SAND2019-8533C

New Directions in Causality and Causal Modeling

International System Dynamics Conference

July 22, 2019

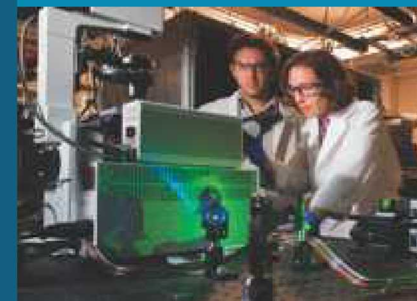

PRESENTED BY

Asmeret Naugle, Felix Wang, Lauren Hund, Fred Rothganger

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Why would we study causality?

The world is causal

Understanding causality can improve our ability to:

- Explain the world
- Predict
- Influence

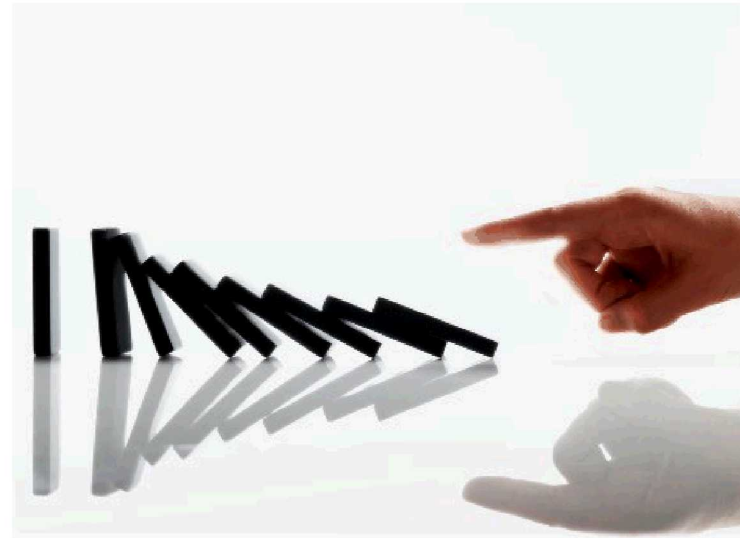

Dangers of untested causal assumptions

- Superstition and the illusion of controllability

Moving forward with causality in science

Theory-based modeling already incorporates causality

- But are the assumptions correct?

"Recent" academic interest in causality

- Symbolic and logical language to study causality
- Interest in incorporating into AI
  - Pearl: "strong AI" will necessitate causal understanding
    - Judea Pearl: One of the well-known causality researchers/writers, 2011 winner of ACM Turing award
    - "Strong AI": Hypothesis that machines could act intelligently by actually thinking (as opposed to simulated thinking)
      - Artificial Intelligence, A Modern Approach. 2nd Edition. Russell & Norvig, 2003. page 947ff

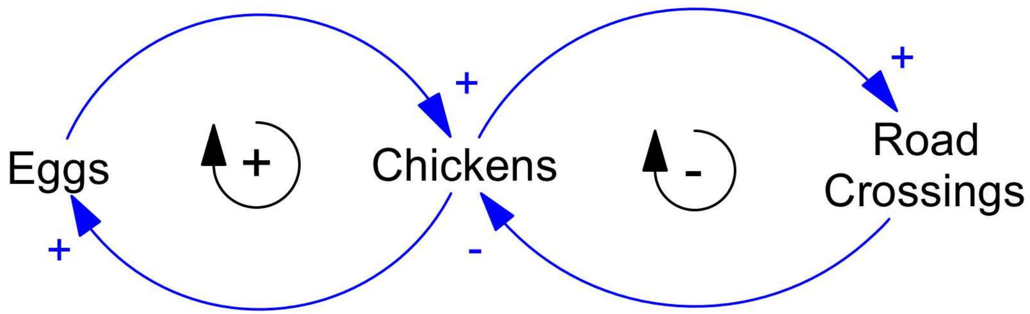

System Dynamics and Causality

What does causality mean for System Dynamics modelers?

- Assumptions, conceptual models
- Concepts of causality are different!

Why does it matter what else is going on in causality and causal modeling?

- Put system dynamics in context
- Improve/combine methods
- New opportunities

A few intuitive definitions (from Wikipedia)

Causality: How one process or state (the *cause*) relates to another (the *effect*), where the first is partly responsible for the second, and the second is partly dependent on the first

- Examples:
  - If it rains, then the ground gets wet.
  - If you study, then you will get a better grade on the test

Counterfactual: A conditional containing an if-clause which is contrary to fact

- Examples:
  - If it had rained yesterday, then the party would have been cancelled.
  - If you had not taken a wrong turn, then you would have gotten there on time.

Confounder: A variable that influences both the dependent variable and independent variable causing a spurious association (correlation, not causation).

