CAT: The Causal Attribution Tool

Steven Mascaro, Kevin B Korb, Erik P Nyberg & Yang Li (Kelvin)

Contents

KK: Introduction
Causal v Probabilistic Reasoning: Tenure Example
Introduction to Causal Criteria
SM: CAT Demo — walk through main features & how it integrates with ABNMS DB
EPN: Adding observations — interpretation and pitfalls

CAT

To try out CAT see http://causalattribution.org:3000/.

To read more about CAT (the Explainer) see http://causalattribution.org:3000/what_is_cat

KK

CAT Introduction

BN tools are quite familiar by now; they've been around for decades.

They have been used for a great variety of tasks:

- Assessing evidence (Fenton, et al., 2016
- Argument Analysis (Nyberg, et al., 2022)
- Modeling & Prediction (Marcot & Penman, 2019; Arora, et al., 2019))
- Explanation
- Hypothetical Reasoning (Glymour & Danks, 2007)

Which of these involve causal reasoning?

KK

CAT Introduction

BN tools are quite familiar by now; they've been around for decades.

They have been used for a great variety of tasks:

- Assessing evidence (Fenton, et al., 2016
- Argument Analysis (Nyberg, et al., 2022)
- Modeling & Prediction (Marcot & Penman, 2019; Arora, et al., 2019))
- Explanation
- Hypothetical Reasoning (Glymour & Danks, 2007)

Which of these involve causal reasoning? Marked in green.

KK

CAT Introduction

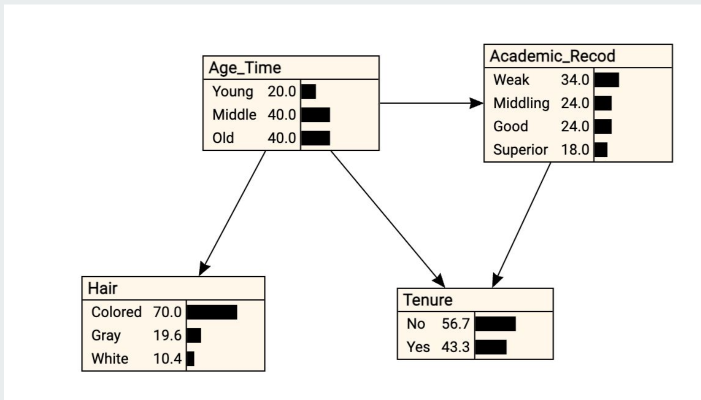

With this model we can reason probabilistically about all kinds of things. E.g., what's the average age of tenured academics?

This is a causal BN, but can't answer causal questions (using ordinary BN tools).

KK

CAT Introduction

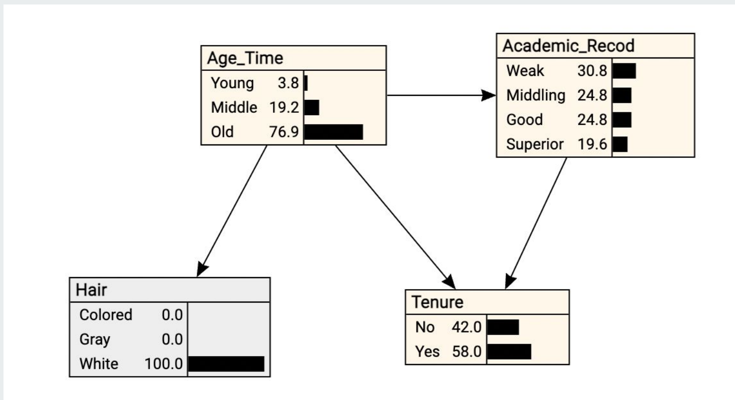

This BN can answer: What’s the probability of tenure given white hair?

What we can’t ask, and get a sensible answer for, is: How will bleaching my hair white affect my chances of getting tenure? (without special hacks)

KK

CAT Introduction

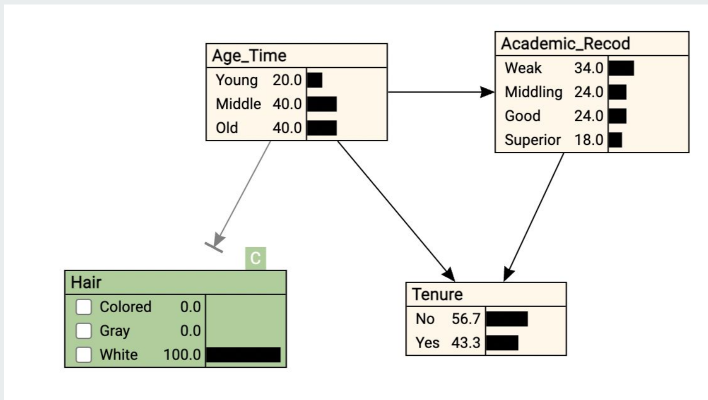

This is the causal model that answers that question.

You can get it by:

- Hacking a BN (but has to be done just right!!)
- Using CAT (simples)

KK

CBNs v BNs

Importantly, CAT can give you the verdict on causal attribution questions:

- Did A cause B?
- How much did A contribute to B? More than C?

According to a variety of causal criteria.

THE UNIVERSITY OF
MELBOURNE

Caveat Emptor: CAT requires its models to be ***causal***. With non-causal networks you will get nonsensical results.

Sourcing, and validating, a causal BN (e.g., machine learning, expert elicitation) is an issue that must precede use of CAT.

KK

CBNs v BNs

It's critical that you establish the model for CAT is ***causal***. Otherwise, CAT's answers will be nonsensical. E.g.,

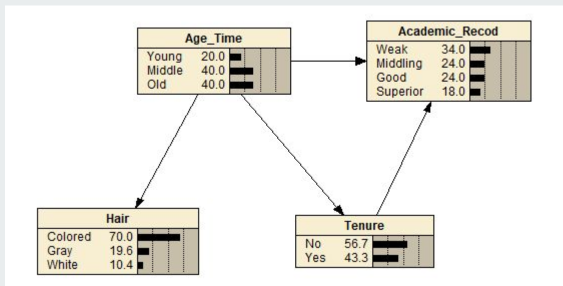

KK

Causal Criteria

Association v Causation

How can we distinguish these?

Scientific verdict usually depends on randomized experiments

Prescientific verdict comes down to observations versus personal interventions. See the psychology of causal reasoning.

KK

Causal Criteria

Association v Causation

How can we distinguish these?
    Scientific verdict usu depends on randomized experiments

With CBNs we can do other things:
    Eyeball
- Does Tenure cause Age?
- Does Age cause Tenure?

    Apply formal criteria

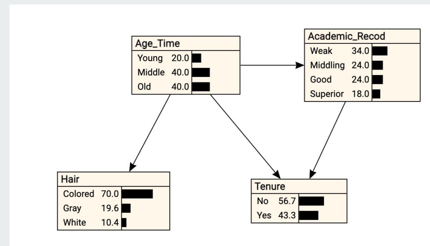

KK

Causal Criteria

David Lewis's Counterfactual Criterion (Lewis, 1973):

> *If A and B are distinct events that actually occur, then A caused B if and only if, were A not to have occurred, B would not have occurred.*

KK

Causal Criteria

David Lewis's Counterfactual Criterion (Lewis, 1973):

> *If A and B are distinct events that actually occur, then A caused B if and only if, were A not to have occurred, B would not have occurred.*

The underlying intuition is close to universal: any cause makes a difference to its effect.

KK

Causal Criteria

David Lewis's Counterfactual Criterion (Lewis, 1973):

> *If A and B are distinct events that actually occur, then A caused B if and only if, were A not to have occurred, B would not have occurred.*

The underlying intuition is close to universal: any cause makes a difference to its effect.

