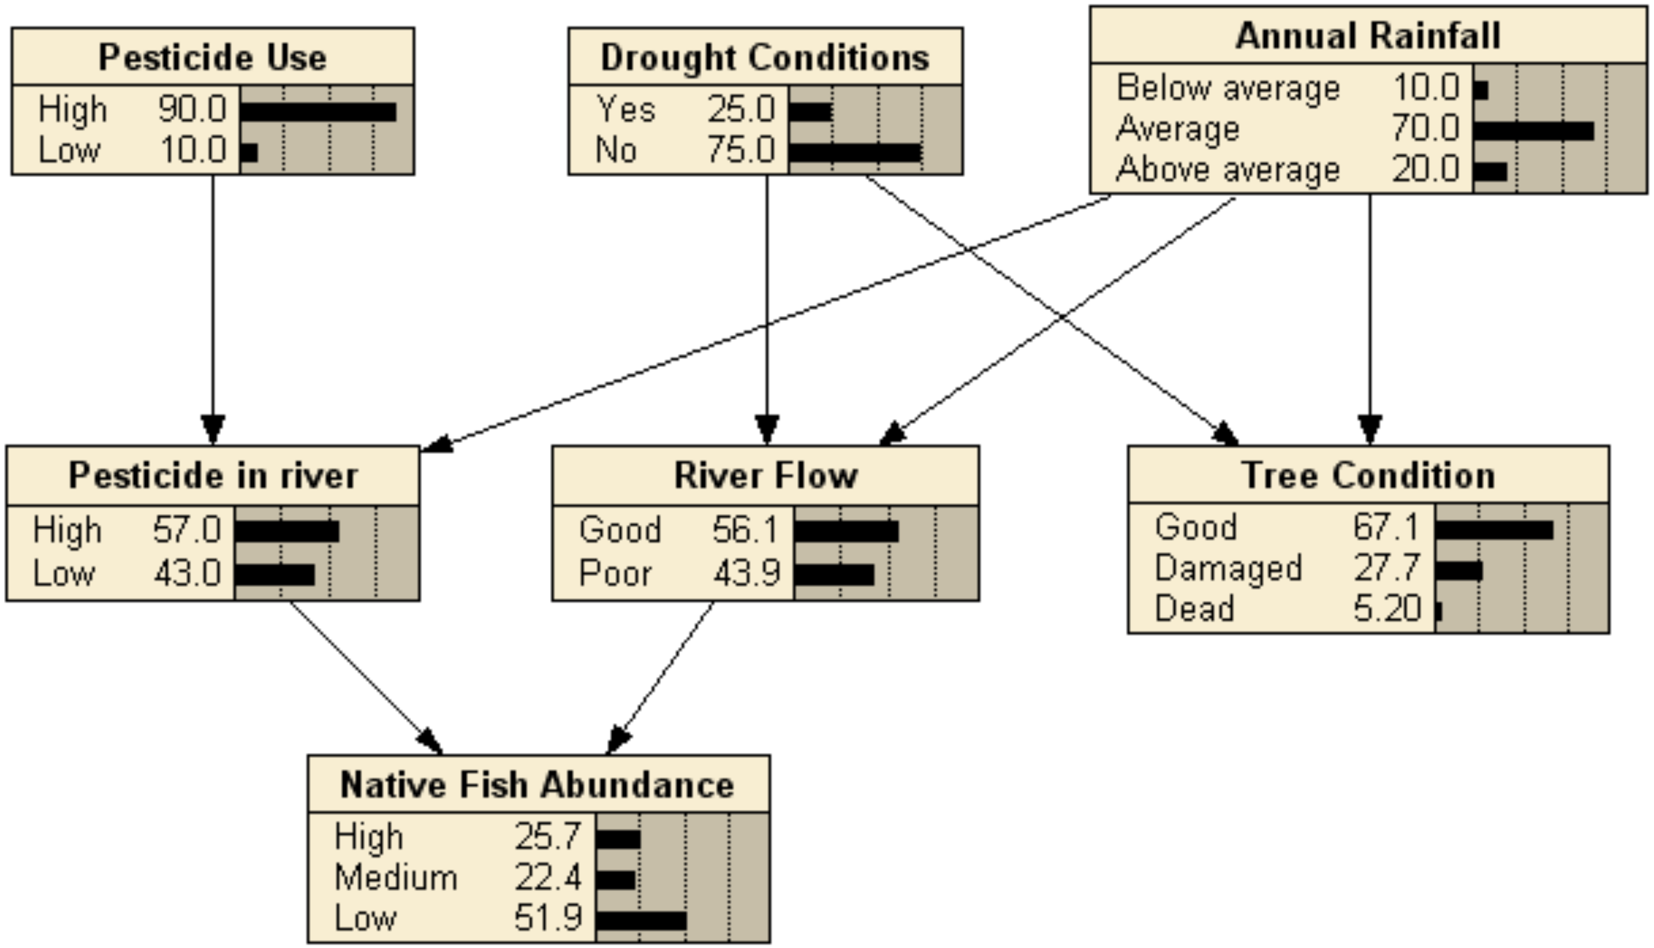

(Screen shots from the Netica BN software)

Diagnosis

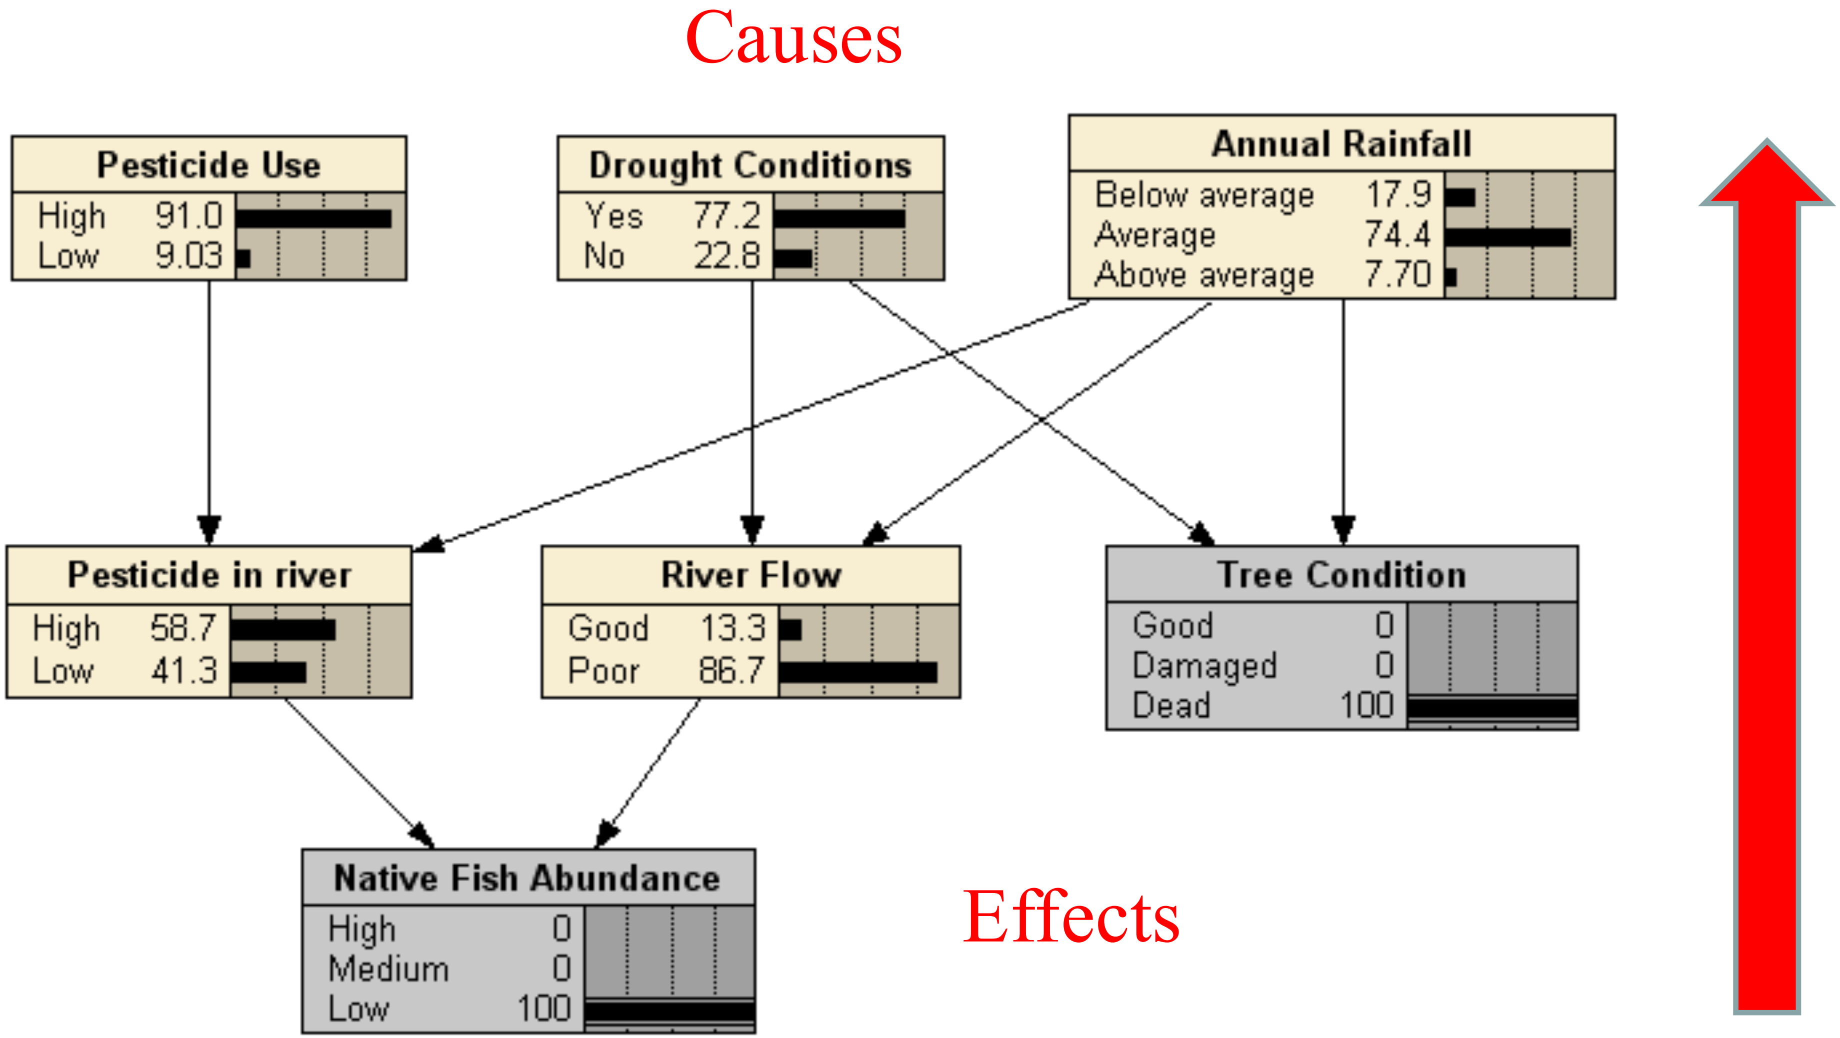

Causes
Pesticide Use
High 91.0
Low 9.03
Drought Conditions
Yes 77.2
No 22.8
Annual Rainfall
Below average 17.9
Average 74.4
Above average 7.70
Pesticide in river
High 58.7
Low 41.3
River Flow
Good 13.3
Poor 86.7
Tree Condition
Good 0
Damaged 0
Dead 100
Native Fish Abundance
High 0
Medium 0
Low 100
Effects