- Examples:
  - Weather affects murder rates and ice cream sales
  - Economic development affects piracy and climate change

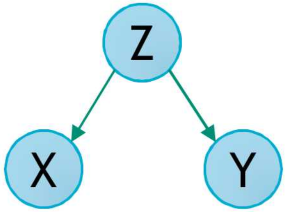

Plan for today

1. Felix Wang:
   - How can causality enable transportability and inference between application areas?
2. Lauren Hund:
   - How can causal modeling make sense of imperfect data?
3. Ryan Dellana:
   - How can machine learning best incorporate causality?

Goals

Help you understand what's going on in causality research to:

Put system dynamics in context

Improve/combine methods

New opportunities

Inference Using Causal Models

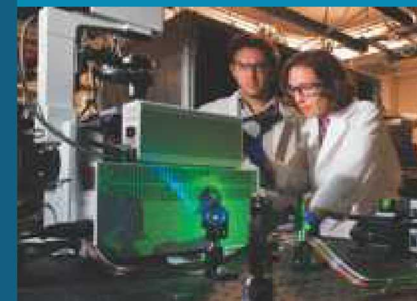

PRESENTED BY

Felix Wang

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2019-2336 PE

Causal Models as Data-Generating Processes

Causal models can be thought of as an "inference engine" taking inputs

- *Causal assumptions* and a structural model to encode these assumptions
- *Queries* about variables of interest and their relationships
- *Data* from observations (seeing) and interventions (doing)

… and producing outputs

- *Logical implications* of our assumptions
- *Statistical implications* of our assumptions
- *Claims* about the answers to our queries

Some examples of causal queries

- Was it move/action X that led to winning/losing the game?
- Will a proposed policy on X have the intended results Y?
- How will the release of product X impact the sales of product Y?
- Does the evidence prove, beyond a shadow of a doubt, the crime Y was committed by X… Or was it an accident that would have happened anyway, because of Z?

Encoding the Structure of Causal Models

Structural Equations Models (SEMs)

- Variables of interest and their background factors
- Probability distributions over the background factors
- Functions that map from variables (and background factors) to variables
- Visually represented as a Directed Acyclic Graph (DAG)

$$M: \begin{array}{l} z \Leftarrow f_Z(u_Z) \\ x \Leftarrow f_X(z, u_X) \\ y \Leftarrow f_Y(x, u_Y) \end{array}$$

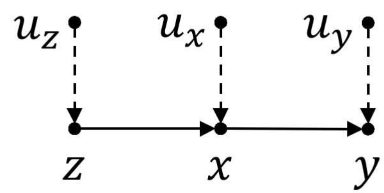

Changes to the model can be reflected as changes in the graph

- Operations (e.g. interventions $do(x)$) add/remove edges

$$M_x: \begin{array}{l} z \Leftarrow f_Z(u_Z) \\ x \Leftarrow x_0 \\ y \Leftarrow f_Y(x, u_Y) \end{array}$$

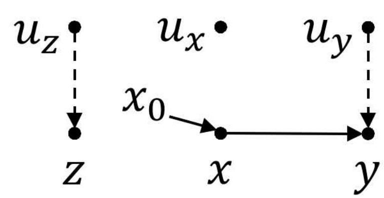

$$P_M(y|do(x)) = P_{M_x}(y)$$

Inference using Causal Models

Many tools available for working with causal models

- D-separation, belief propagation, adjustment criteria and formulae, etc.
- Graphical representation supplements the discovery of solutions

Variables: $x =$ smoking, $w =$ tar, $y =$ cancer, $z =$ genes (unobserved)

M:

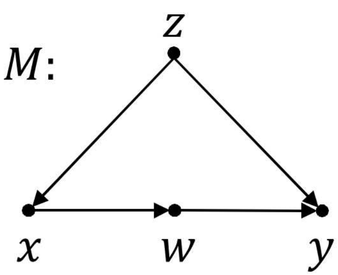

$$
\begin{aligned}
z &\Leftarrow f_Z(u_Z)\\
x &\Leftarrow f_X(z, u_X)\\
w &\Leftarrow f_W(x, u_W)\\
y &\Leftarrow f_Y(w, z, u_Y)
\end{aligned}
$$

Query: $P(y|do(x))$

$= \Sigma_w P(y|do(x), w) P(w|do(x))$ Probability Axioms

$= \Sigma_w P(y|do(x), do(w)) P(w|do(x))$ 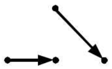 Observation to Intervention

$= \Sigma_w P(y|do(x), do(w)) P(w|x)$ 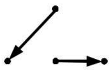 Intervention to Observation

$= \Sigma_w P(y|do(w)) P(w|x)$ 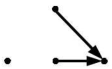 Delete Intervention