One problem is how to formalize this intuition. CAT provides a platform for doing so, in many different ways.

KK

Causal Criteria

Three formal Causal Criteria are implemented already in CAT:

- Fraction of Attributable Risk
- Cheng’s Causal Power Theory
- Wiggle Theories & Causal Information Theory

SM

ABNMS BN Repository Compatibility

CAT is compatible with the ABNMS BN repository

There are two options:

- Integrated on the same server (buttons allow moving between BN repo and CAT)
- Allow imports from the ABNMS BN repository and vice versa

How CAT will integrate into ABNMS
Uploading CBNs
Loading ABNMS BNs

CAT Demo

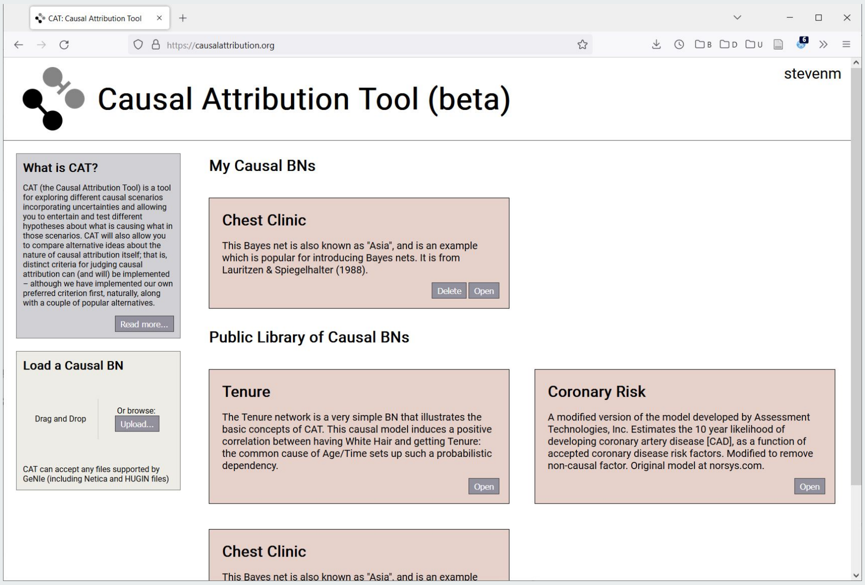

CAT: Causal Attribution Tool
https://causalattribution.org
stevenm
Causal Attribution Tool (beta)
What is CAT?
CAT (the Causal Attribution Tool) is a tool for exploring different causal scenarios incorporating uncertainties and allowing you to entertain and test different hypotheses about what is causing what in those scenarios. CAT will also allow you to compare alternative ideas about the nature of causal attribution itself; that is, distinct criteria for judging causal attribution can (and will) be implemented – although we have implemented our own preferred criterion first, naturally, along with a couple of popular alternatives.
Read more...
Load a Causal BN
Drag and Drop
Or browse:
Upload...
CAT can accept any files supported by GeNIe (including Netica and HUGIN files)
My Causal BNs
Chest Clinic
This Bayes net is also known as "Asia", and is an example which is popular for introducing Bayes nets. It is from Lauritzen & Spiegelhalter (1988).
Delete
Open
Public Library of Causal BNs
Tenure
The Tenure network is a very simple BN that illustrates the basic concepts of CAT. This causal model induces a positive correlation between having White Hair and getting Tenure: the common cause of Age/Time sets up such a probabilistic dependency.
Open
Coronary Risk
A modified version of the model developed by Assessment Technologies, Inc. Estimates the 10 year likelihood of developing coronary artery disease [CAD], as a function of accepted coronary disease risk factors. Modified to remove non-causal factor. Original model at norsys.com.
Open
Chest Clinic
This Bayes net is also known as "Asia", and is an example

CAT Demo

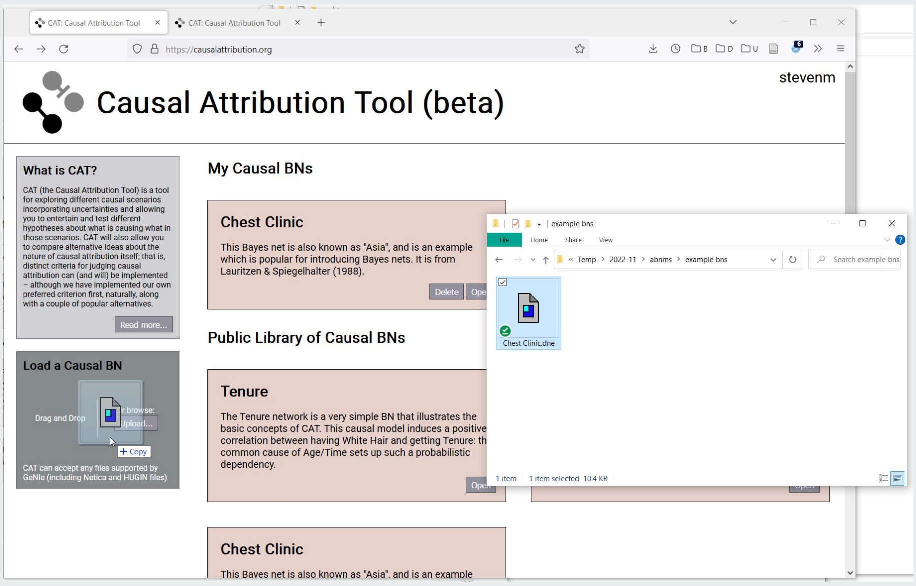

CAT: Causal Attribution Tool
https://causalattribution.org
stevenm
Causal Attribution Tool (beta)
What is CAT?
CAT (the Causal Attribution Tool) is a tool for exploring different causal scenarios incorporating uncertainties and allowing you to entertain and test different hypotheses about what is causing what in those scenarios. CAT will also allow you to compare alternative ideas about the nature of causal attribution itself; that is, distinct criteria for judging causal attribution can (and will) be implemented – although we have implemented our own preferred criterion first, naturally, along with a couple of popular alternatives.
Read more...
Load a Causal BN
Drag and Drop
+ Copy
CAT can accept any files supported by GeNIe (including Netica and HUGIN files)
My Causal BNs
Chest Clinic
This Bayes net is also known as "Asia", and is an example which is popular for introducing Bayes nets. It is from Lauritzen & Spiegelhalter (1988).
Delete
Public Library of Causal BNs
Tenure
The Tenure network is a very simple BN that illustrates the basic concepts of CAT. This causal model induces a positive correlation between having White Hair and getting Tenure: the common cause of Age/Time sets up such a probabilistic dependency.
Chest Clinic
This Bayes net is also known as "Asia", and is an example
example bns
File
Home
Share
View
Temp
2022-11
abnms
example bns
Search example bns
Chest Clinic.dne
1 item
1 item selected 10.4 KB