Prediction

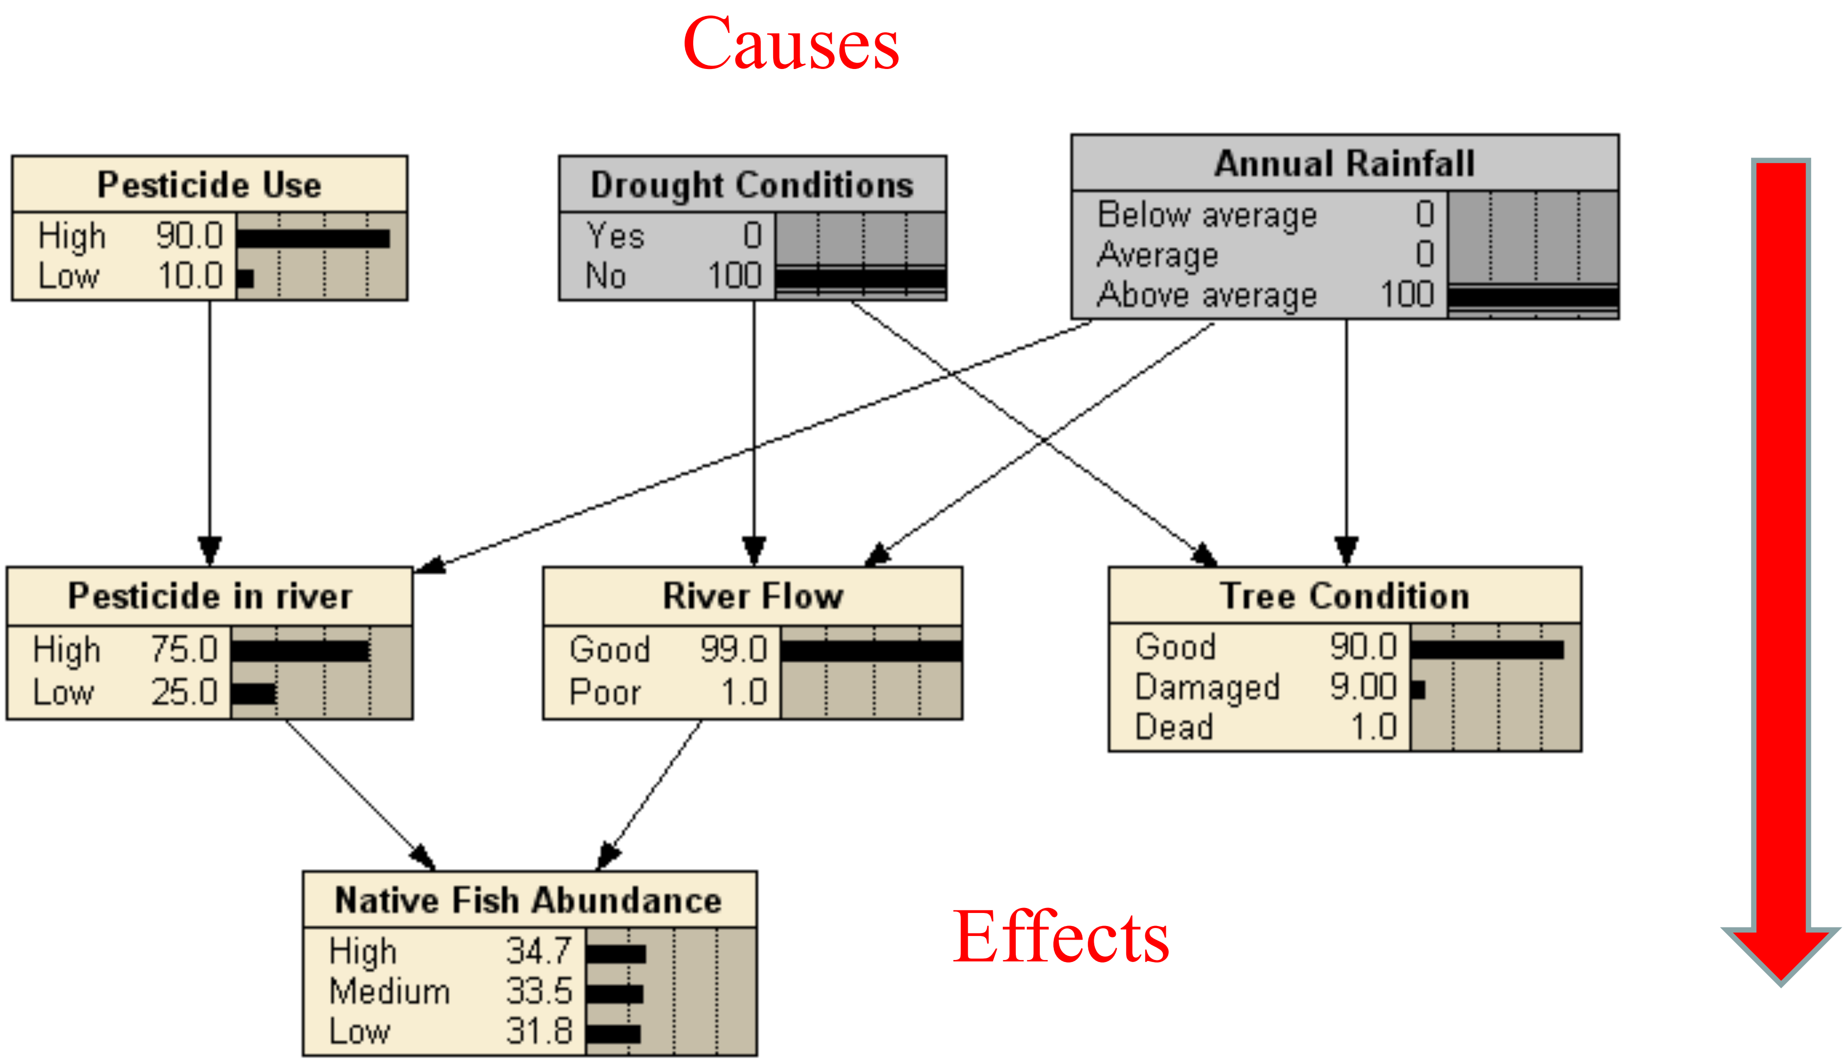
Causes
Pesticide Use
High 90.0
Low 10.0
Drought Conditions
Yes 0
No 100
Annual Rainfall
Below average 0
Average 0
Above average 100
Pesticide in river
High 75.0
Low 25.0
River Flow
Good 99.0
Poor 1.0
Tree Condition
Good 90.0
Damaged 9.00
Dead 1.0
Native Fish Abundance
High 34.7
Medium 33.5
Low 31.8
Effects

What next?

- Have model
- Have estimates of posterior probabilities

Q. How do we use these probabilities to inform decisions (about action or interventions)?

Risk Assessment – the decision-theoretic view

Risk = Likelihood x Consequence

P(Outcome|Action,Evidence) Utility(Outcome|Action)

Definition (Expected Utility)

$$EU(A|E) = \sum_i P(O_i|E,A) \times U(O_i|A)$$

Decision making is about reducing risk
or "maximising expected utility"

Decision network = BN + decisions + utilities

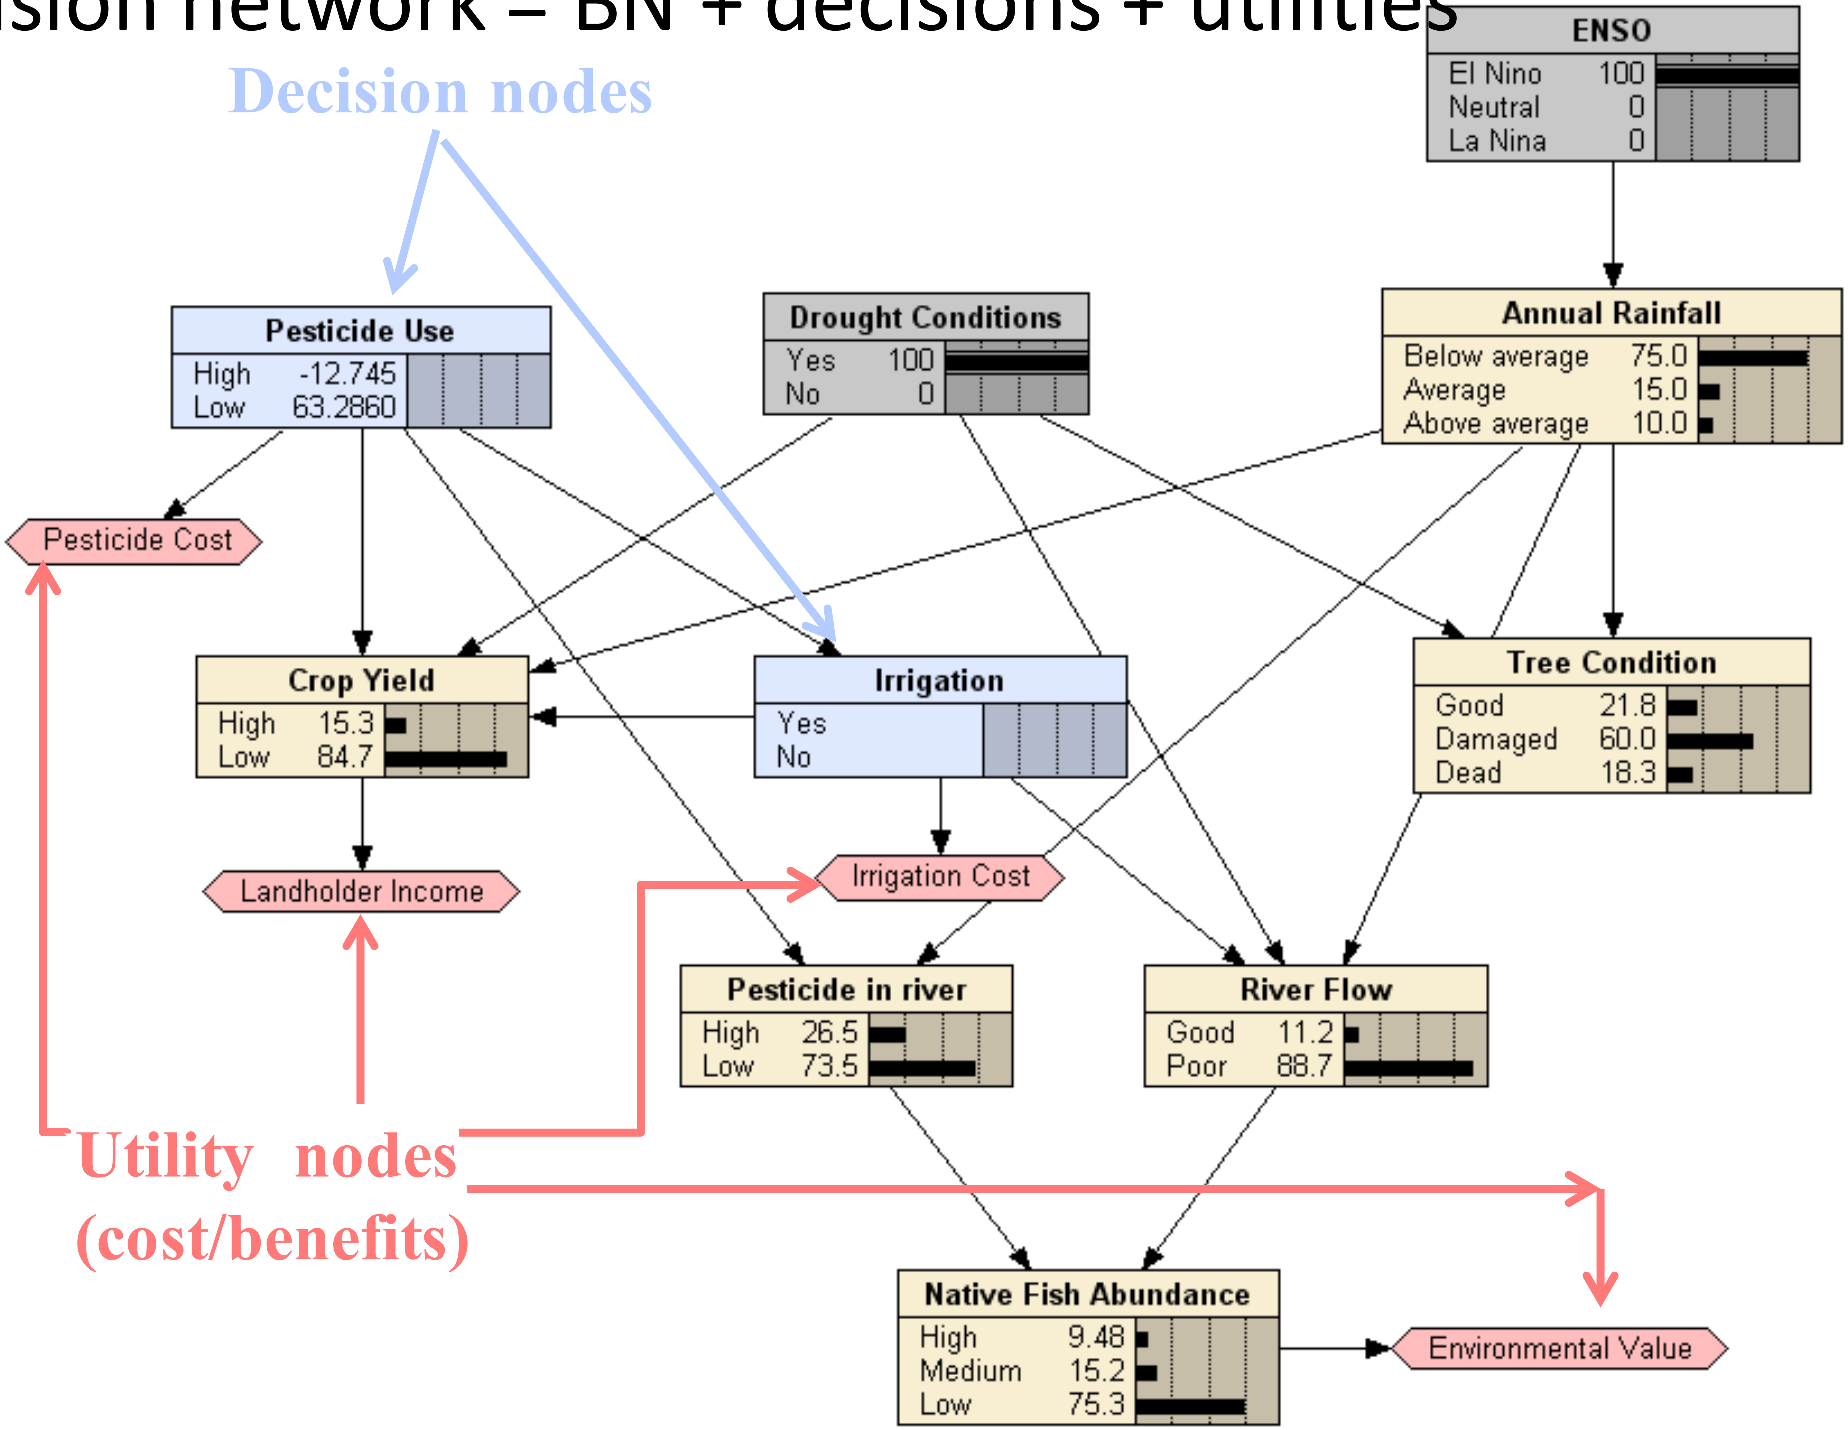

Our methodology

1: Build a model

- E.g. Variables: patient's details, diseases, symptoms, interventions
- Costs/benefits: eg. $, QALY

2: Embed model in decision support tool

- Diagnosis
- Prognosis
- Treatment
- Risk assessment
- Prevention

Experts

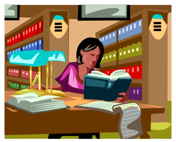

Literature

Data

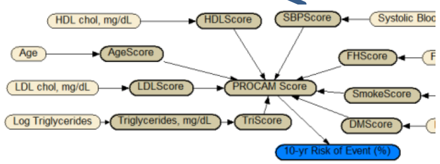

Structure

Pesticide in River	River Flow	P(FishAbundance | PesticideInRiver, RiverFlow) High	Medium	Low
High	Good	0.2	0.4	0.4
High	Poor	0.01	0.1	0.89
Low	Good	0.8	0.15	0.05
Low	Poor	0.05	0.15	0.8

Parameters

Review

Test

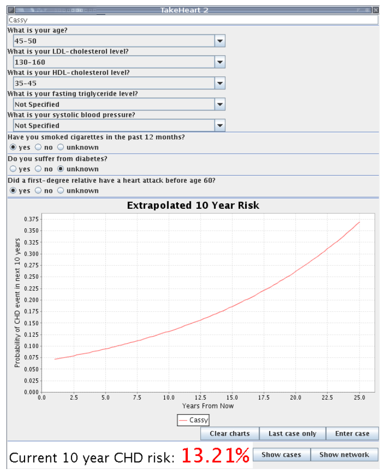

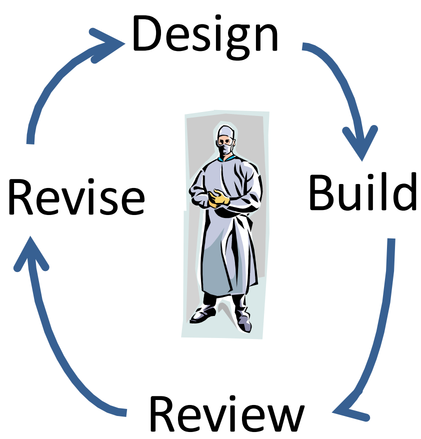

Bayesian networks for fog forecasting
(Collaboration with Aus Bur. of Meteorology)