$= \Sigma_{x'} \Sigma_w P(y|do(w), x') P(x'|do(w)) P(w|x)$ Probability Axioms

$= \Sigma_{x'} \Sigma_w P(y|w, x') P(x'|do(w)) P(w|x)$ 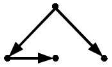 Intervention to Observation

$= \Sigma_{x'} \Sigma_w P(y|w, x') P(x') P(w|x)$: Estimand 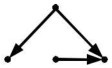 Delete Intervention

Selecting Between Causal Models

Selection diagrams identify where two domains of interest may differ

- E.g.: distribution of background factors, mapping functions, causal structure
- Represented graphically through selection variables (adding an edge $s_i \rightarrow v_i$)

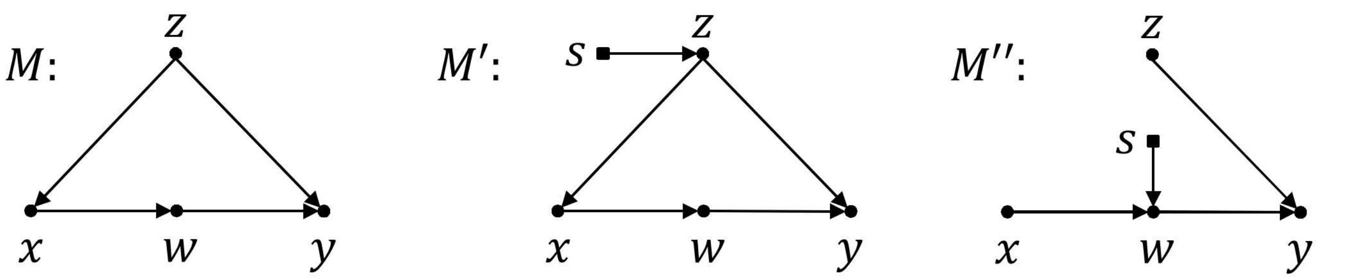

$P(z) \neq P'(z)$ $\quad$ $w \Leftarrow f(x) \neq f''(x)$

Example: Given studies in M', M'' how effective will marketing be in population M?

- Variables: x =advertisements, y = purchase, z = age, w = click-through rate
- Study in M' differs from target population in age
- Study in M'' was done as a randomized trial, and on a platform with high click-through rates

Transportability of Causal Models

Transportability allows us to make inferences between domains of interest

- Determine if generalizations are valid or invalid
- "License assumptions" about the transfer of causal relationships
- Different assumptions will yield different routes for transporting information
- Synthesize observational and interventional information from multiple domains

Some more examples highlighting transportability

- How do I navigate/find where I need to go when traveling in a new city?
- Why don't my physics simulations N match the new experimental measurements M?
- Should I have expected them to match in the first place?
- It's difficult to measure Y because of cost, what can I use as a proxy?
- Will a government program X in country N be effective if applied to country M?
- What experiments should be conducted to fill in the gaps of understanding?

Applications to Machine Learning and Artificial Intelligence

Dealing with data in a smarter way for machine learning

- How the data gets generated is important
- Model-based ML produces specific solutions to specific problems
- Combining human expertise in generating assumptions with ML methods for data analysis

Leveraging causal structure in AI agents

- Transfer learning during domain shift
- Sample efficiency and learning from counterfactuals
- Giving AI "free will"

References

- Pearl, Judea and Bareinboim, Elias. "External Validity: From Do-Calculus to Transportability Across Populations" *Statistical Science*, 2014, Vol. 20. No. 4, 579-595
- Pearl, Judea. *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge Univ. Press, 2009
- Pearl, Judea and Mackenzie, Dana. *The Book of Why: The New Science of Cause and Effect*, Basic Books, New York, 2018

Using causal models to analyze imperfect data

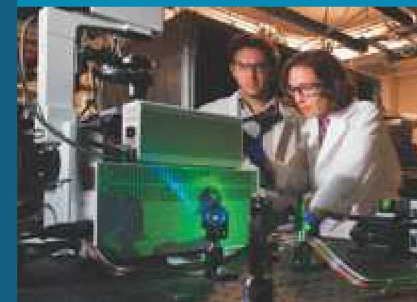

PRESENTED BY

Lauren Hund

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

How does causal modeling work?

Recall that causal queries (counterfactuals) are outcomes from hypothetical interventions on a system.

In an ideal world:

- Conduct the intervention on the system and see what happens – randomization (design of experiments, clinical trials)

In a data-limited world:

- Analyze observational data (you don't control the design – what you see is what you get)
- Common approach to causal modeling: fake randomization.