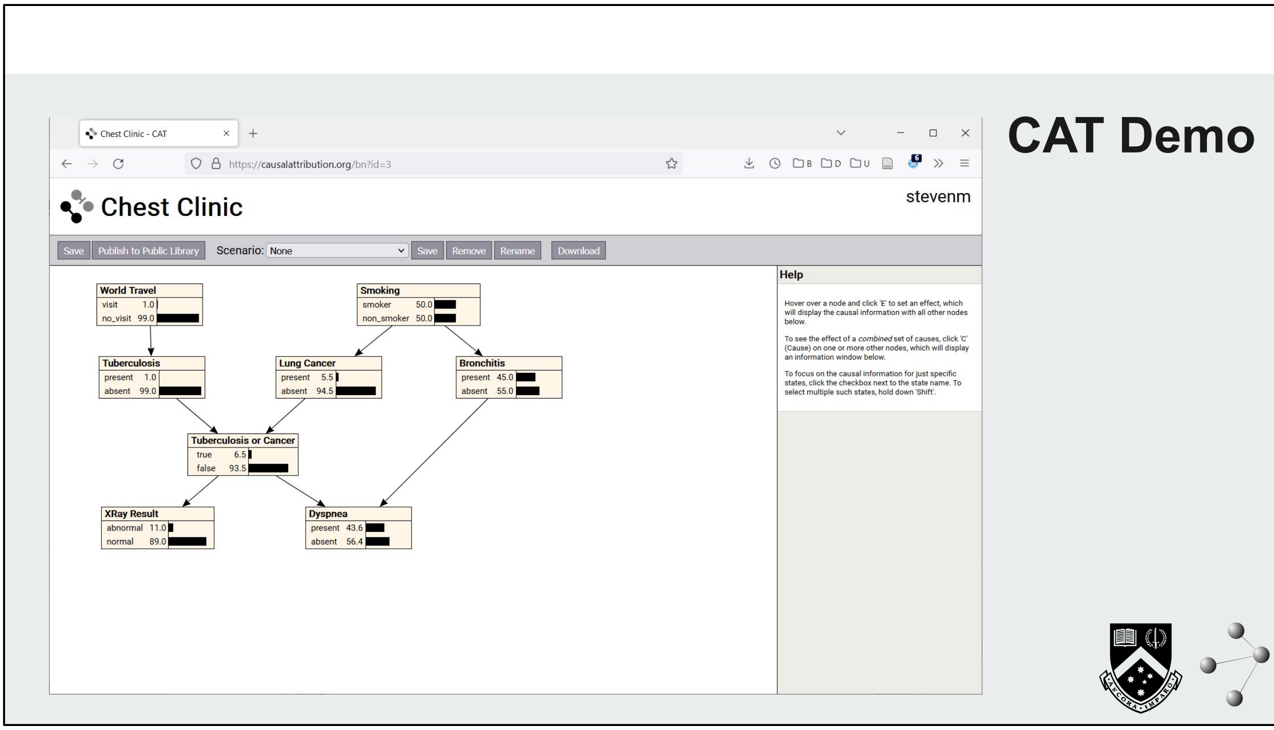

Using CAT for BN updating/reasoning

Using CAT for causal reasoning

Hypothetical interventions

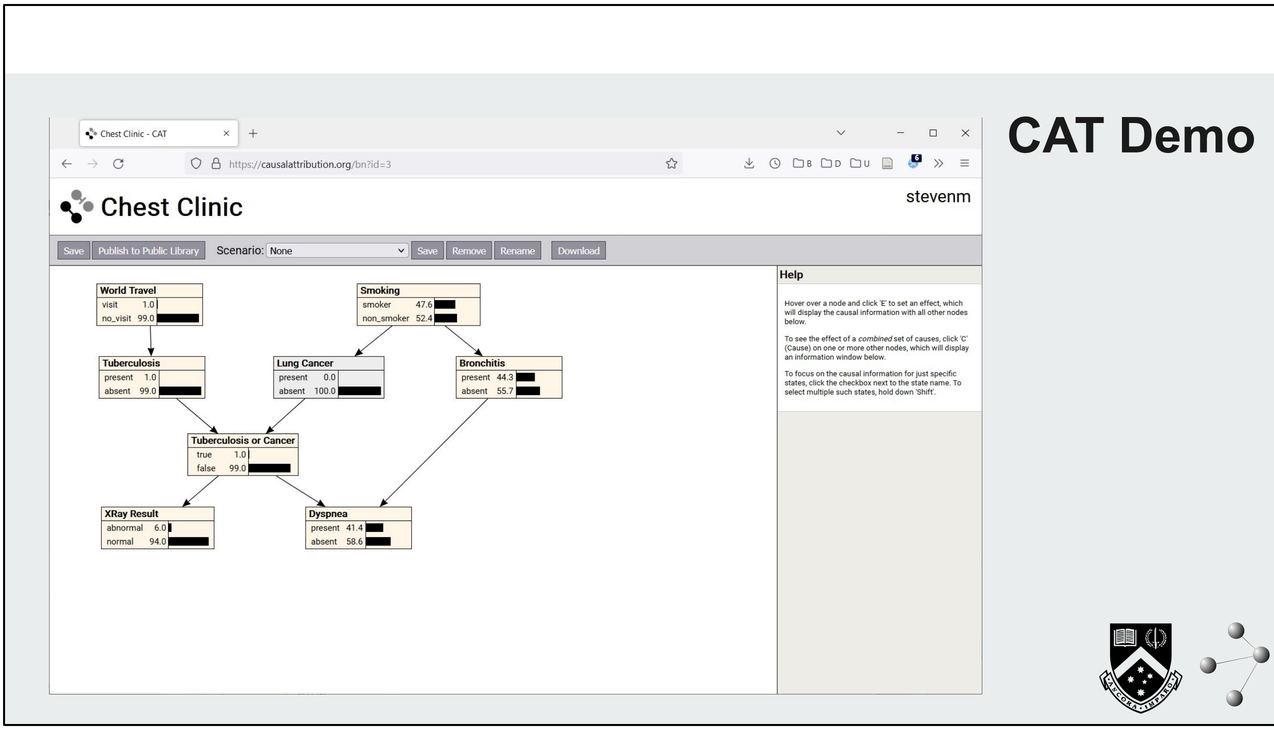

Using CAT for BN updating/reasoning

Using CAT for causal reasoning

Hypothetical interventions

CAT Demo

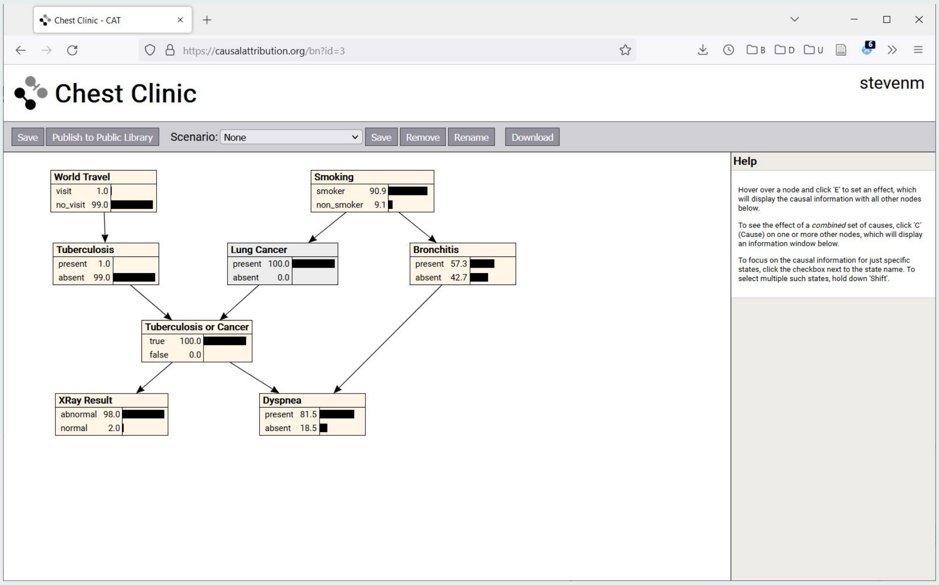

Using CAT for BN updating/reasoning
Using CAT for causal reasoning
- Hypothetical interventions