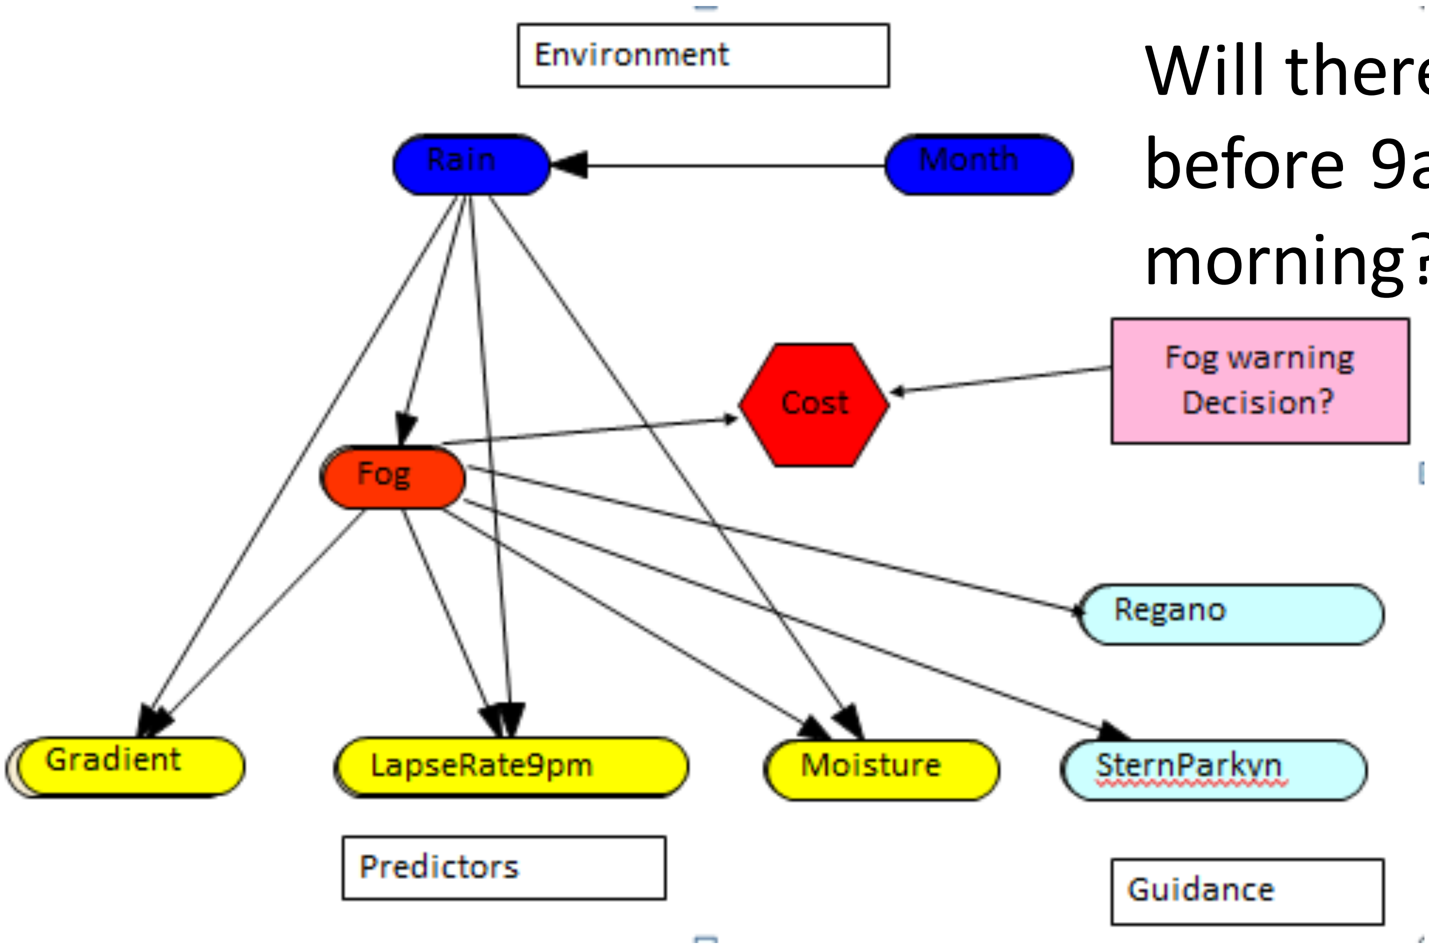

Will there be a fog before 9am tomorrow morning?

Phase 1: (Boneh et al, 2015)

In use by weather forecasters in Melbourne since 2006

Bayesian networks for fog forecasting

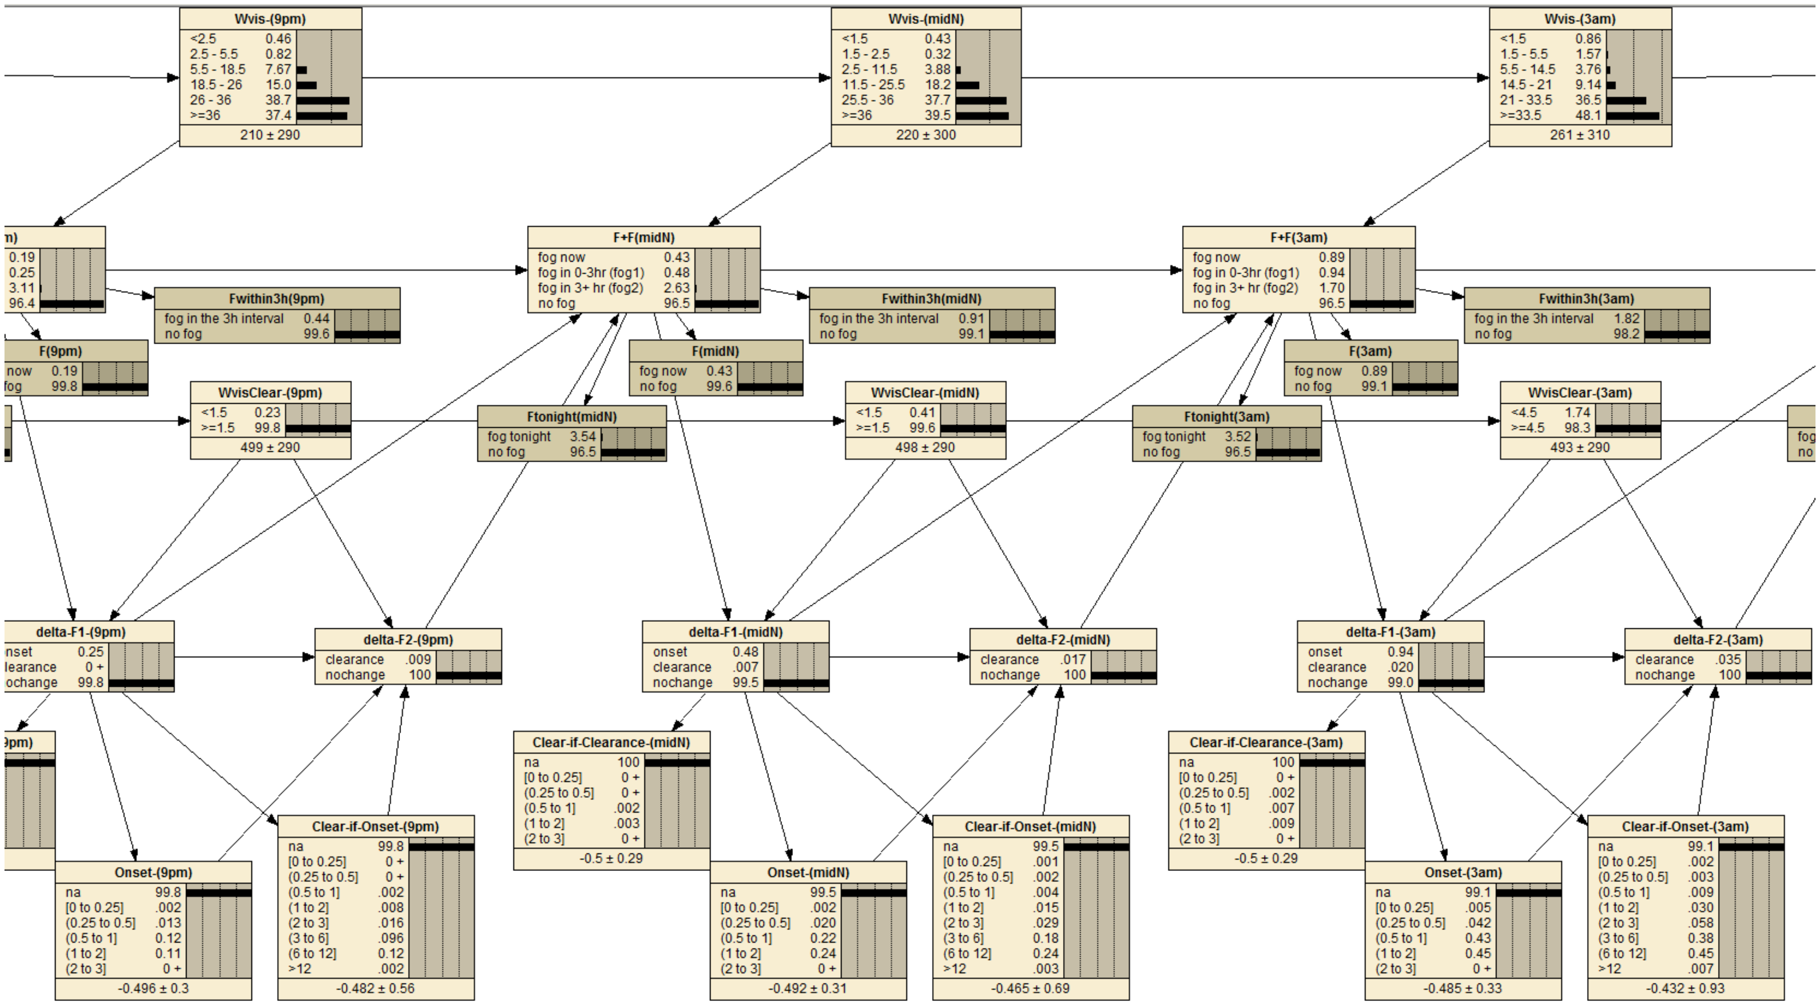

Stage 2: 2013-15 Research project (prototype) (Boneh et al. In preparation)

- Explicit temporal modelling, predicting time of onset and clearance

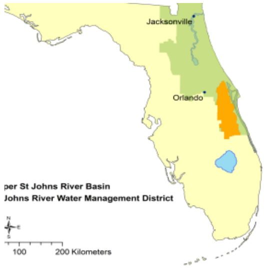

Case Study: Modelling Willows (in St. John’s River Basin, Florida)

Experiment -2011

Plant and seed collecti

Transplanting

Greenhouse expts

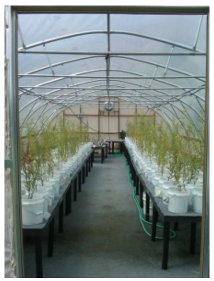

C

tificial Islands

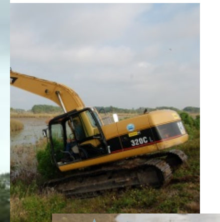

Ponds

Grazing

DBN Example: Modelling Willows (in St. John's River Basin, Florida)

(Wilkinson et al., 2013)

