Information-based Causal Modeling (IBCM) and Applications

Tony Cox

tcoxdenver@aol.com

April 10, 2019

Goals

- Provide Hill-like criteria for assessing consistency of data with the hypothesis of a *manipulative* causal exposure-response, using information rather than association
  - How much will *changing* exposure X *change* subsequent probability distribution of effect Y, given levels of other direct causes of Y (e.g., age, sex, smoking, income)?
    - *Not* association-based (Hill/IARC/regression), attributive (burden of disease), predictive (Granger), counterfactual/ potential outcomes (propensity score, MSM), structural (Simon), mechanistic, or but-for (legal) causation
  - Needed for decisions and policy recommendations
  - Information in bits = reduction in conditional entropy of Y | X
- Minimize or eliminate untested assumptions
  - Use causal graphs to articulate and test causal hypotheses
    - Direct effects, total effects, CPTs, pathways of mechanisms
  - Test invariant causal predictions (ICP) vs. make counterfactual assumptions

Example of an information-based DAG: What exactly does it mean?

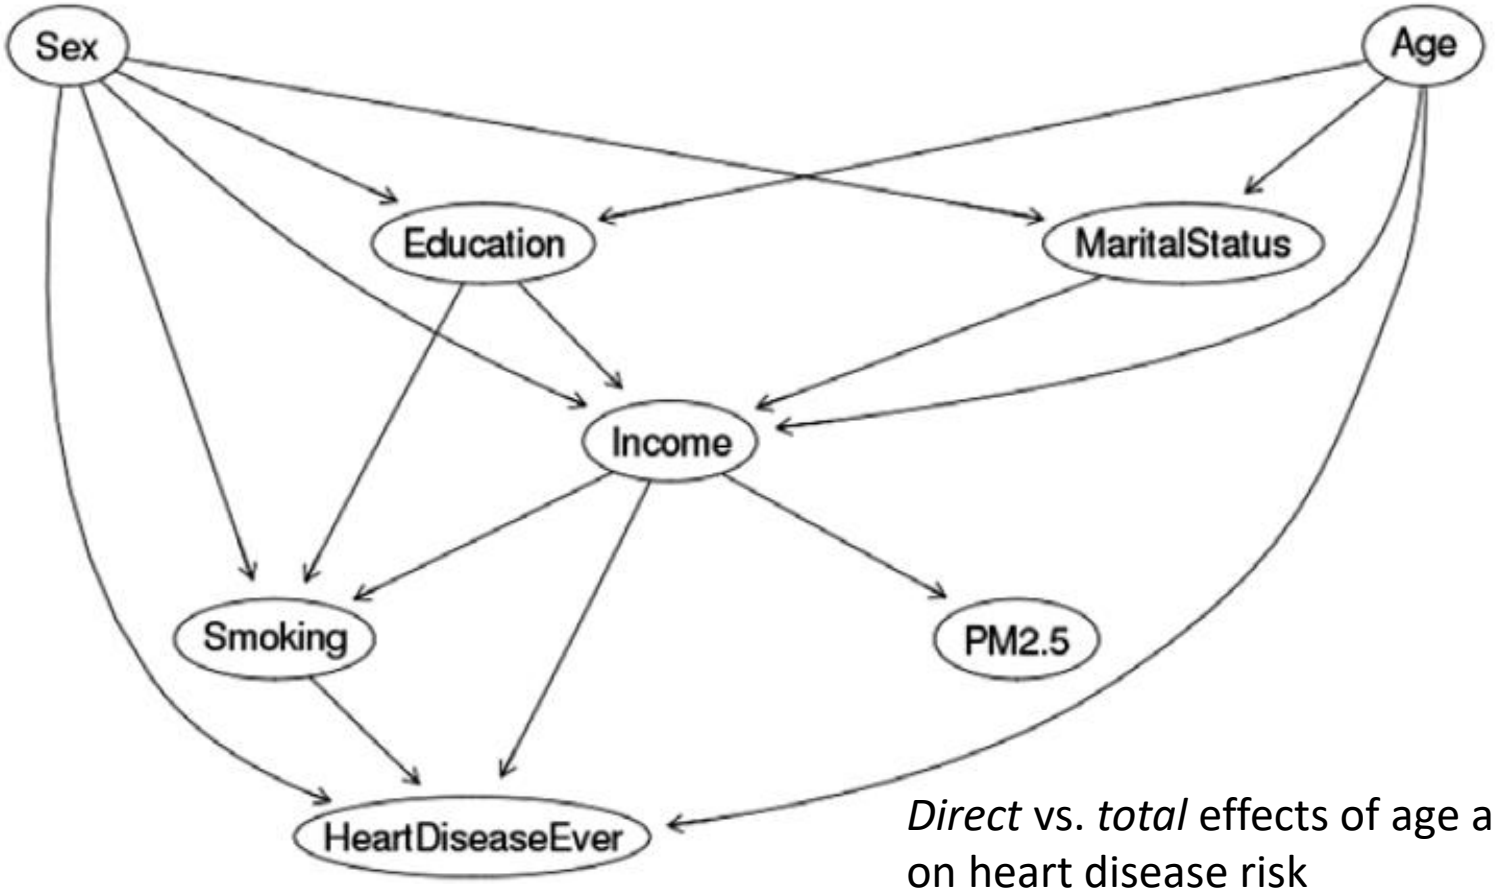

Direct vs. *total* effects of age and income on heart disease risk

Arrow = not conditionally independent of

Ambiguity of counterfactual income levels

Cox LA. Jr. 2018. Socioeconomic and particulate air pollution correlates of heart disease risk. *Environ Res*. 167:386–392. CDC BRFSS data set

Many packages provide algorithms and principles to identify (causal) DAGs from data

- Conditional independence (constraint-based algorithms): Causes are informative about their effects
  - *dagitty*, *bnlearn* packages; *CompareCausalNetworks* package
- Likelihood principle (score-based algorithms): Valid causal models explain the data (i.e., make it not too unlikely).
  - Choose DAG model to maximize likelihood of data (*bnlearn* package)
- Composition principle: If $X \rightarrow Z \rightarrow Y$, then $dy/dx = (dz/dx)*(dy/dz)$