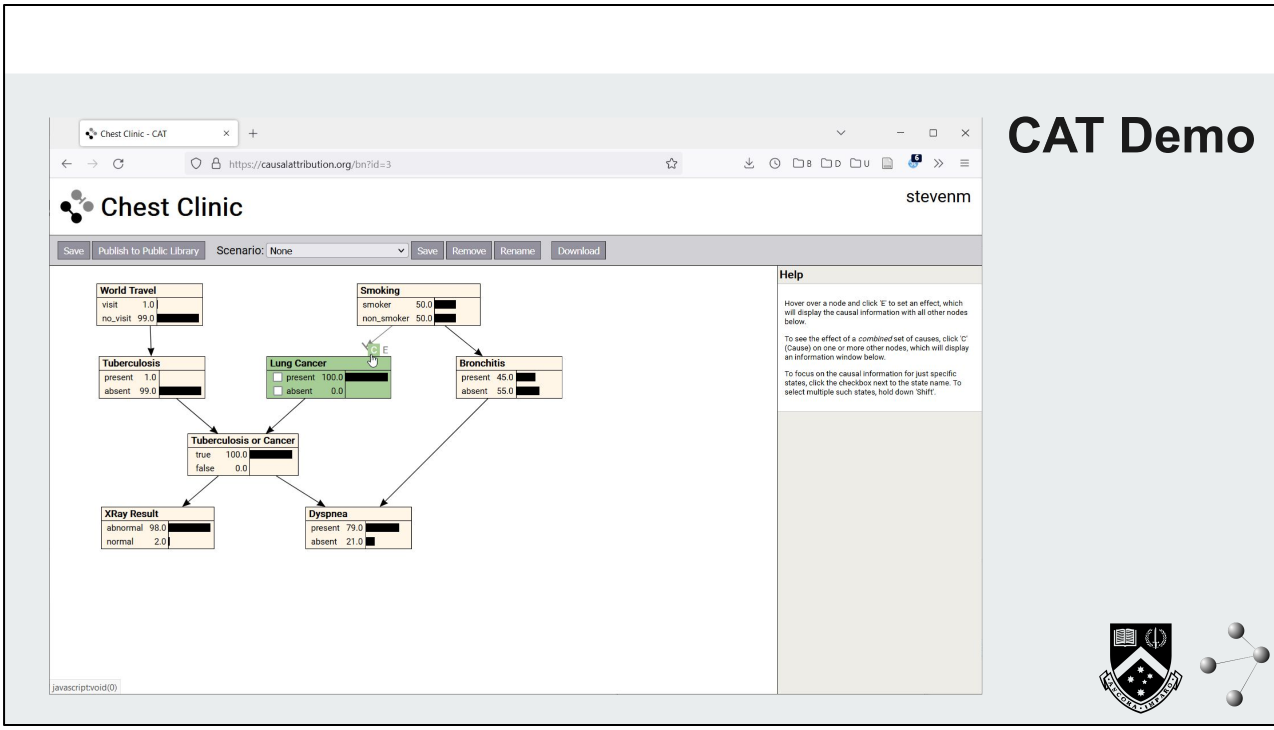

Using CAT for BN updating/reasoning

Using CAT for causal reasoning

Hypothetical interventions

CAT Demo

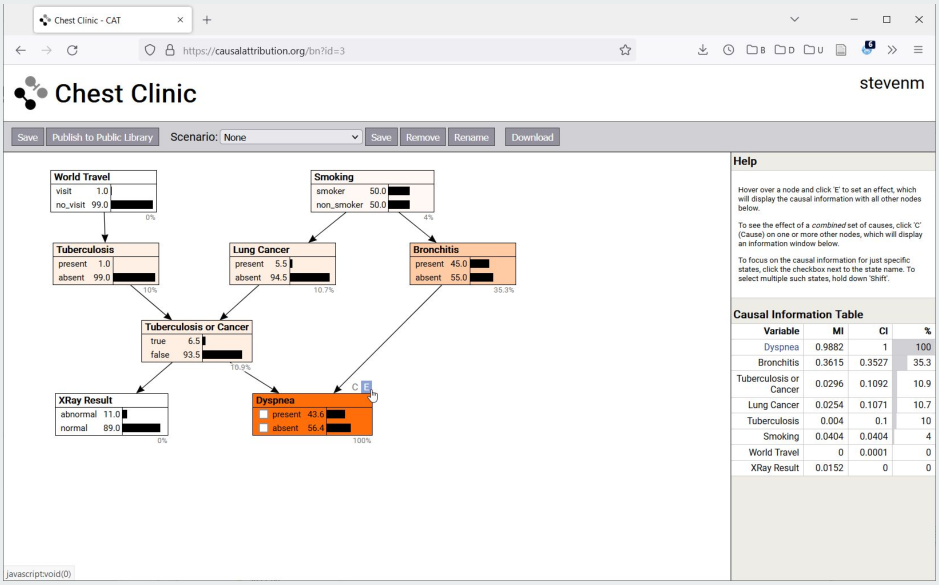

CAT measures

- Mutual Information
- Cheng's Causal Power
- FAR
- CI

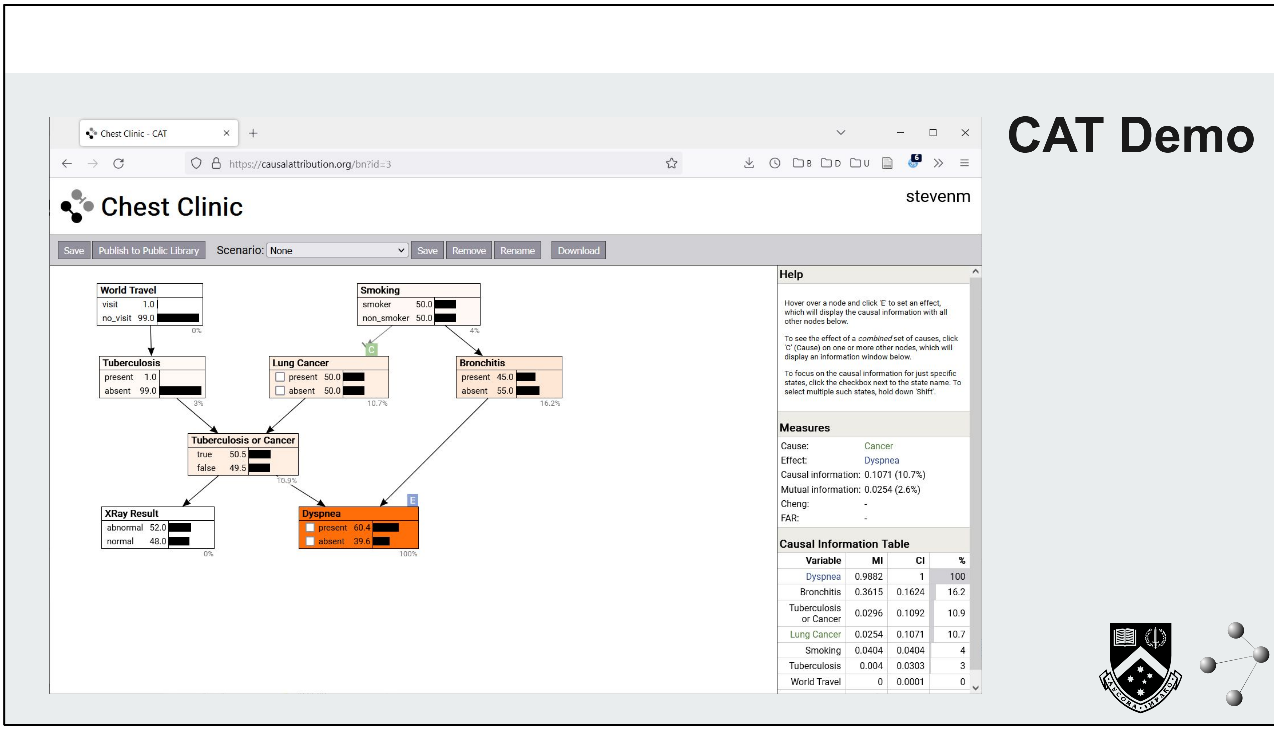

CAT measures

- Mutual Information
- Cheng’s Causal Power
- FAR
- CI

CAT Demo

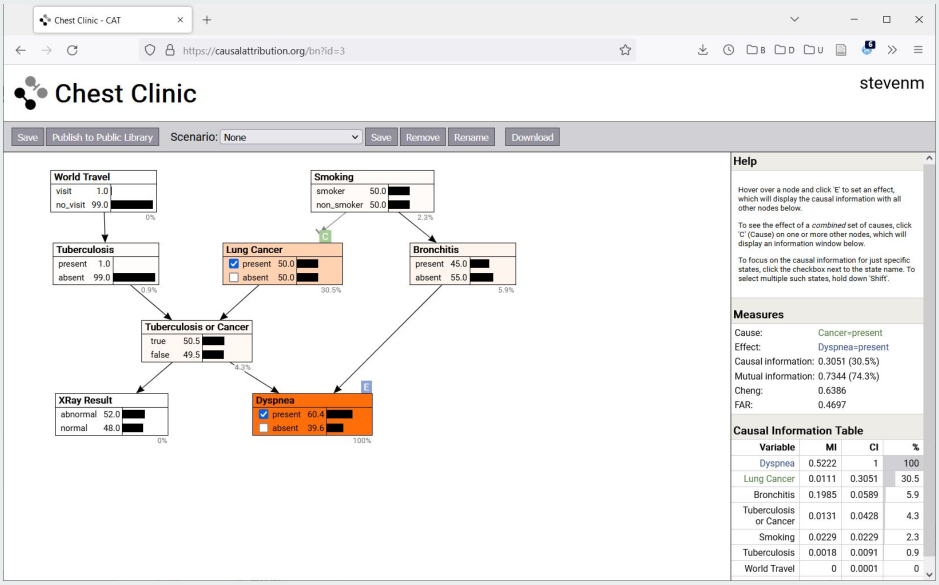

CAT measures

- Mutual Information
- Cheng’s Causal Power
- FAR
- CI

CAT Demo

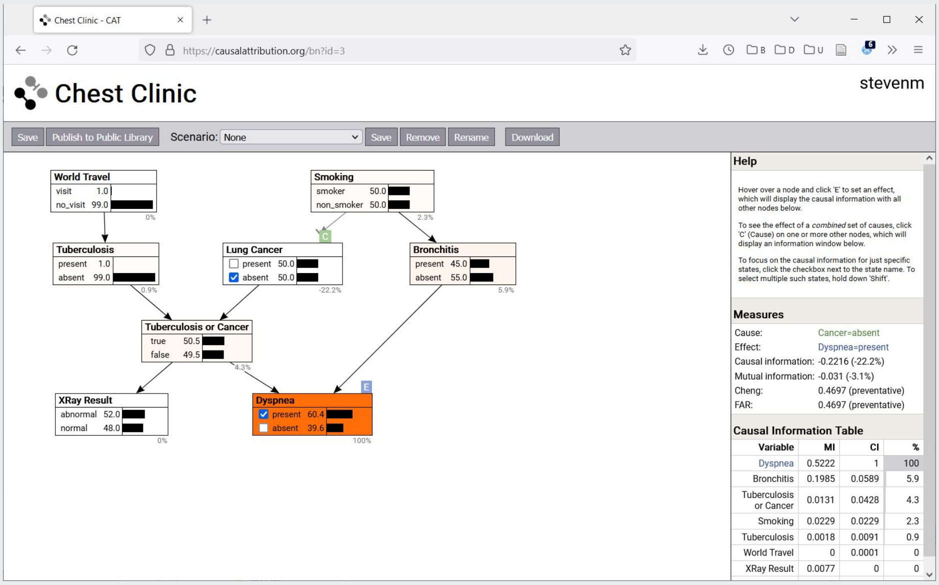

CAT measures

- Mutual Information
- Cheng's Causal Power
- FAR
- CI

CAT Demo

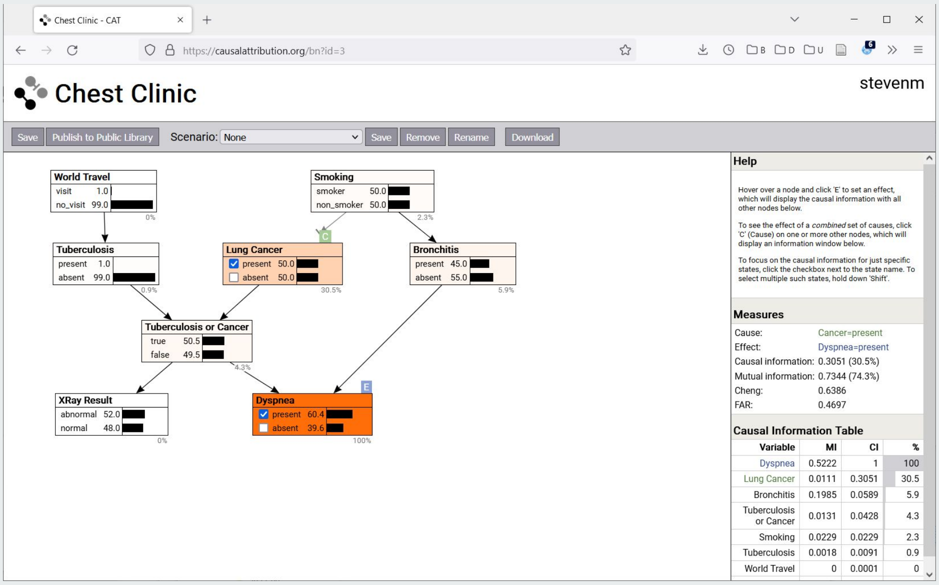

CAT measures

- Mutual Information