Ideal

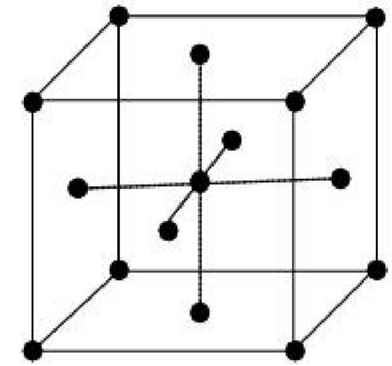

Actual

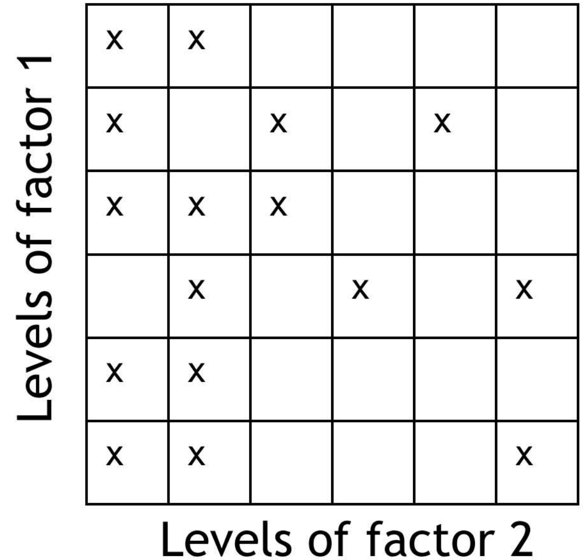

Sandia is all about counterfactuals!

Sandia problems commonly concern "extrapolative prediction."

- Generating a predictive distribution for an unobserved outcome – prediction of a counterfactual.

Some examples:

- Nuclear weapons are the ultimate counterfactual predictions - without full-system tests, we certify weapons.
- Weapon components: How will a component perform across a variety of conditions (temperature, environments, age?)
- Computer models: Run model and predict to setting without data (counterfactual).

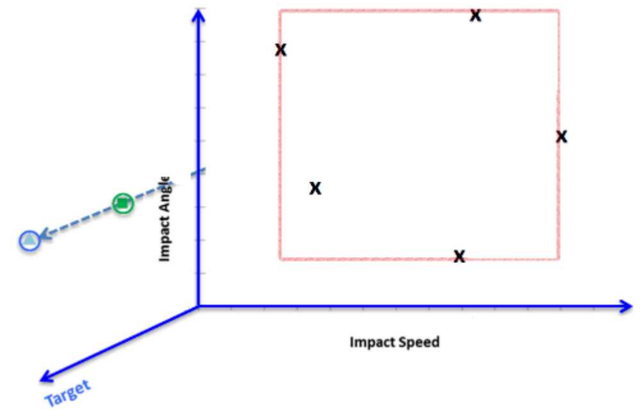

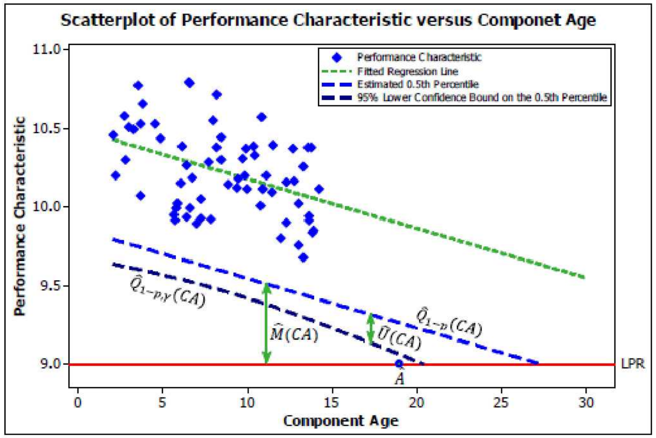

Dealing with data insufficiencies

Causality gives a language to talk about credibility of a prediction given less-than-ideal data.

- Much of causal inference is simply ensuring that your data analysis methods accurately reflect the "**data generating mechanism**," i.e. how your data were generated.
- Under what set of assumptions is my counterfactual prediction valid?

Structural causal modeling is one language of causality.

Steps for causal analysis:

1. Define a causal query.
   - Often a function of a counterfactual.
2. Determine how the collected data relates to the true underlying structural causal model.
   - Make a DAG!
3. Check if sufficient data to estimate query.
4. **Estimate the query from the data.**

Step 1: Causal query

$$P\big(Y = do(X = x)\big)$$

Step2: DAG

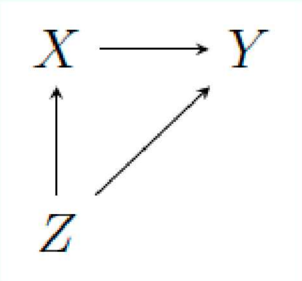

Step 3: Check criteria

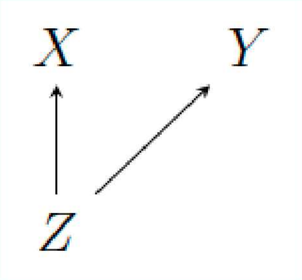

Estimate QoI

In practice, we want to move from ***qualitative*** DAG model to ***quantitative*** statistical model in order to estimate a causal query.