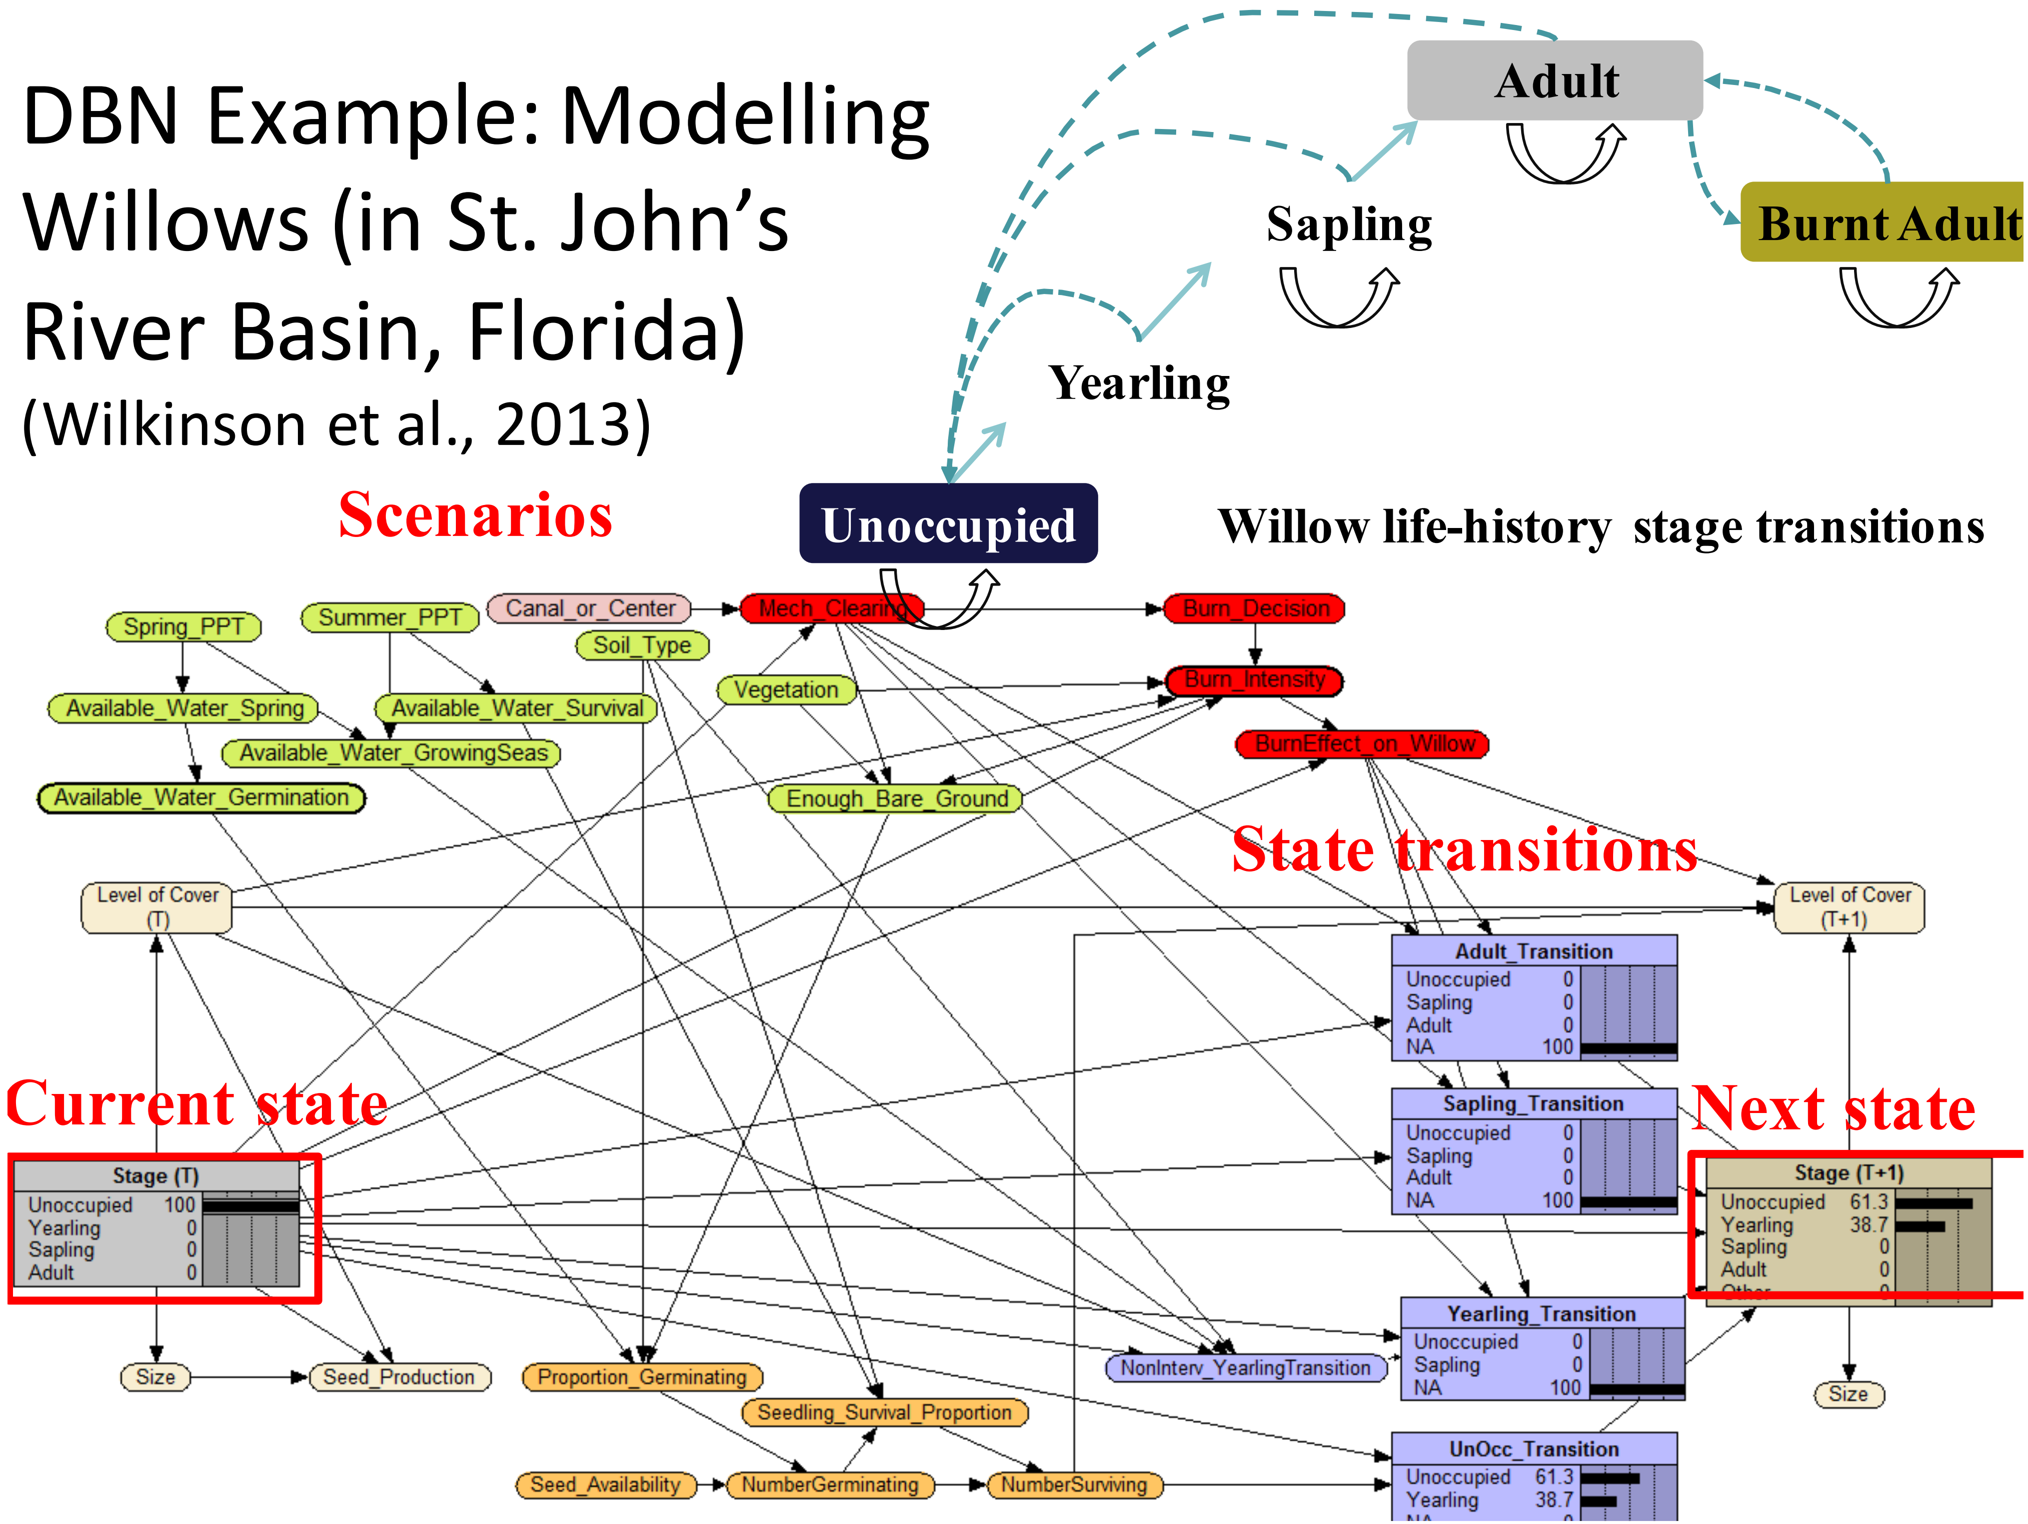

Architecture of the Integrated Management Tool

(Chee et al., 2016)