- Cheng's Causal Power
- FAR
- CI

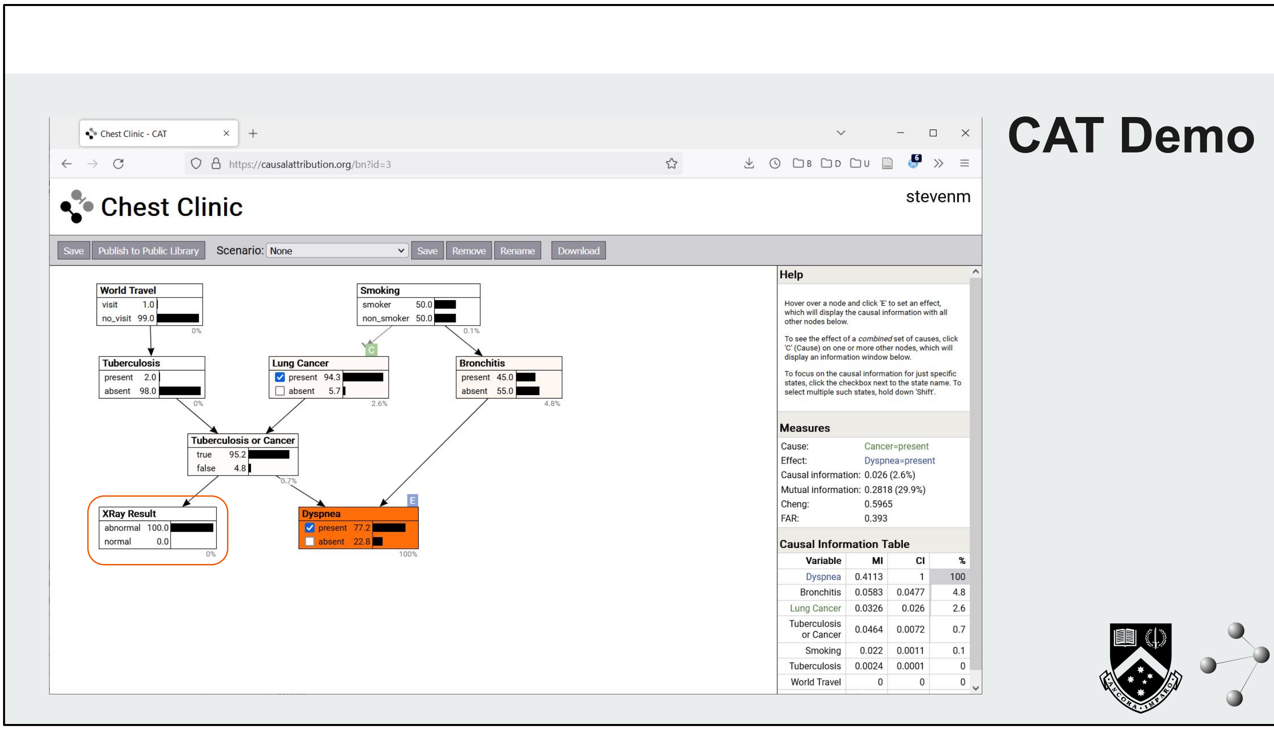

CAT measures

- Mutual Information
- Cheng's Causal Power
- FAR
- CI

CAT Demo

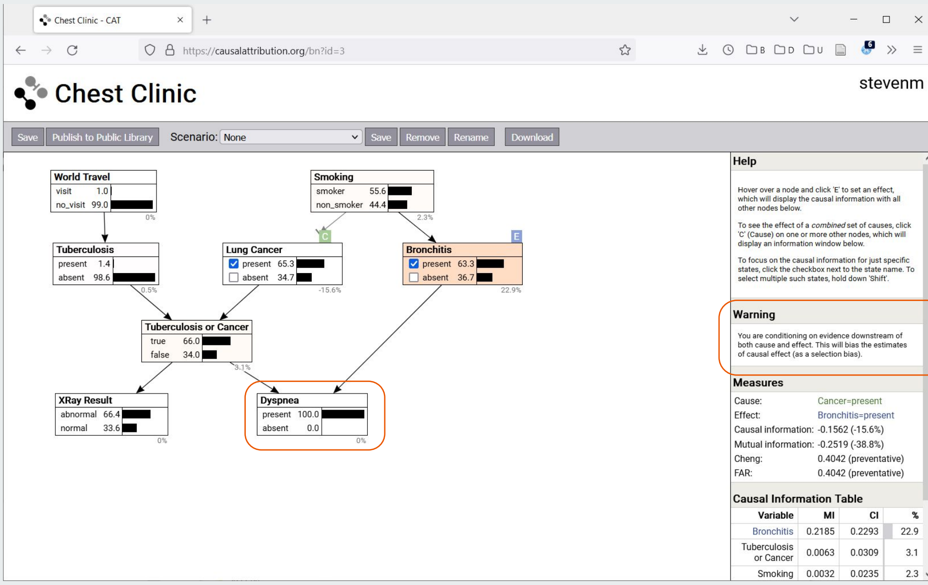

CAT measures
- Mutual Information
- Cheng's Causal Power
- FAR
- CI

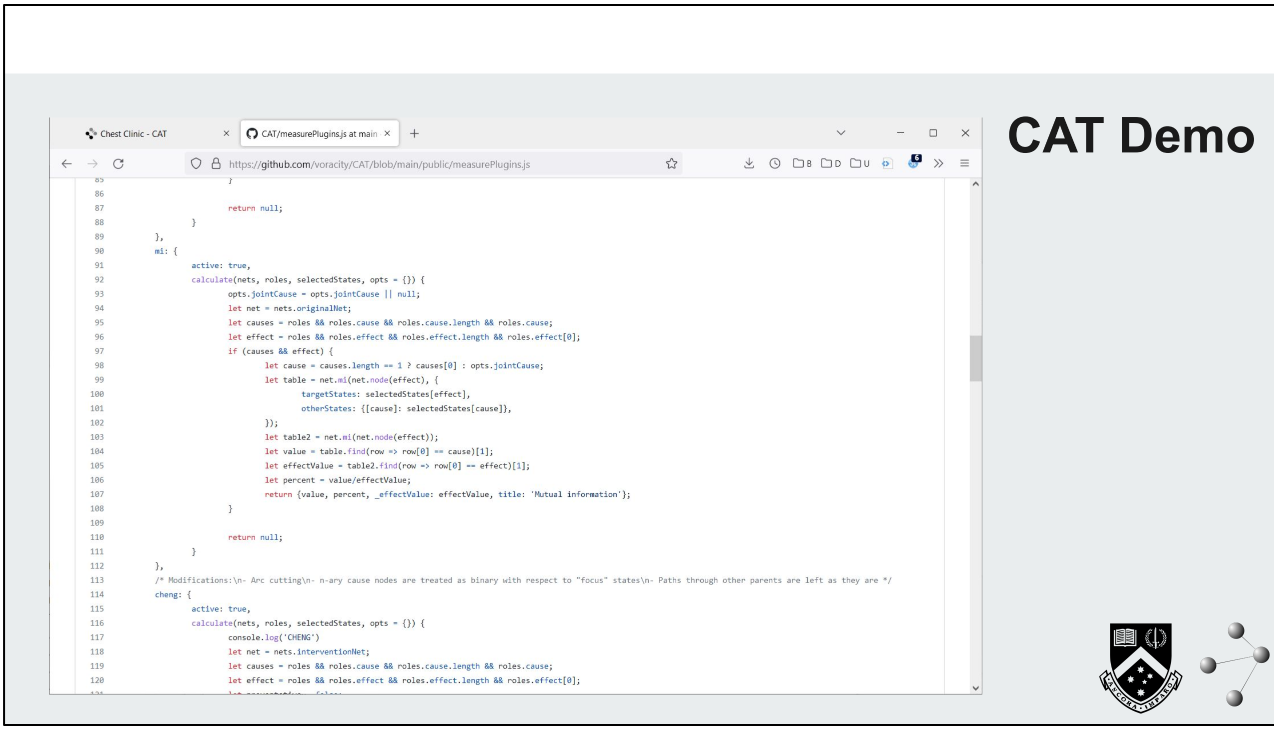

How to add new measures via GIT Hub
See https://github.com/voracity/CAT

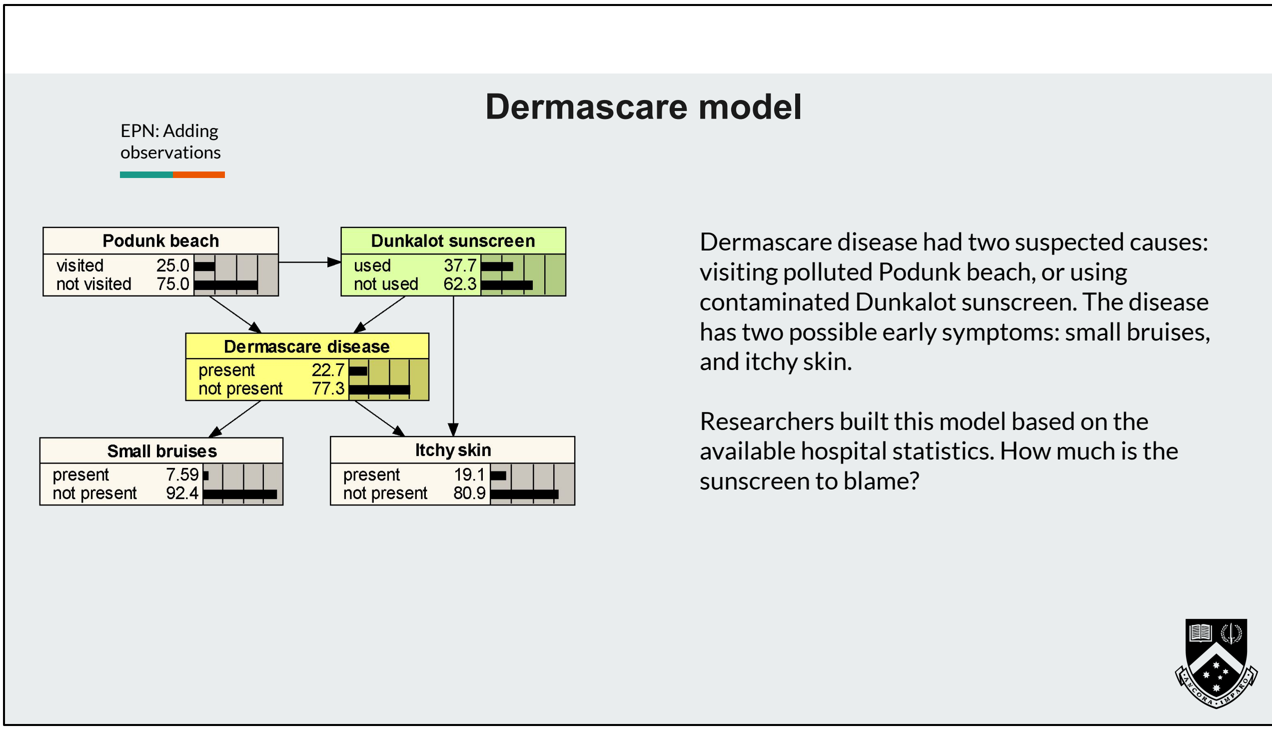

We've been talking about measuring the causal power of one variable over another by simulating an intervention on the putative cause. This is structurally distinct from just entering an observation about the cause. But now I want to make a few remarks about doing both at once: intervening on the cause where we have observed some other variables.