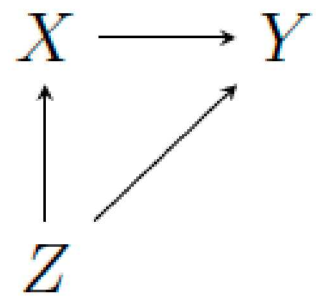

Adjustment formula:

$$P\big(Y = do(X = x)\big) = \sum_z P(Y|X = x, Z = z)P(Z = z)$$

Unobserved counterfactual

Observed in data

Stratifying on Z, we can estimate the counterfactual of interest from the data.

- Other formulations of the adjustment formula exist, e.g. for selection variables and for the front-door criterion.

Estimate causal query

Causal inference is all about models and assumptions.

- What assumptions are you willing to make?
- Do you have enough data to fit a "good" statistical model under those assumptions?

Fundamental assumptions of causality: Given a random sample from a population:

- ***Exchangeability:*** no unmeasured confounding
  - Measure enough variables?

- ***Positivity:*** enough data to estimate $P(Y|X = x, Z = z)$.
  - Have enough data?

- ***Consistency:*** no multiple versions of treatment
  - Treatment can be hypothetically manipulated in a consistent manner
  - Example: drug; counterexample: BMI.

Role of modeling and assumptions

Curse of dimensionality – use statistical models to approximate distribution of Y|X,Z.

"Art" of statistical modeling.

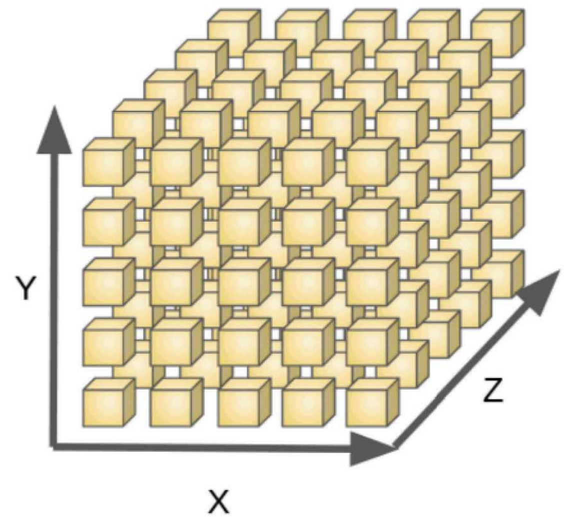

Picture taken from: https://medium.freecodecamp.org/the-curse-of-dimensionality-how-we-can-save-big-data-from-itself-d9fa0f872335

There is an implicit fourth assumption needed for causal estimation: correct model specification.

- Causal methods are often 'model-agnostic': how you model is separate from how you calculate causal estimands given the model.
- The 'modeling' stage is where good statistical and ML models come into play.

Thank you!

"Use of technical causal language, a good use, in our estimation, must be recognized as simply a shorthand for better versus worse analyses, as judged by the author, and not a metaphysical statement about causation per se..."

Lipton and Odegaard (2005)

Generative Models in Machine Learning

PRESENTED BY

Fred Rothganger

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Definition of "Generative Model"

As a probabilistic expression

X = State variable
Y = Observable variable

Discriminative model:
$P(X|Y)$

Generative model:
$P(X,Y)$
or $P(Y|X)$, where $P(X)$ is known or estimated

Definition of "Generative Model"