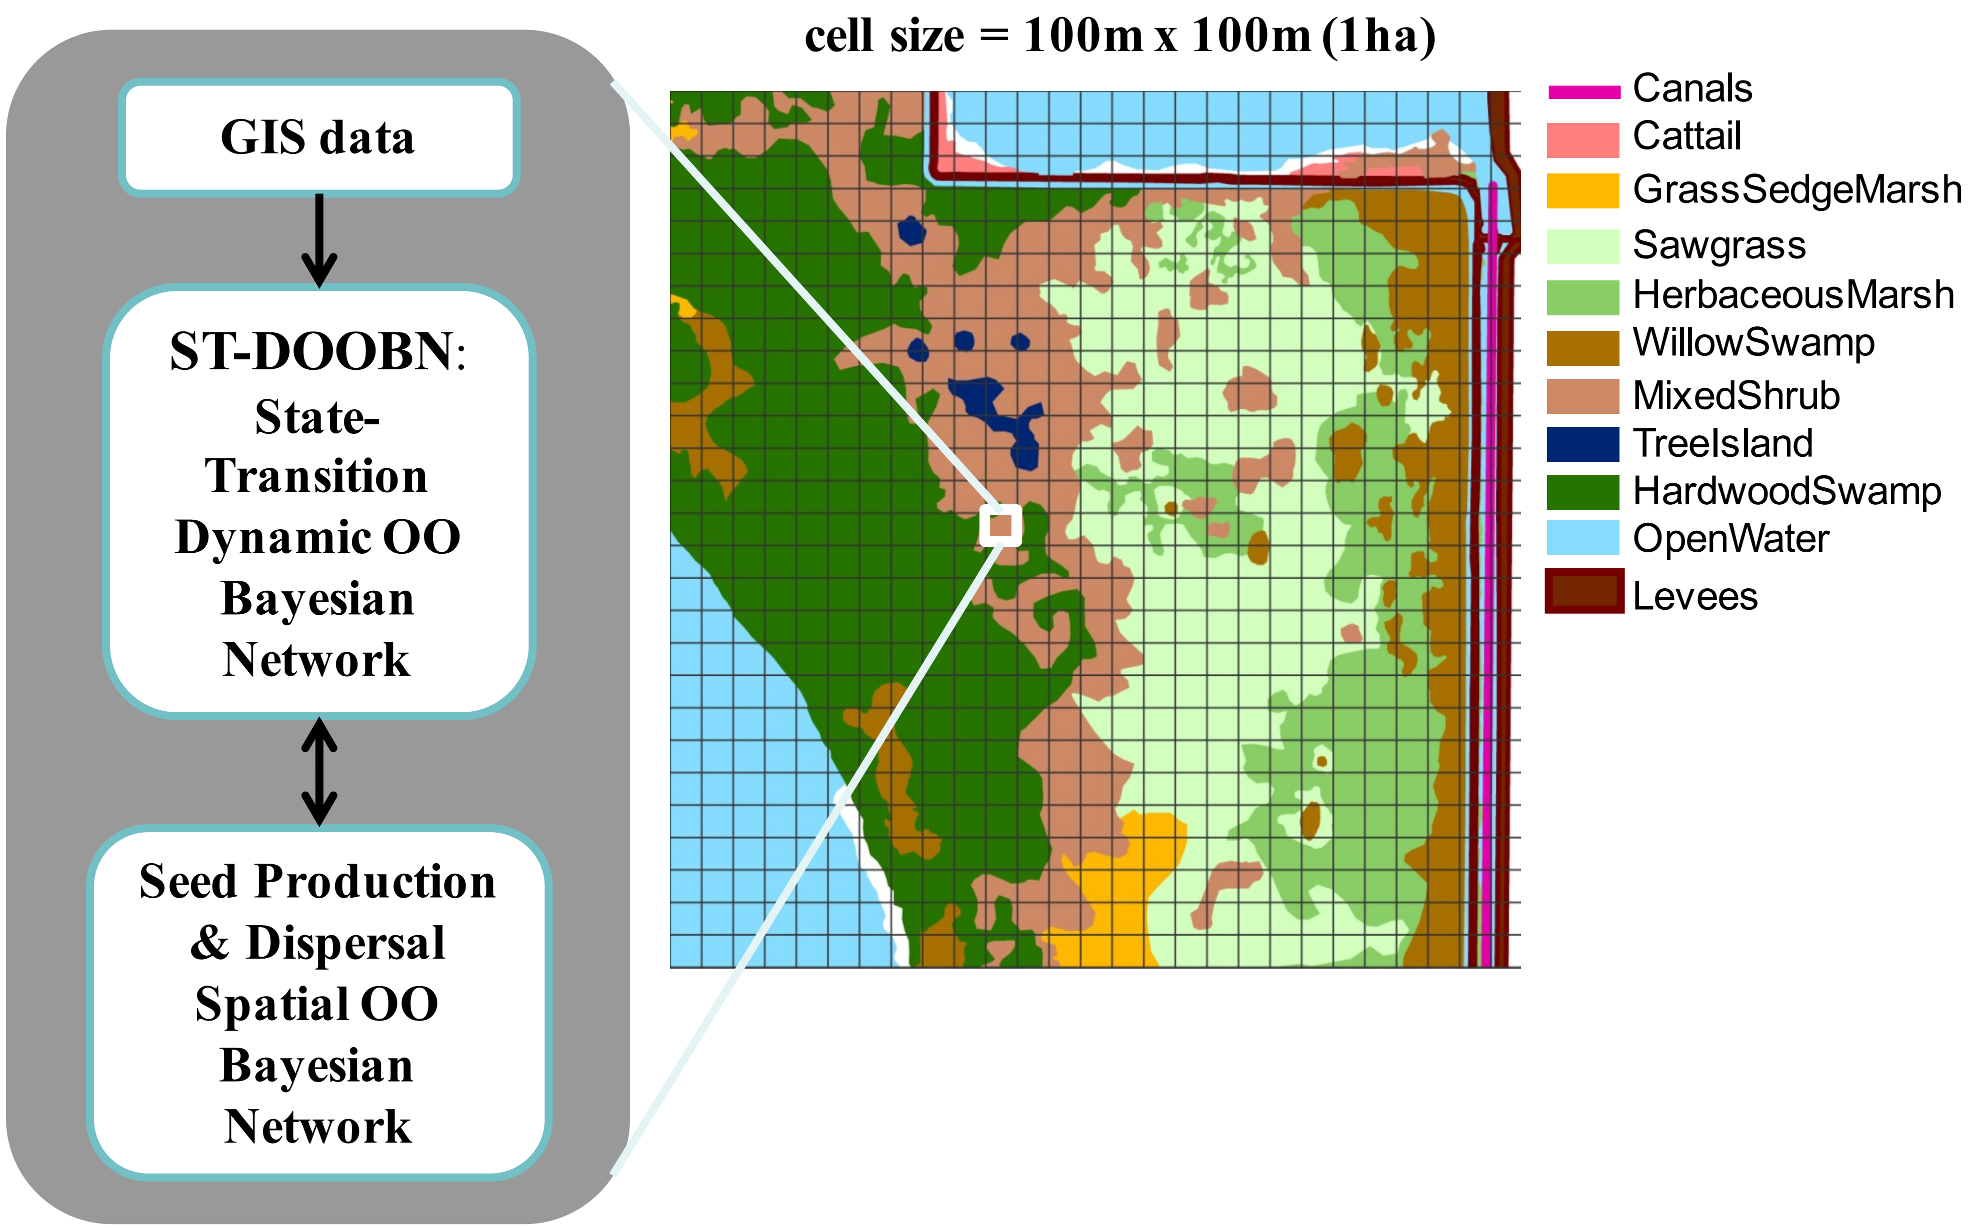

Willows ST-OODBN$_{x,y,t}$

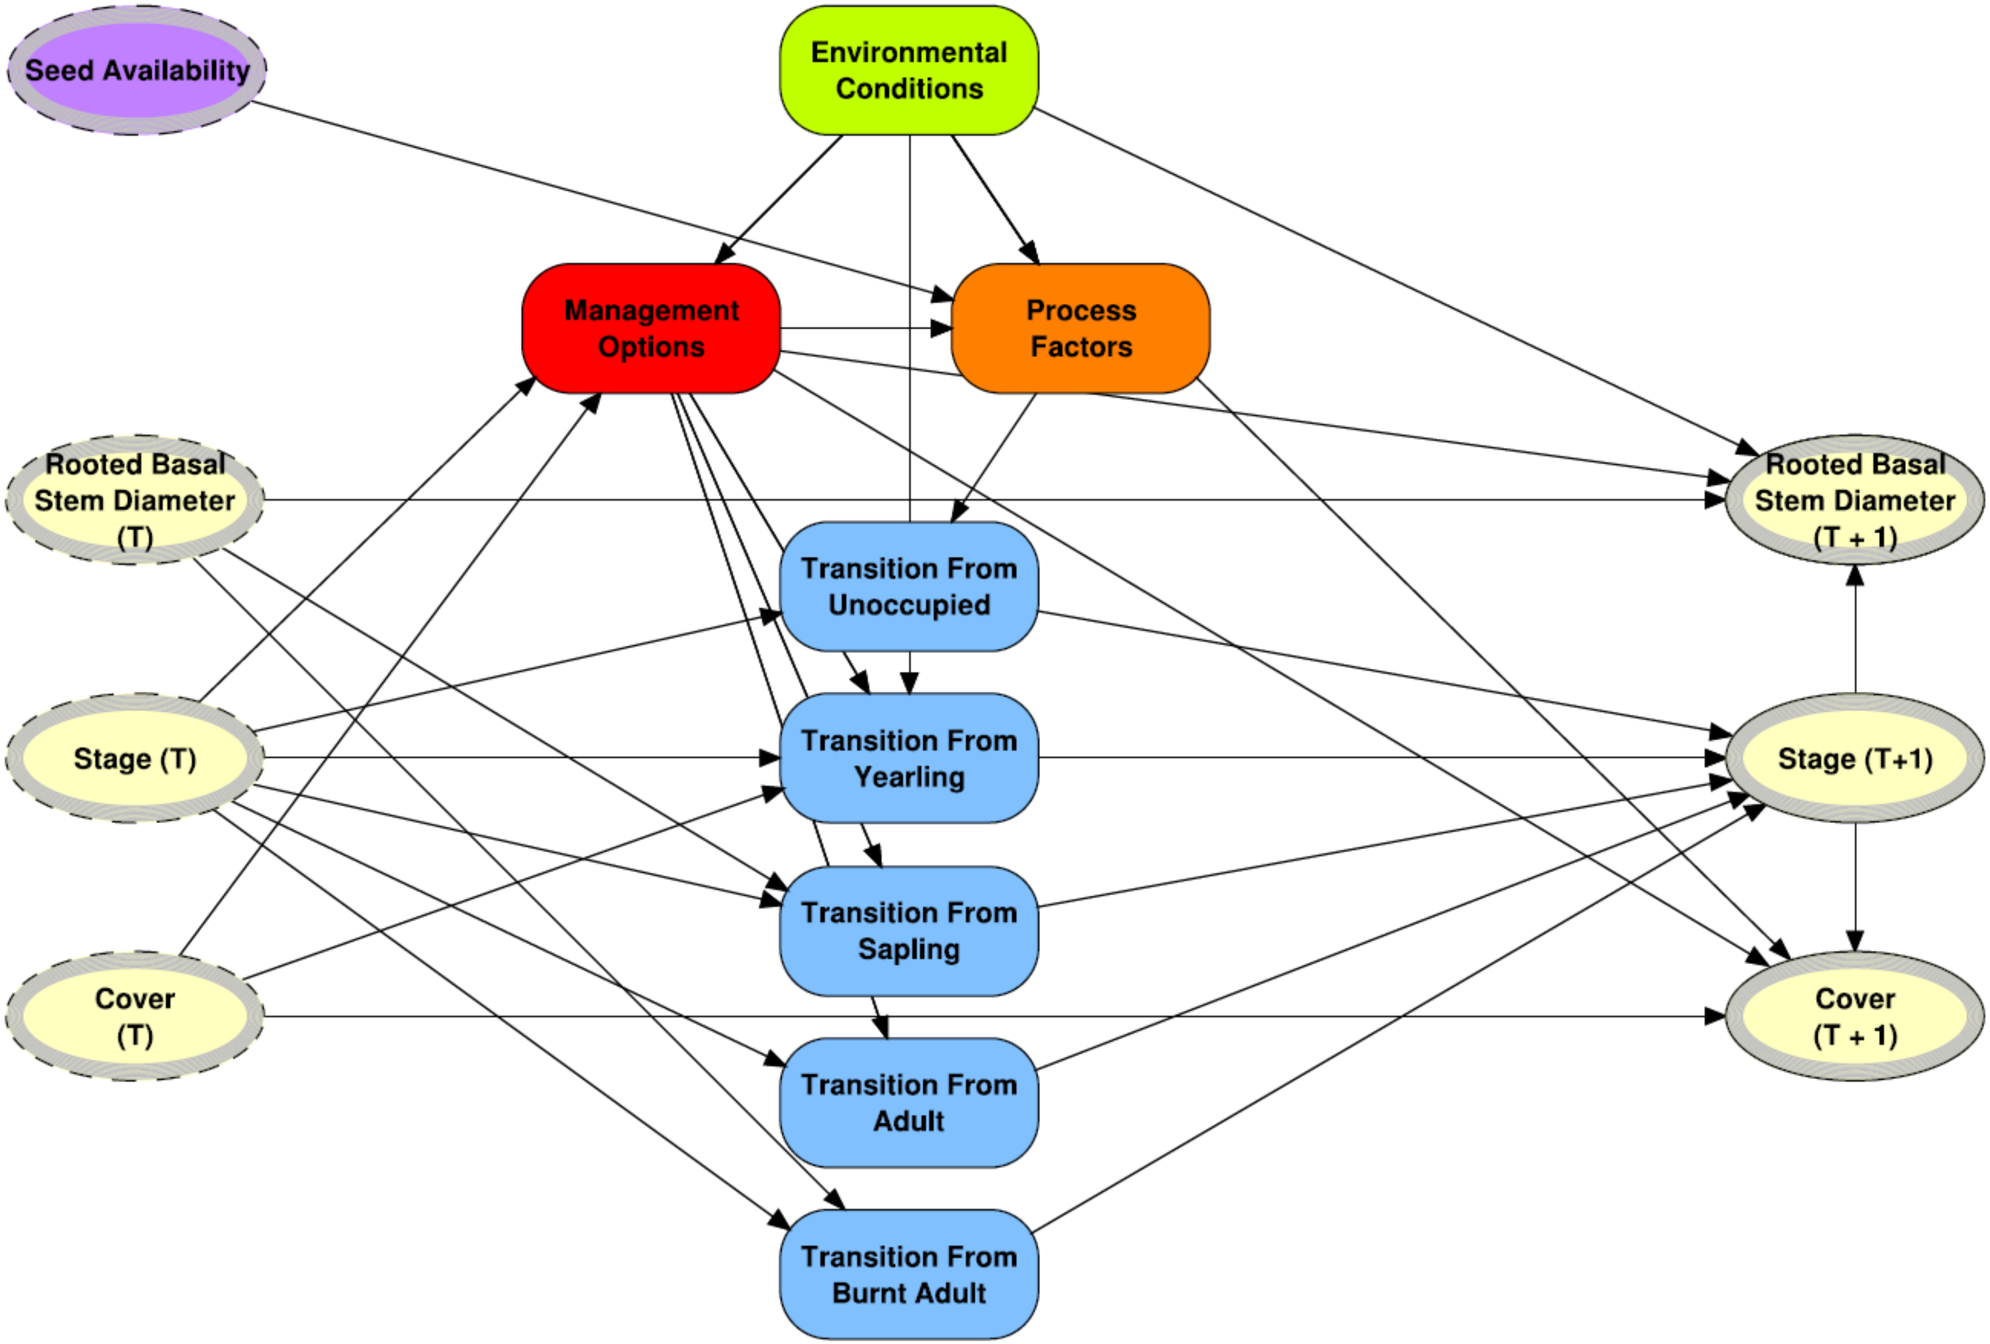

Modelling Seed Production

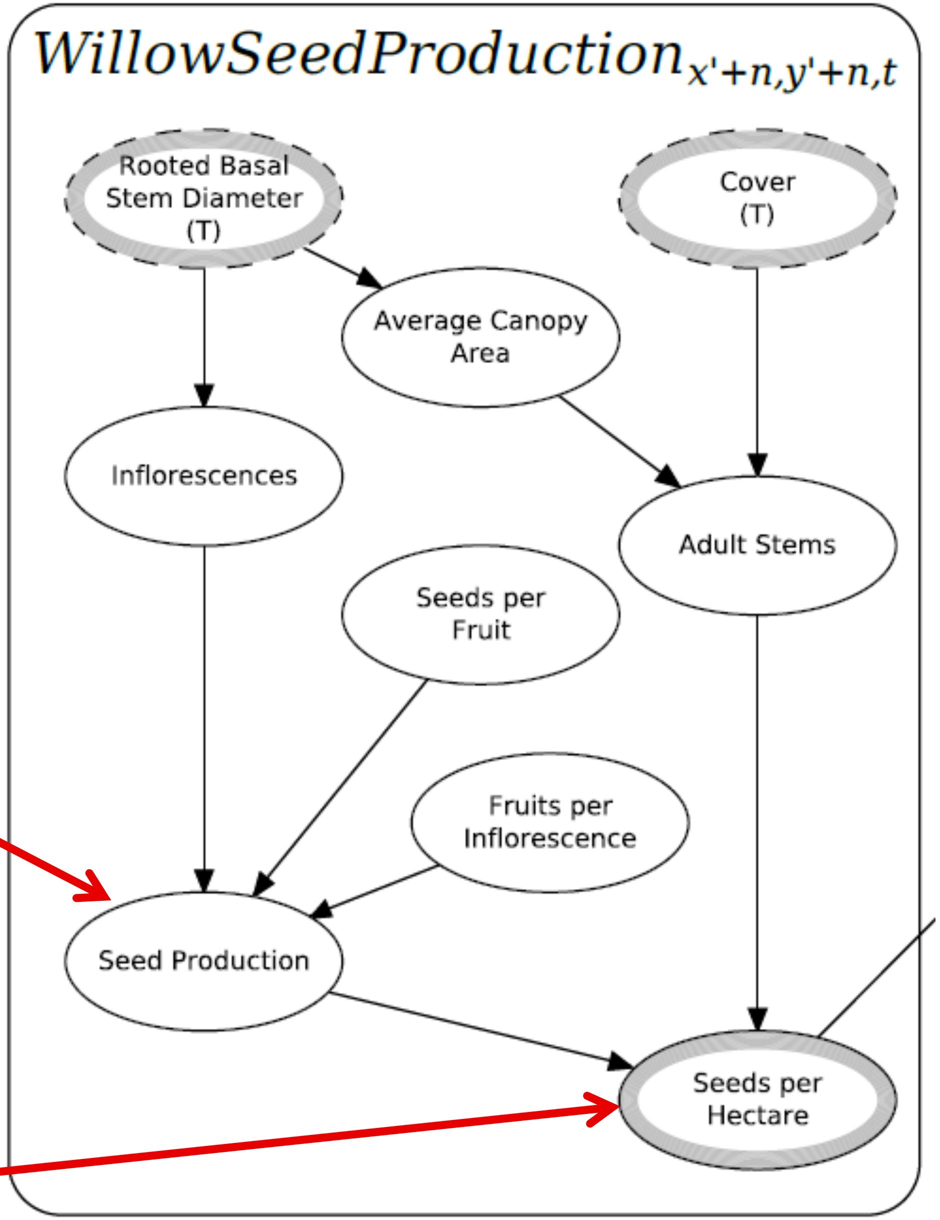

SeedsPerStem = I * F * S

where

I is number of ***inflorescences***

F is the number ***fruits per inflorescence***

S is number of ***seeds per fruit***

SeedsPerHA = SeedsPerStem * NumberOfStemsPerHA

Application of BNs for complex environmental management: Western Grasslands Reserves

(DSE Project 2012-2013) (Sinclair et al., In Preparation)

- 10,000 ha to be restored to native grasslands over 10-20 years
- Task: build a dynamic BN to evaluate “what-if” scenarios over 20 years
  - a range of management strategies
  - for a variety of land types
  - explicitly representing costs and environmental values

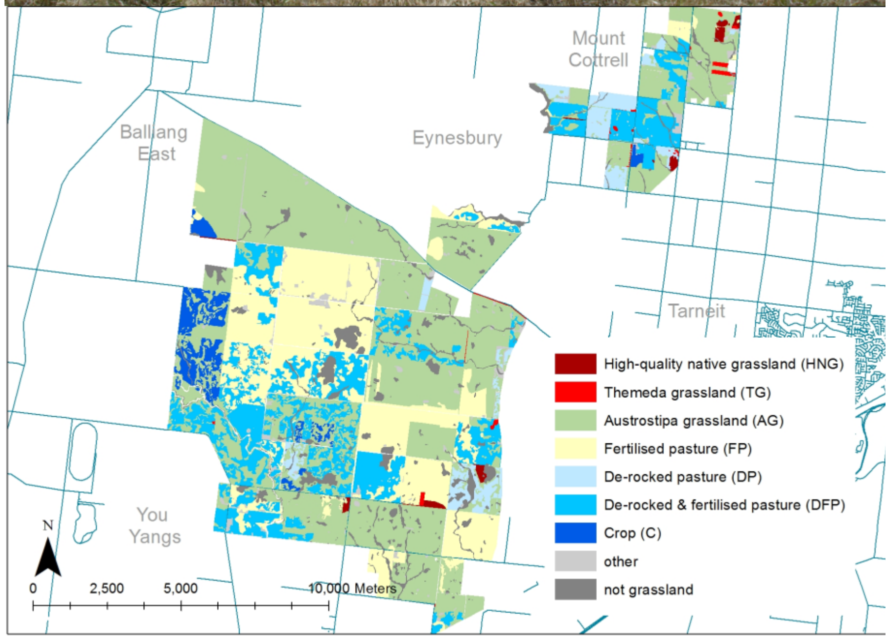

BNs for Risk assessment for Log Exports in NZ

Collaboration with SCION (NZ timber research)