This fictitious example concerns a skin disease called Dermascare. Dermascare disease had two suspected causes: visiting polluted Podunk beach, or using contaminated Dunkalot sunscreen. The disease has two possible early symptoms: small bruises, and itchy skin. Researchers built this model based on the available hospital statistics. How much is the sunscreen to blame?

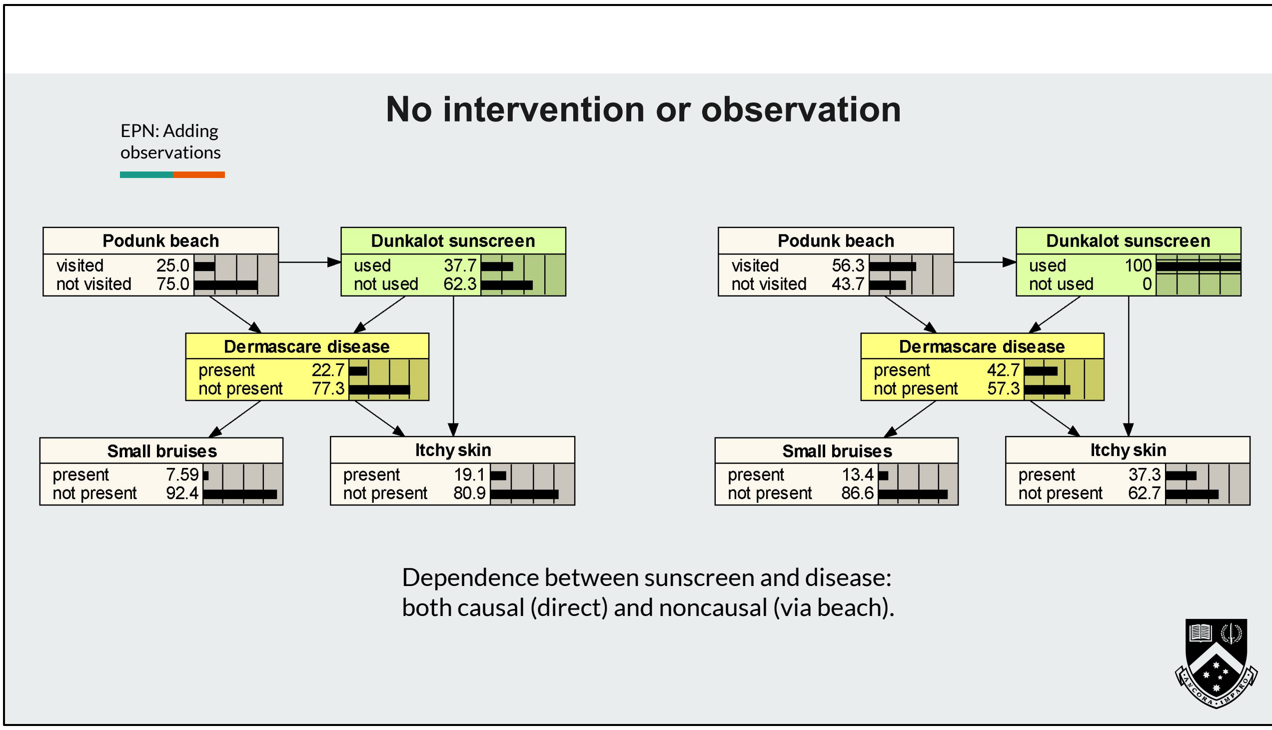

As we've already seen, the problem with just observing the people who used sunscreen is that the change in probability for them having the disease may be partly due to noncausal paths. Here, using sunscreen makes it more likely that