As a procedure

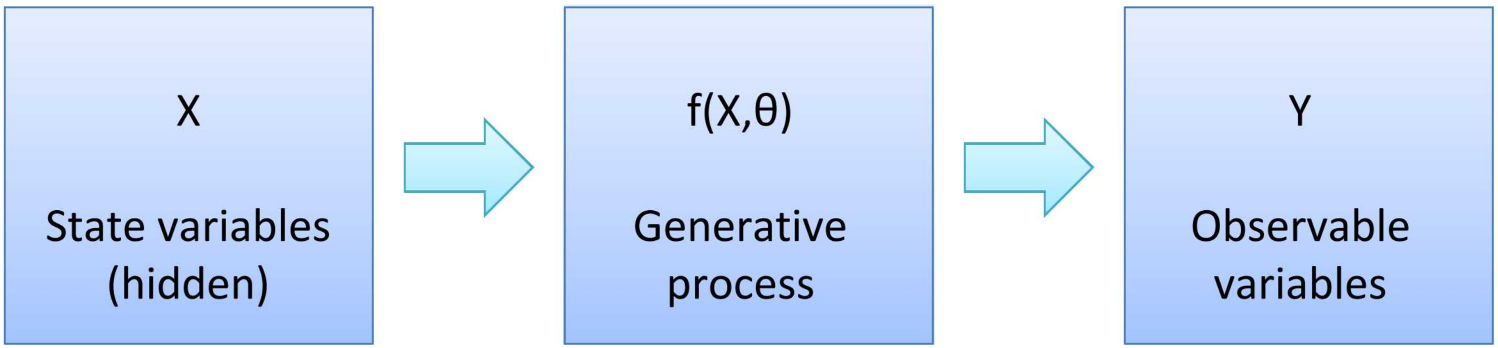

Assumes a chain of causality from X, through mechanisms f(θ), to Y.

Example: 3D graphics (such as video games or special effects in movies)

- X = Position/orientation of various objects in scene
- θ = Shape/surface material of objects
- f = Projection process. In the real world, it is reflected light reaching your eye.
- Ray-tracing and other 3D rendering methods simulate this.
- Y = Resulting pixels on the movie screen.