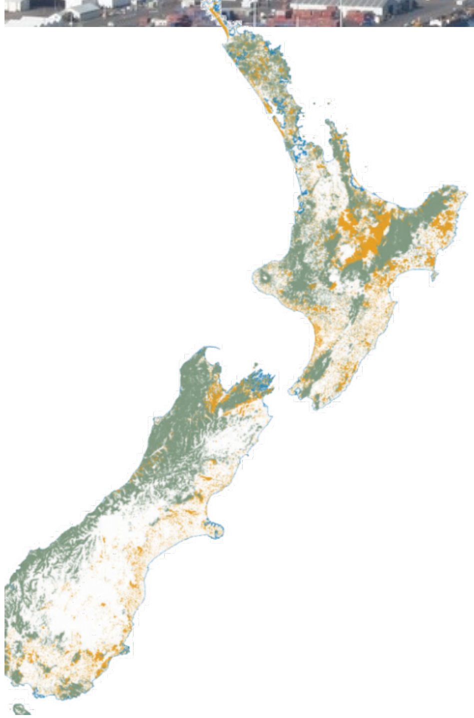

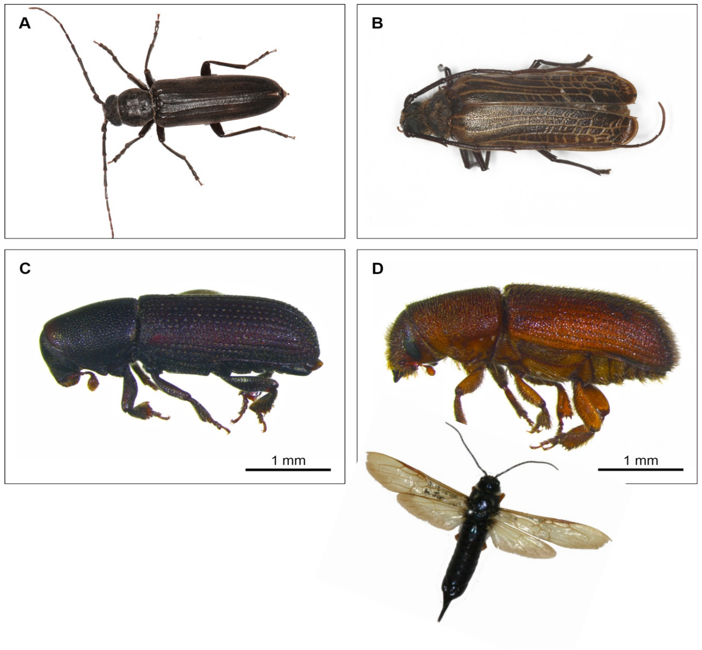

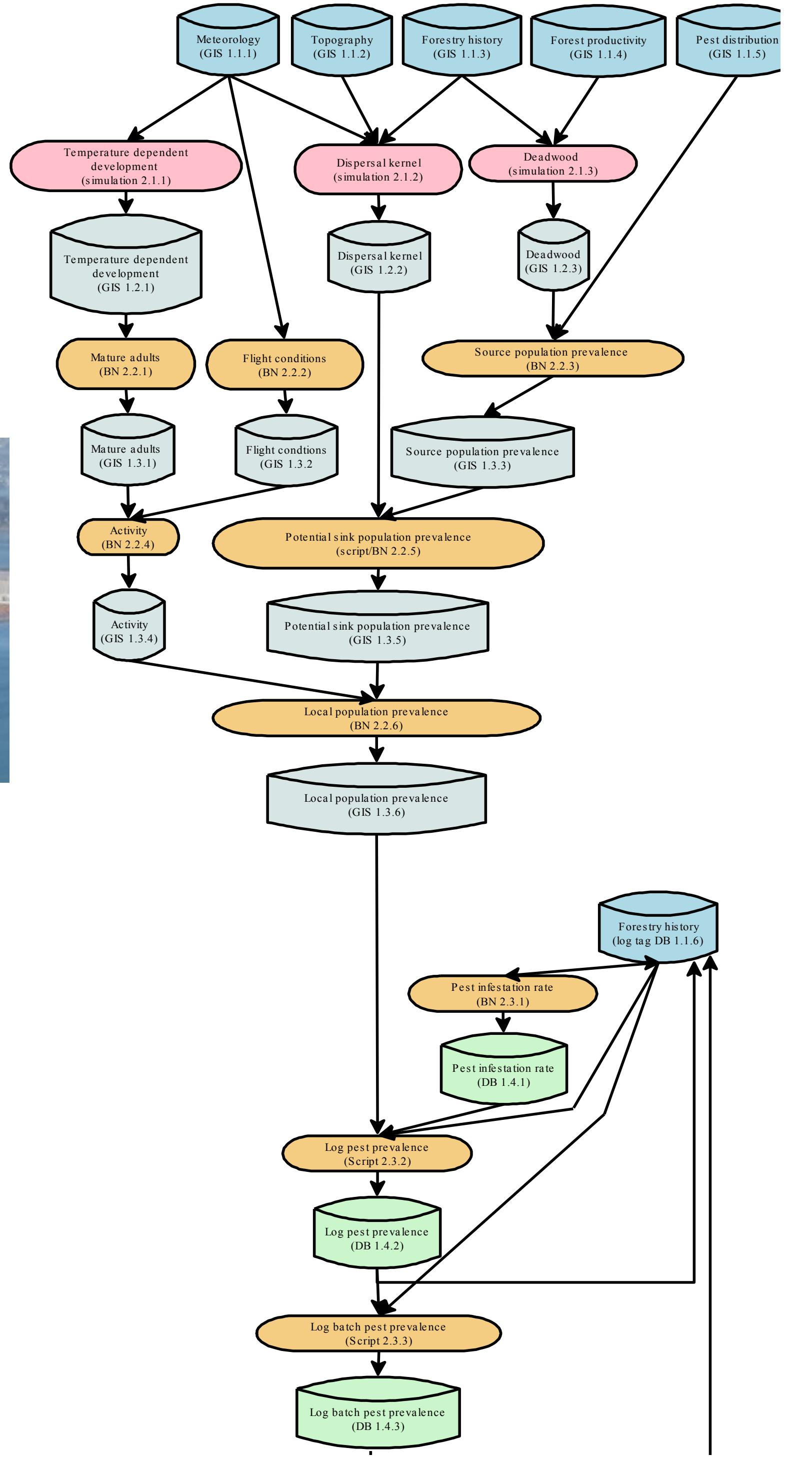

BayesWatch: BNs for anomaly detection in tracking (Collaboration with DSTO) (Mascaro et al., 2014)

- Task: Detect anomalous behaviour of vessels, cars, pedestrians
- Originally used AIS data from vessels in Sydney Harbour

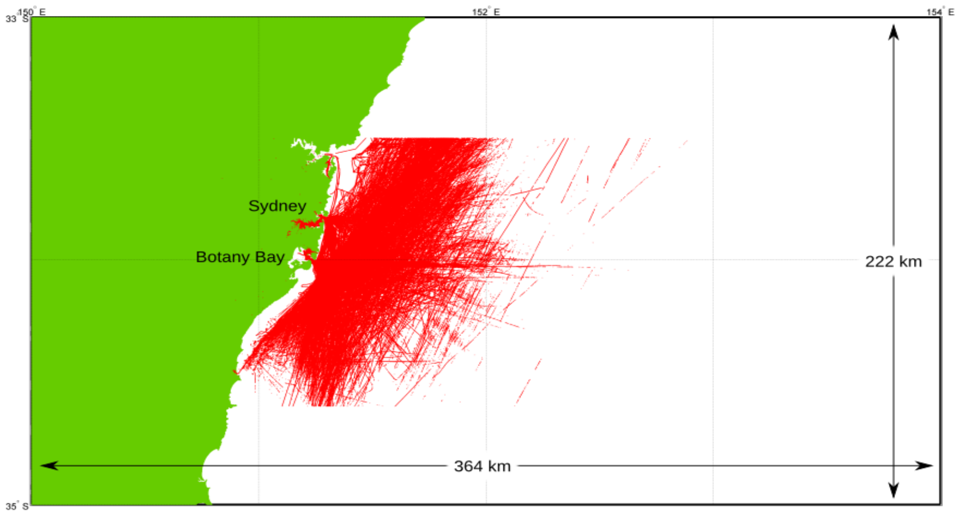

- Apply BN learning to build models of behaviour
- Combines <u>Time series</u> BN + Track <u>Summary</u> BN
- Use metrics to assess anomalous tracks

Modelling bushfire prevention & suppression (Penman et al., 2015)

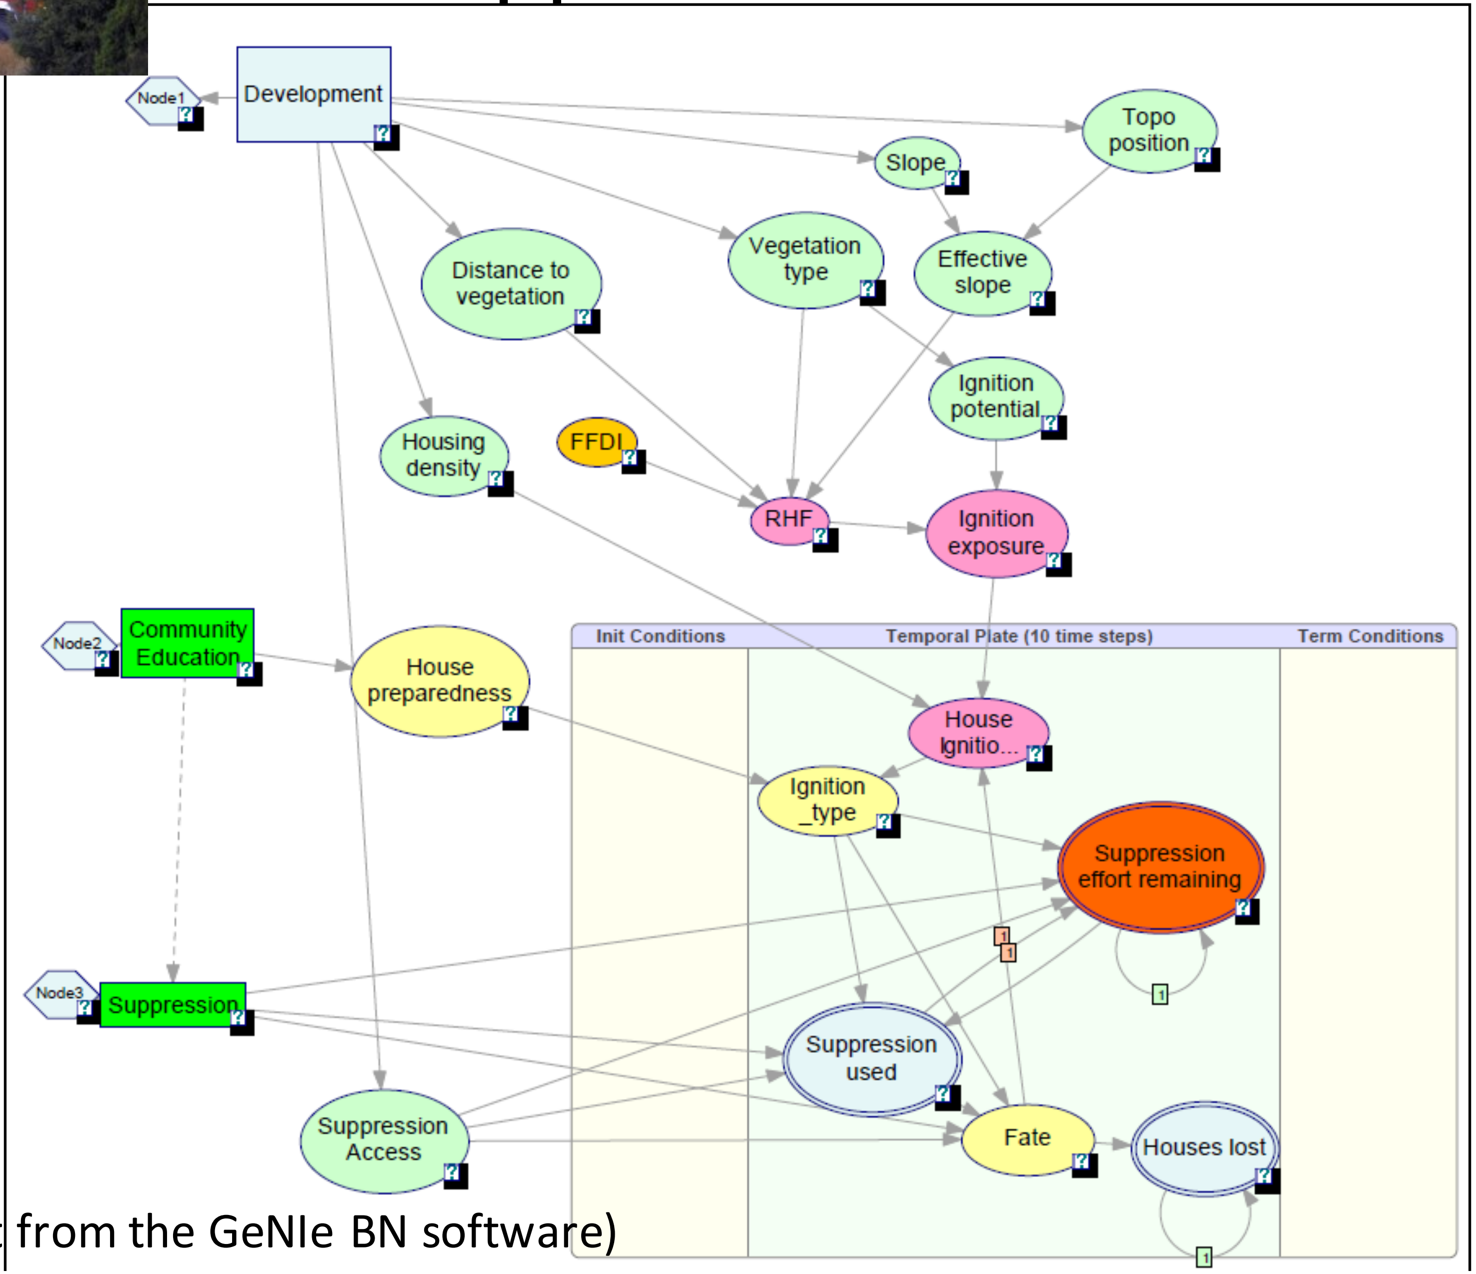

(Screen shot from the GeNIe BN software)

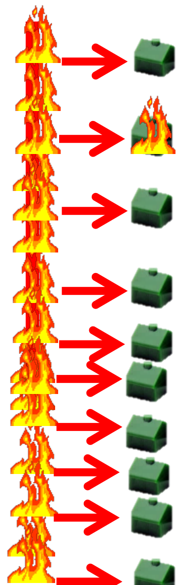

Medical Risk Assessment: Heart Disease

(Nicholson et al., 2008; Flores et al., 2011)