the person visited the beach, which might be contributing to the dramatic increase in the probability of the disease.

Intervening on the cause of interest

EPN: Adding observations

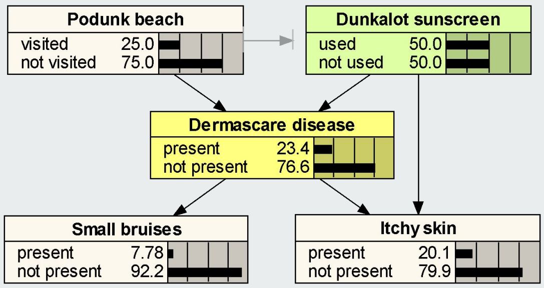

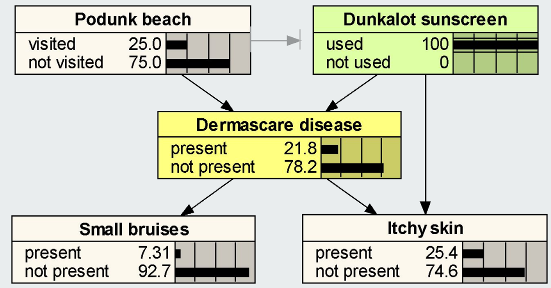

Dependence between sunscreen and disease:
only causal (direct).

When we use CAT to simulate an intervention on sunscreen, such as the randomisation shown here, we can break such noncausal paths, which guarantees that the remaining dependence is entirely due to the causal paths. Notice that using sunscreen has a slightly protective effect against the disease; the dramatic increase we saw previously was entirely due to visiting the beach.