Example: Rao & Ballard

Level 1

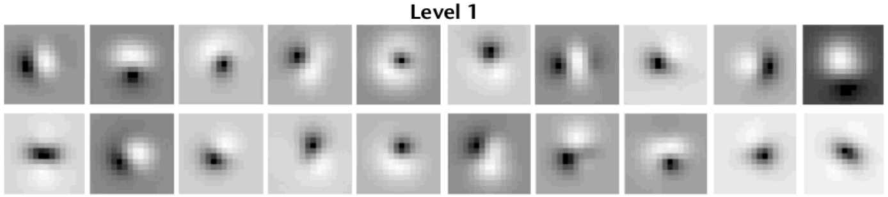

Level 2

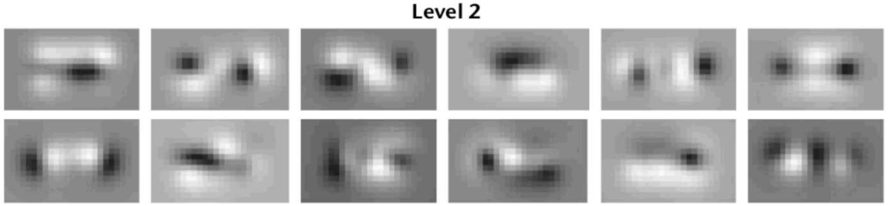

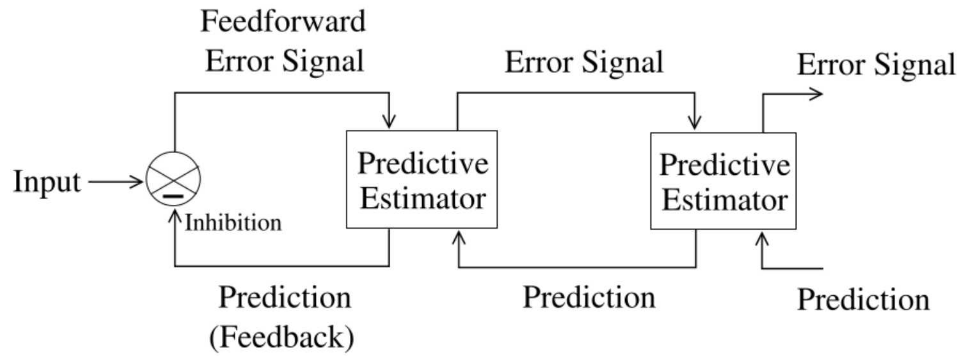

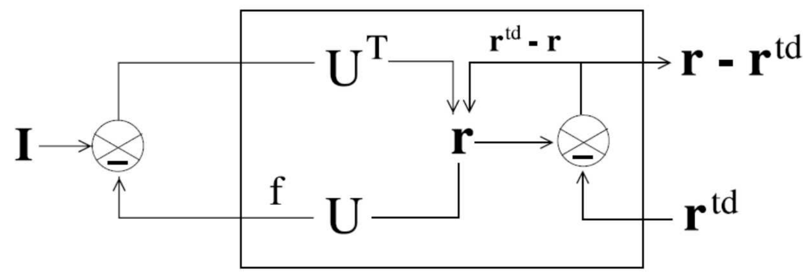

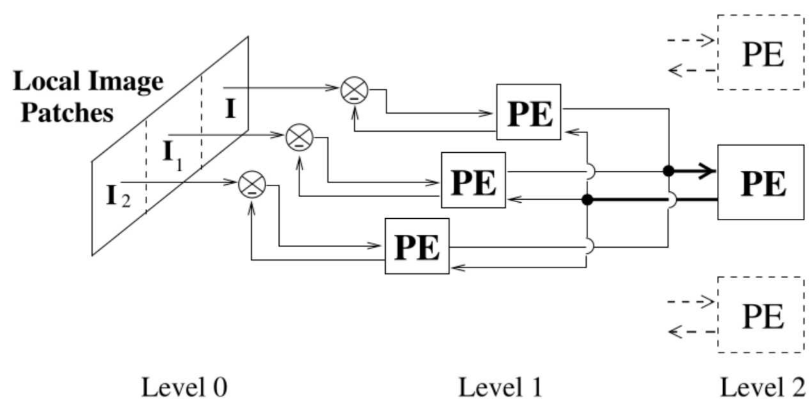

Example: PredNet

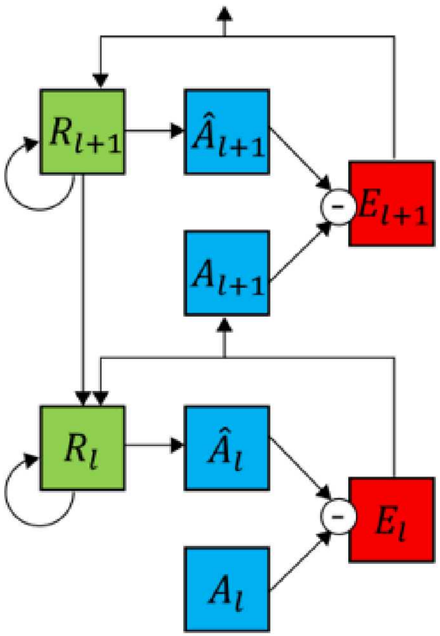

[https://coxlab.github.io/prednet]

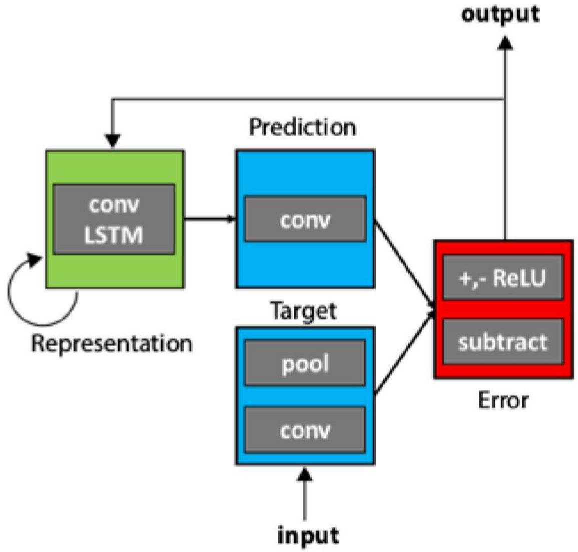

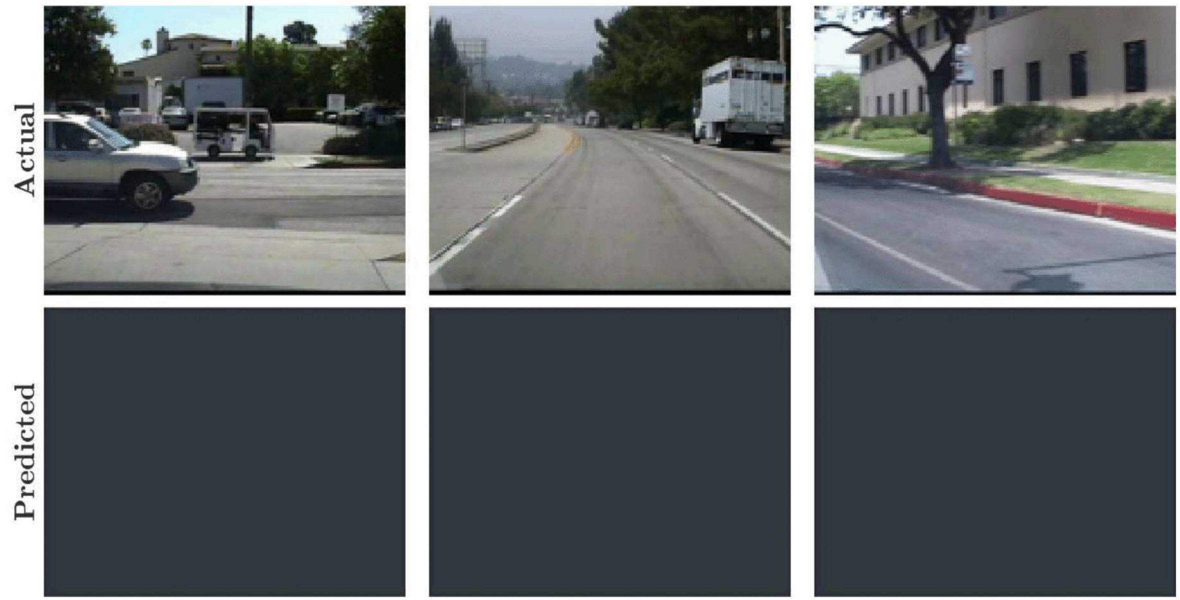

Example: Generative Adversarial Network (GAN)