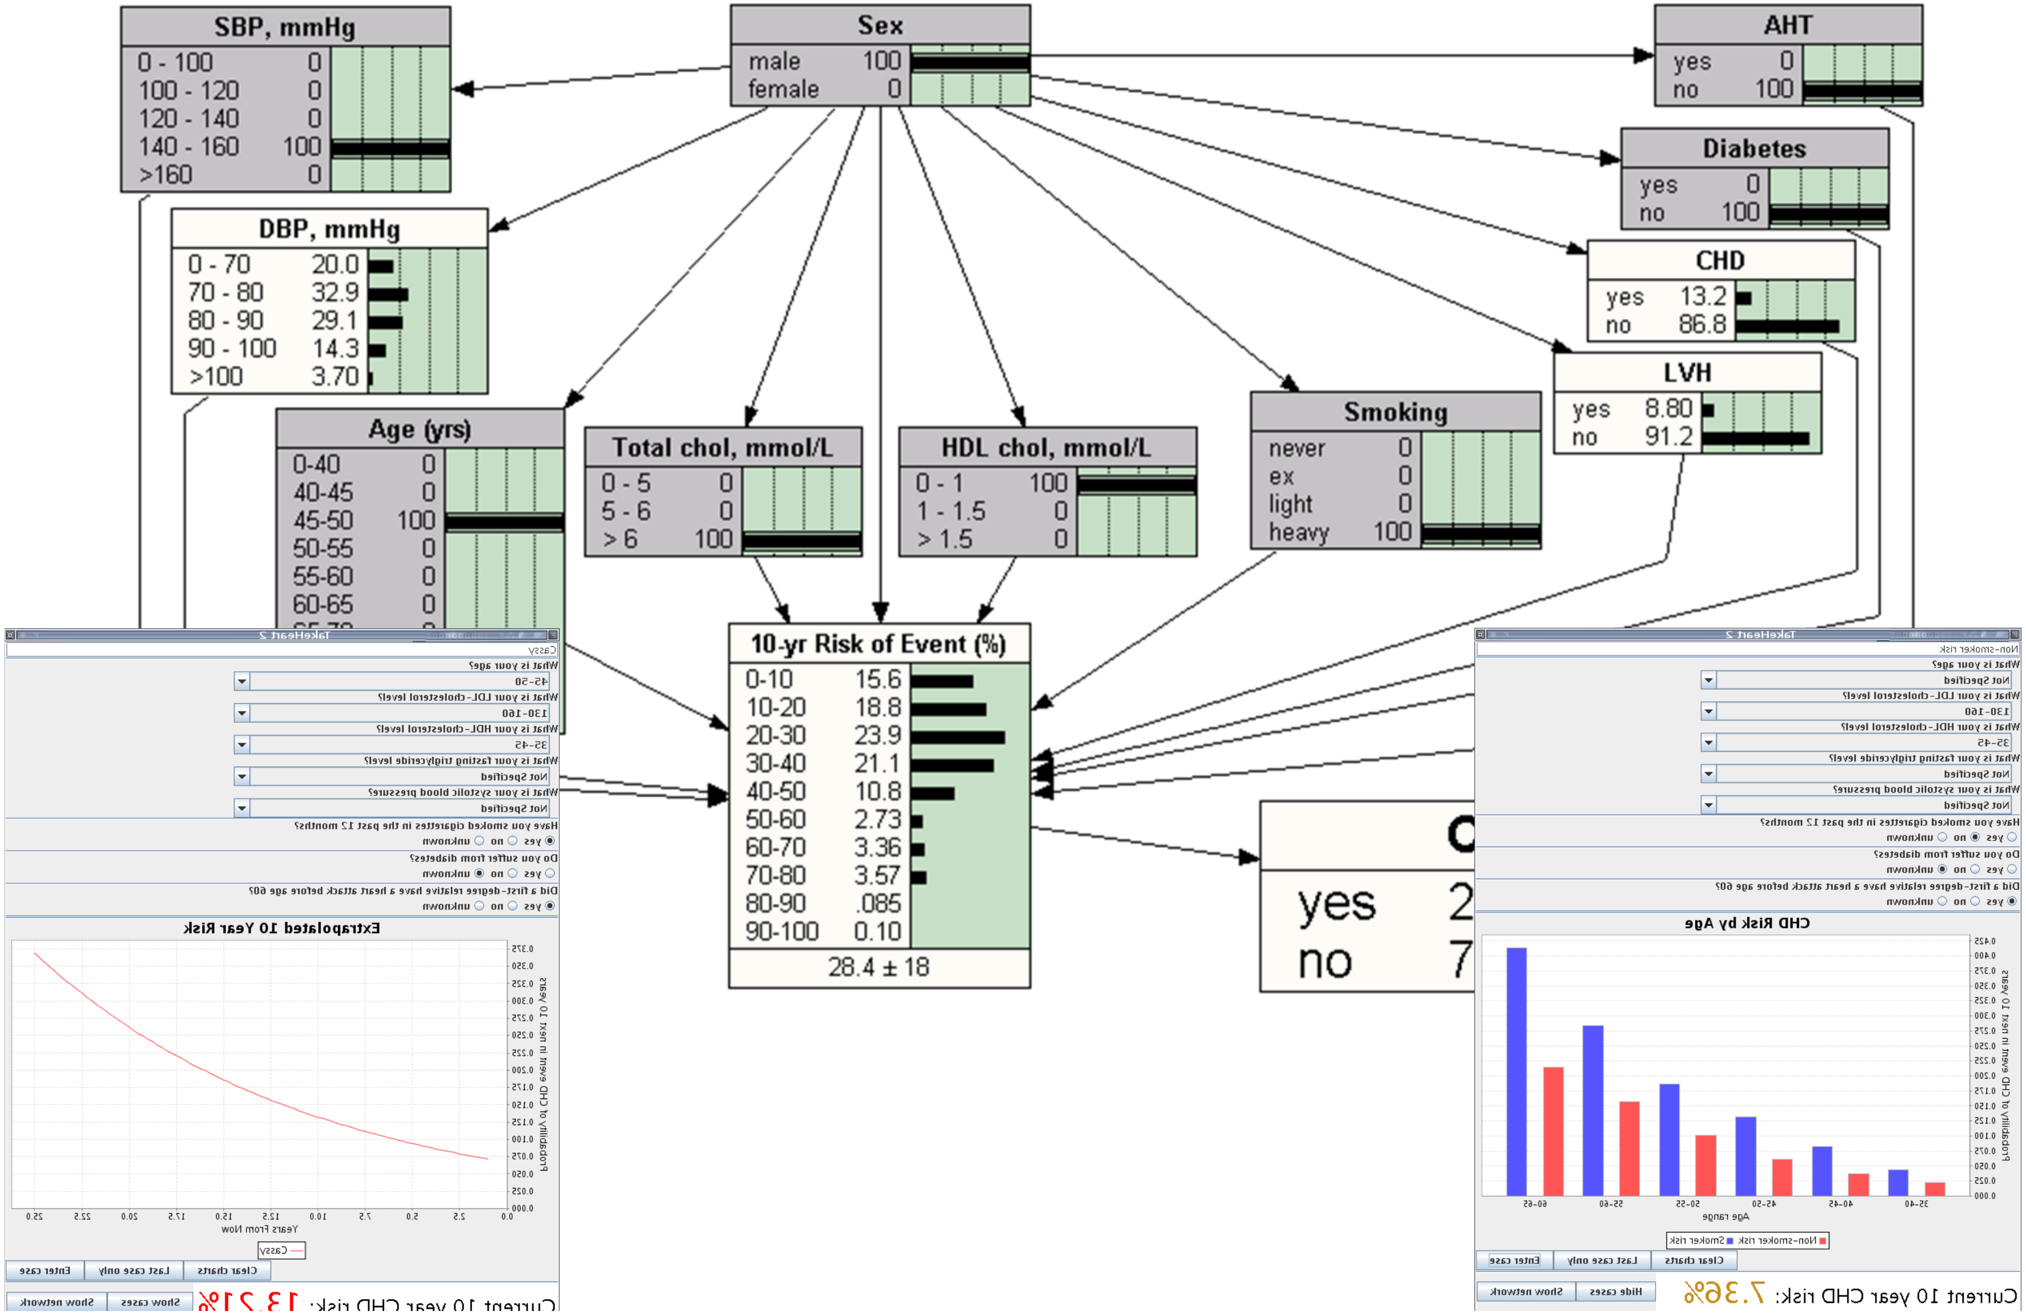

Simple BN model: NZ apples Import Risk Assessment (Wintle et al., 2014). Example of ‘what-if’ reasoning

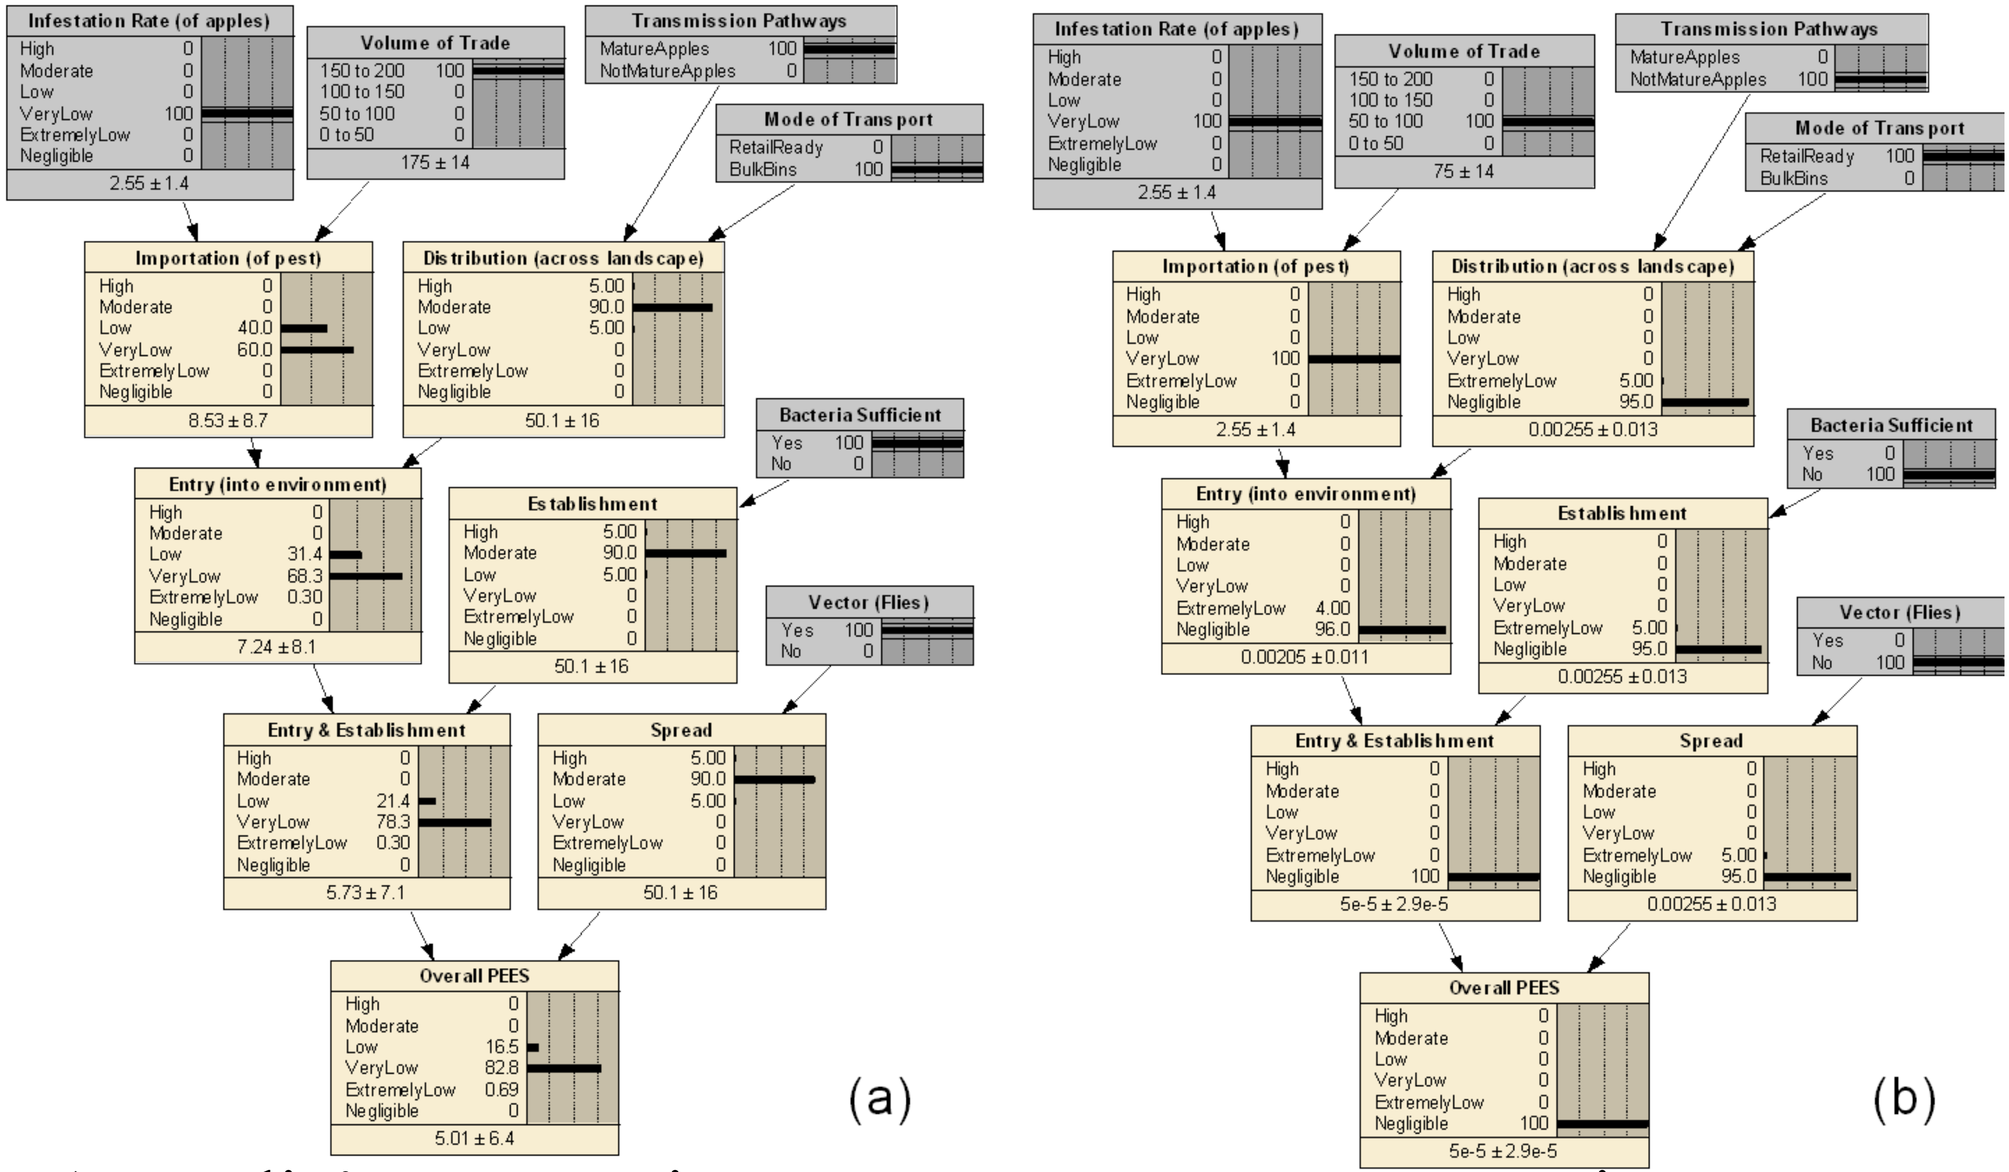

Australia’s assumptions

NZ assumptions

NZ Apples BN (Wintle et al., 2014)
Explicit handling of quantities (of product)