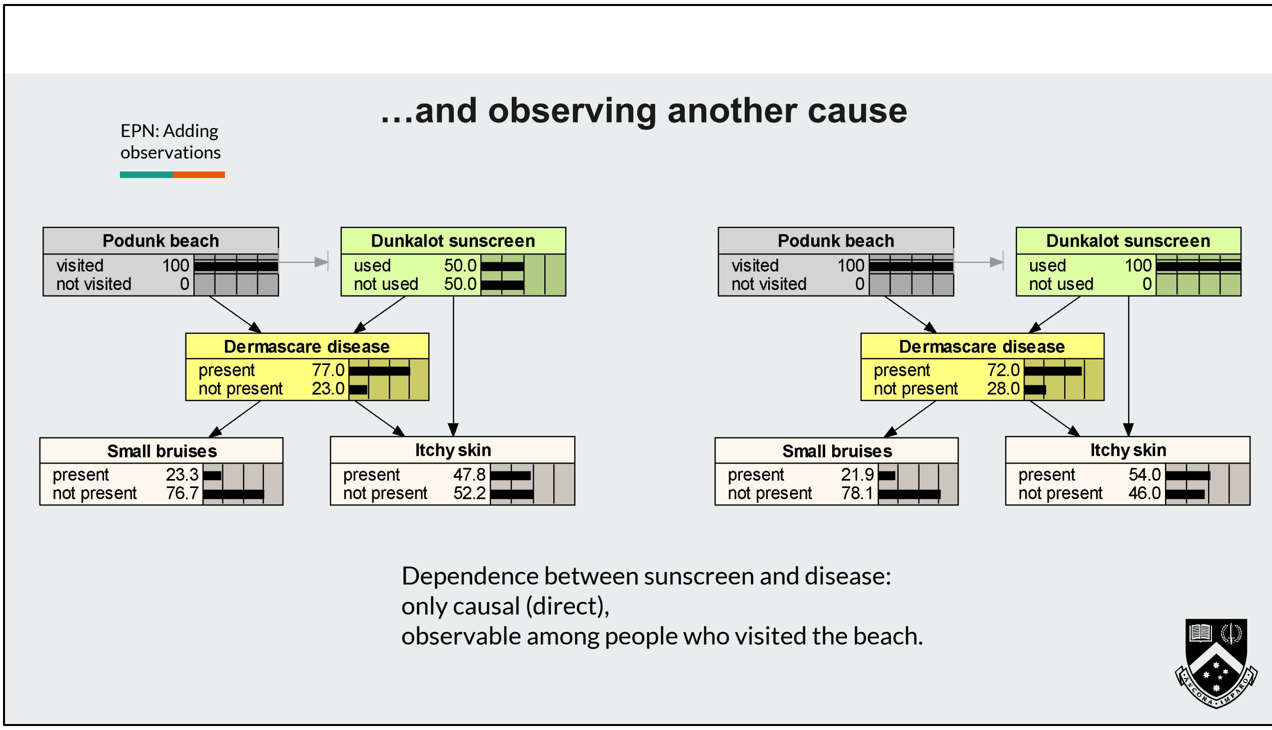

But what if we have also observed that someone has visited the beach? Now, the first mistake one might make in interpreting the result would be to think that the dependence here is the causal power of sunscreen in general. This is equivalent to selection bias: if we want to know the causal power for the population in general, but we have selected (usually inadvertently) a disproportionate number of people who went to the beach, then the results may differ and a naive extrapolation would be misleading.

However, this is a perfectly legitimate measurement provided that it is interpreted correctly: it's the causal power of sunscreen on the disease that is observable among people who visited the beach.

…and observing an effect

EPN: Adding observations

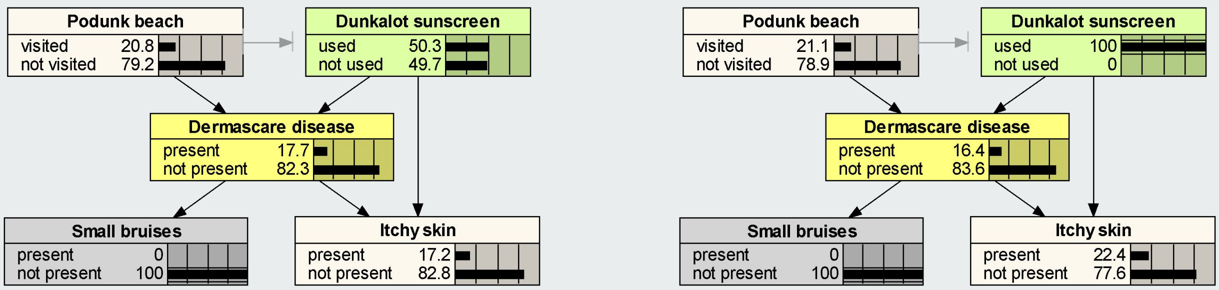

Dependence between sunscreen and disease:
only causal (direct),
observable among people who don't have bruises.

Similarly, we can measure the causal power given the observation that someone has no bruises. A more subtle mistake in interpretation is to say that this is the causal power sunscreen really has (or had) in this subset of the population. But you're not seeing what sunscreen might have done to these people if it resulted in them not having bruises. You're only seeing the causal influence that is observable given that they ended up without bruises.

…and observing a common effect

EPN: Adding observations

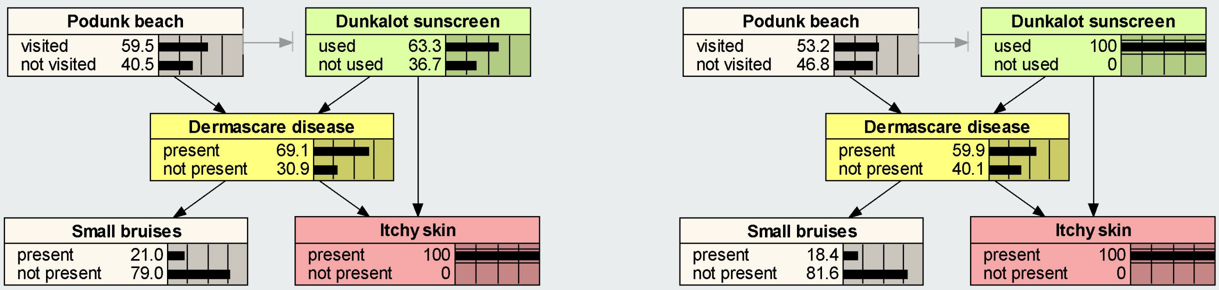

Dependence between sunscreen and disease:
both causal (direct) and noncausal (via itchy skin),
observable among people who have itchy skin.

But even if you interpret the measurement properly, there is a more serious quantitative problem if you observe a common effect, i.e., a descendant of both the cause and the effect variable of interest. Here, the observation actually creates a noncausal path that the intervention on sunscreen doesn't prevent. Specifically, either sunscreen or the disease could cause itchy skin. When we apply sunscreen, we have partly explained the itchy skin observation, so it becomes weaker evidence for the disease, and the probability of the disease drops. Consequently, any measurement of the dependence between sunscreen and disease will be partly due to the noncausal component.

We will add a verbal and visual warning to this tool (e.g., colouring the variable red) if you have added an observation that created a noncausal connection of this kind. This will help prevent users from misinterpreting the result, and might persuade them not to add the observation after all.

References

- Arora, Paul, et al. "Bayesian networks for risk prediction using real-world data: a tool for precision medicine." Value in Health 22.4 (2019): 439-445.
- Chickering, D. M. (1995, August). A transformational characterization of equivalent Bayesian network structures. In Proceedings of the Eleventh conference on Uncertainty in artificial intelligence (pp. 87-98).
- Fenton, Norman, Martin Neil, and Daniel Berger. "Bayes and the law." Annual Review of Statistics and Its Application 3 (2016): 51. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4934658/
- Glymour, Clark, and David Danks. "Reasons as causes in Bayesian epistemology." The Journal of philosophy 104.9 (2007): 464-474.
- Korb, Kevin B.; Nyberg, Erik P. & Hope, Lucas (2011). A new causal power theory. In P. M. Illari, F. Russo, & J. Williamson (Eds.), Causality in the Sciences, pp 628–652. Oxford: Oxford University Press.
- Marcot, Bruce G., and Trent D. Penman. "Advances in Bayesian network modelling: Integration of modelling technologies." Environmental modelling & software 111 (2019): 386-393.
- Nyberg, Erik P., et al. "BARD: A structured technique for group elicitation of Bayesian networks to support analytic reasoning." Risk Analysis 42.6 (2022): 1155-1178.