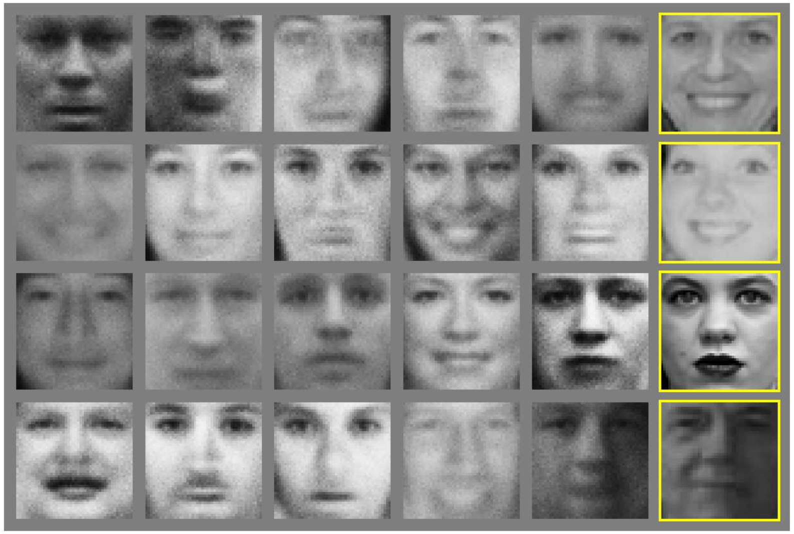

[Goodfellow 2014]

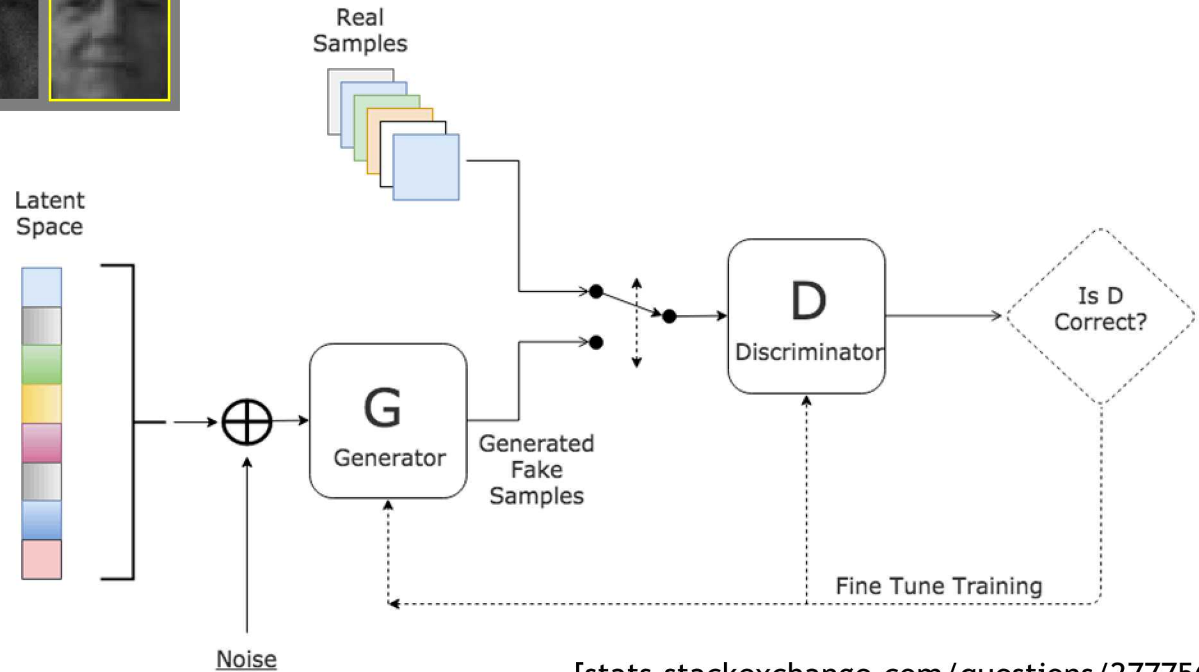

[stats.stackexchange.com/questions/277756]

How to learn model structure

- The previous examples use neural networks and associated learning methods. Despite the generality, these have an implicit, hand-crafted structure.
- An alternative is to represent the generative process in a language suitable for genetic algorithm (GA) style random search.

Thank You!

Goals

Help you understand what's going on in causality research to:

Put system dynamics in context

Improve/combine methods

New opportunities