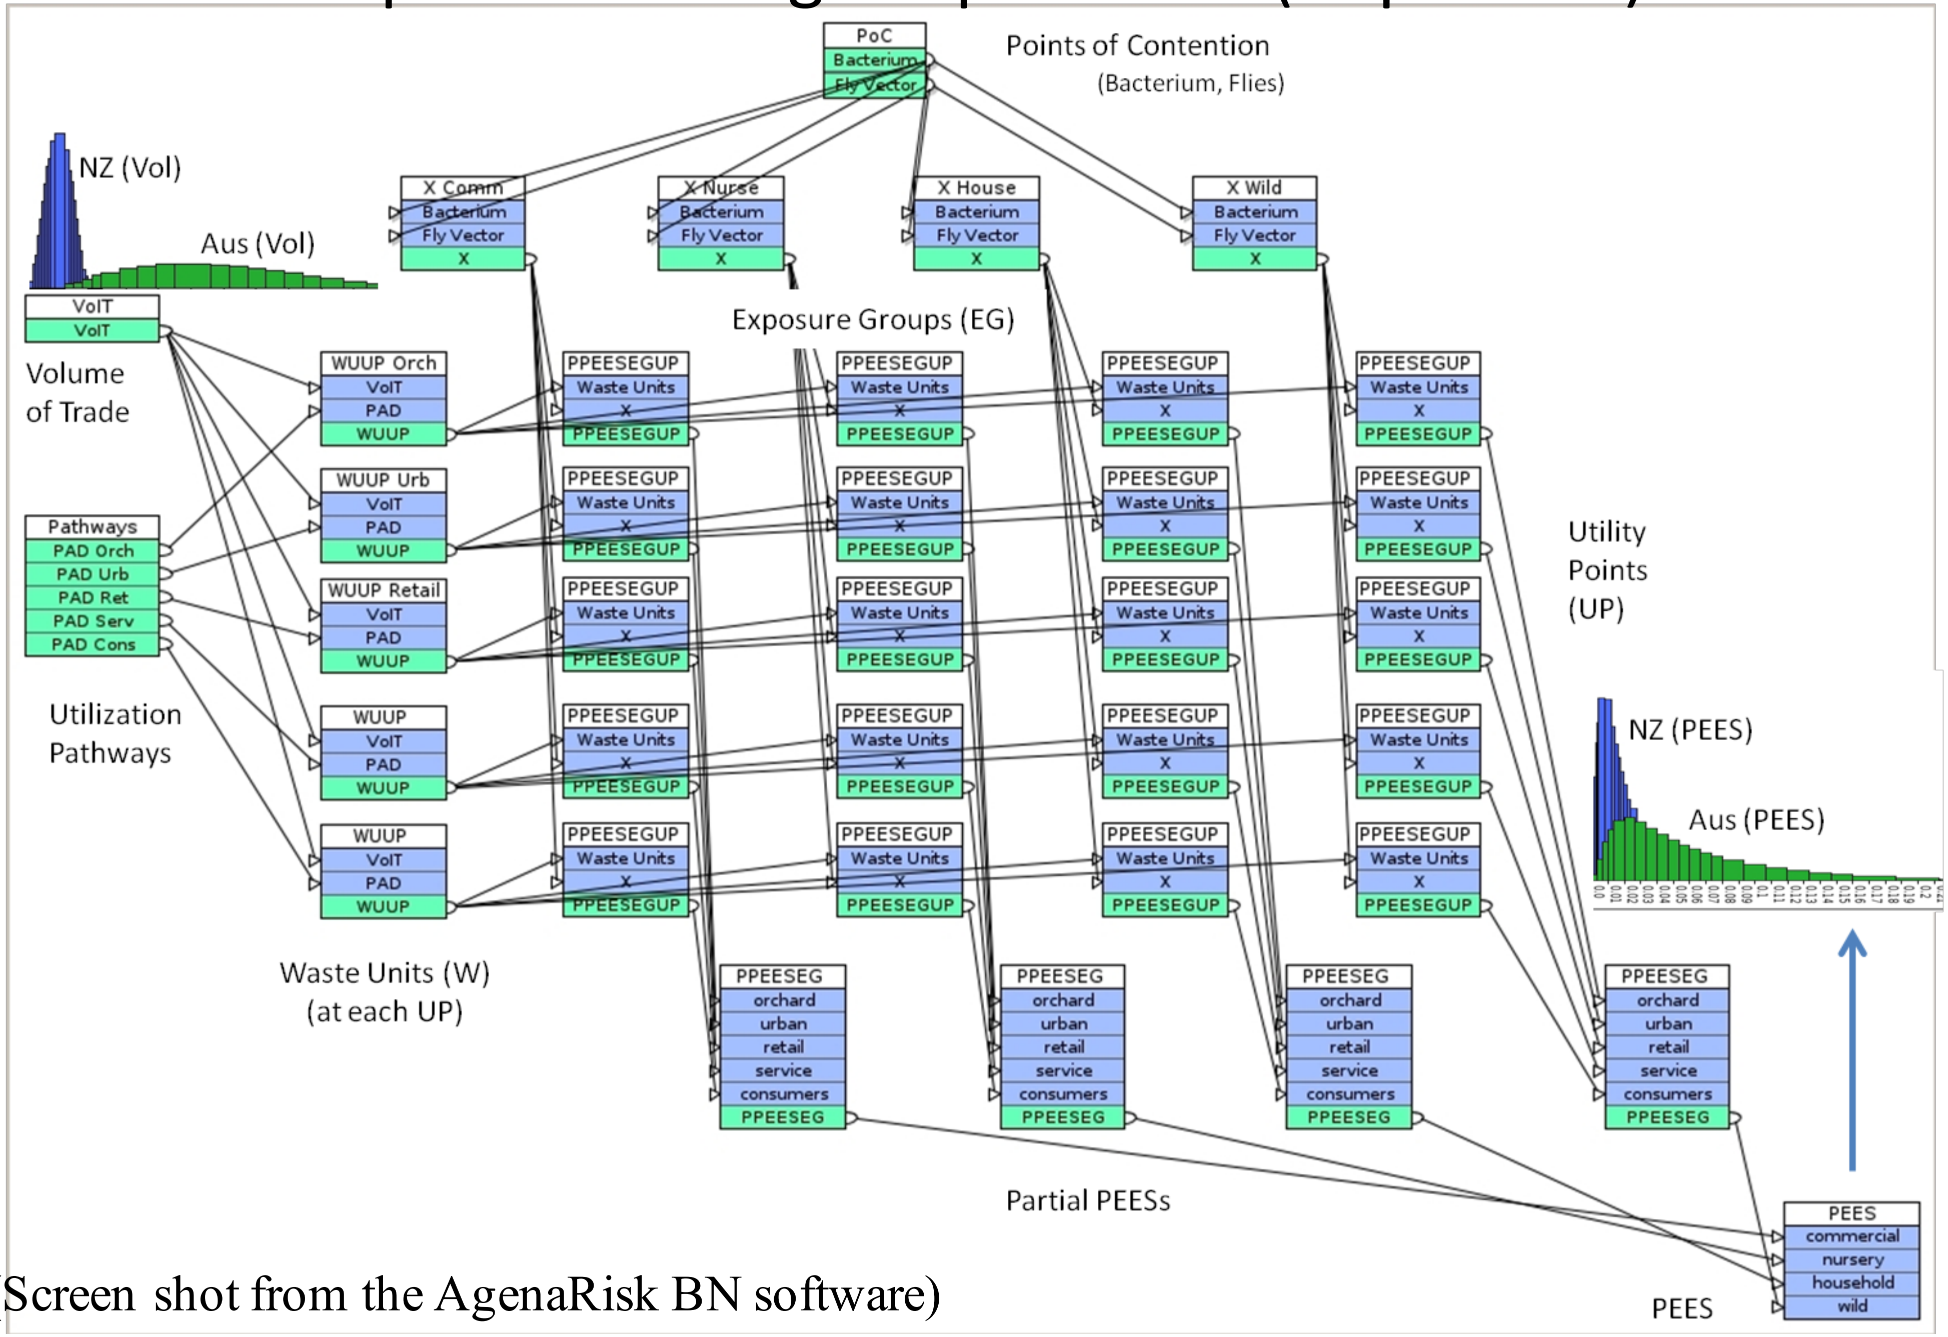

(Screen shot from the AgenaRisk BN software)

Bayesian modelling overview

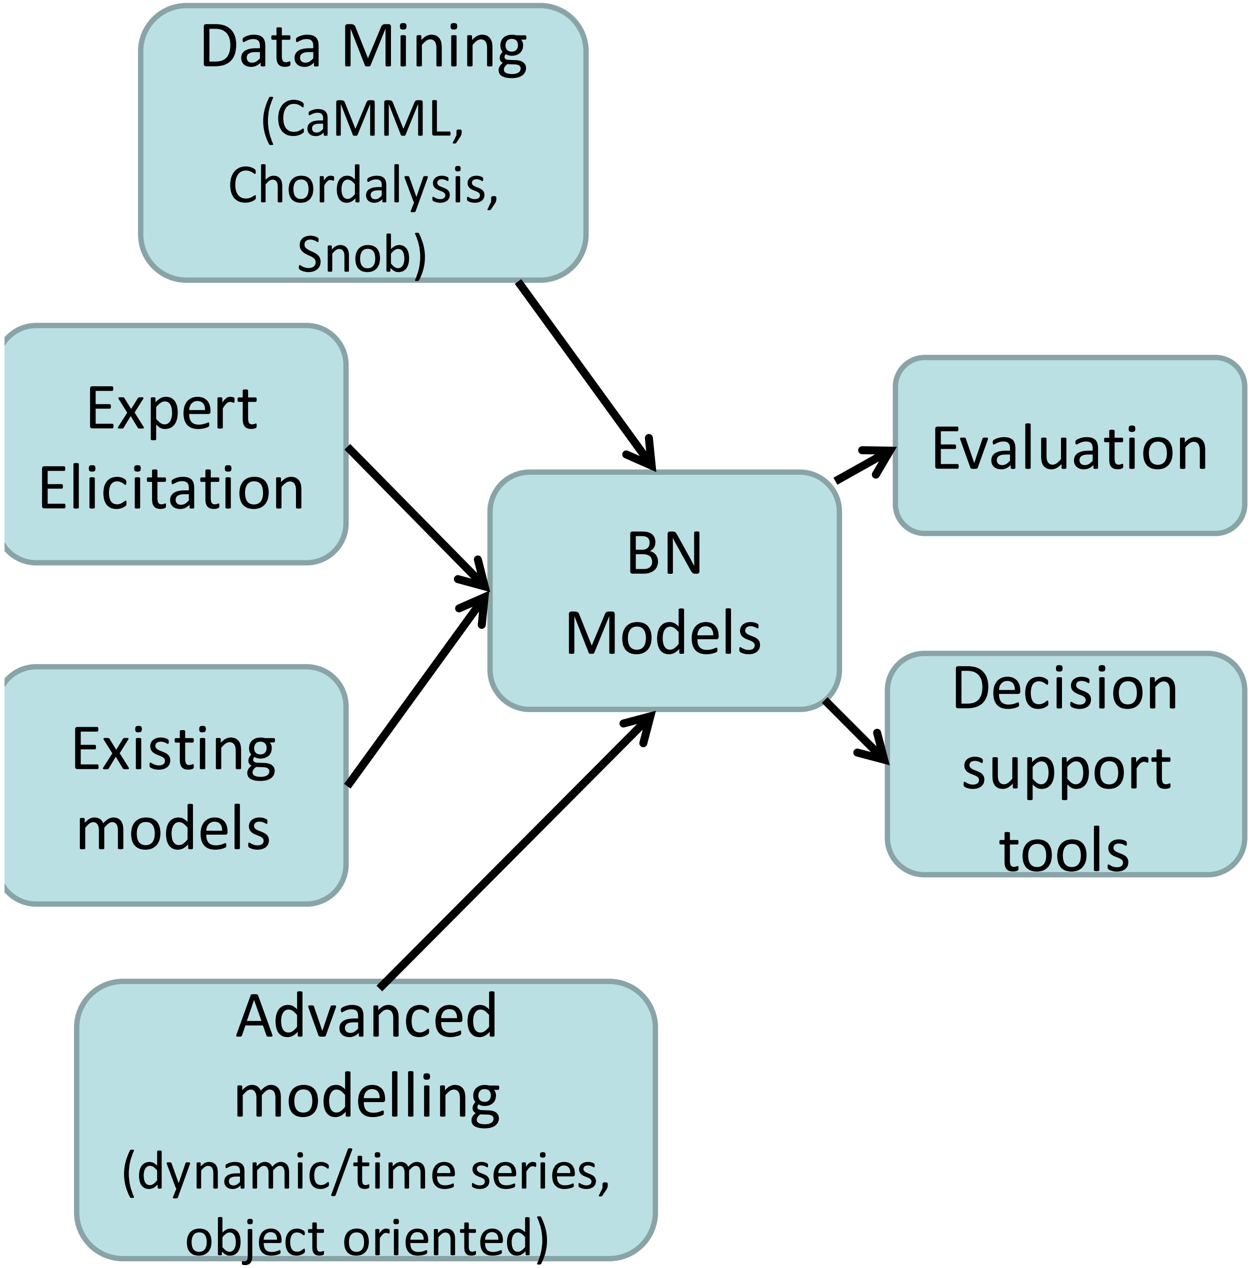

Current research

- Full OOBN framework and methodology
- Learning models with unobserved variables
- Online Delphi-based expert elicitation
- Visualisation of probabilistic outputs