  - Path analysis, *lavaan* package
- Granger principle: Predictively useful information flows from causes to their effects over time (*granger.test*, *bnstruct*)
- Model error specification principle: Causes reveal simplicity
  - effect = f(cause) + error; *LiNGAM* packages
- Invariance of causal CPTs (*InvariantCausalPrediction* package): Completely decribed causal relationships are universal
  - Peter et al. (2015) http://stat.ethz.ch/~nicolai/invariant.pdf

Proposed criteria for consistency with manipulative causal exposure-response

1. *Mutual information:* Causes are informative about their direct effects.
2. *Directed dependence:* Information flows from causes to effects over time
3. *Internal consistency:* Estimates of the same effect using different adjustment sets (e.g., common causes) are not significantly different
4. *External consistency:* Invariant causal prediction (ICP): Response conditional probability table (CPT) does not differ across studies

Proposed criteria for manipulative causal exposure-response

5. *Coherence*: Path from exposure to response through causal biological network
6. *Causal mediation confirmation*: Changes in exposure explain quantitative changes in mediating variables and resulting response(s)
   - Chain of accountability (HEI): $X \rightarrow Z \rightarrow Y$
7. *Refutation:* Data reject alternative (non-causal) explanations

Directed acyclic graph (DAG) model

- What does it mean?
  - Arrows into *Smoking*?
  - Non-causal (BN): factors joint pdf
  - Causal: CPTs also represent invariant causal mechanisms
- When are arrows "causal"?
  - Directed information flow
  - Homogeneous CPTs (latent vars.)
- How to learn DAG from data?
  - Test conditional independence implications, CPT invariance and homogeneity, compositionality
- Use in risk assessment?
  - Partial dependence plots for direct and total causal effects of interest
- How trustworthy are the results?
  - Non-parametric model ensembles
  - Robustness to different algorithms and principles

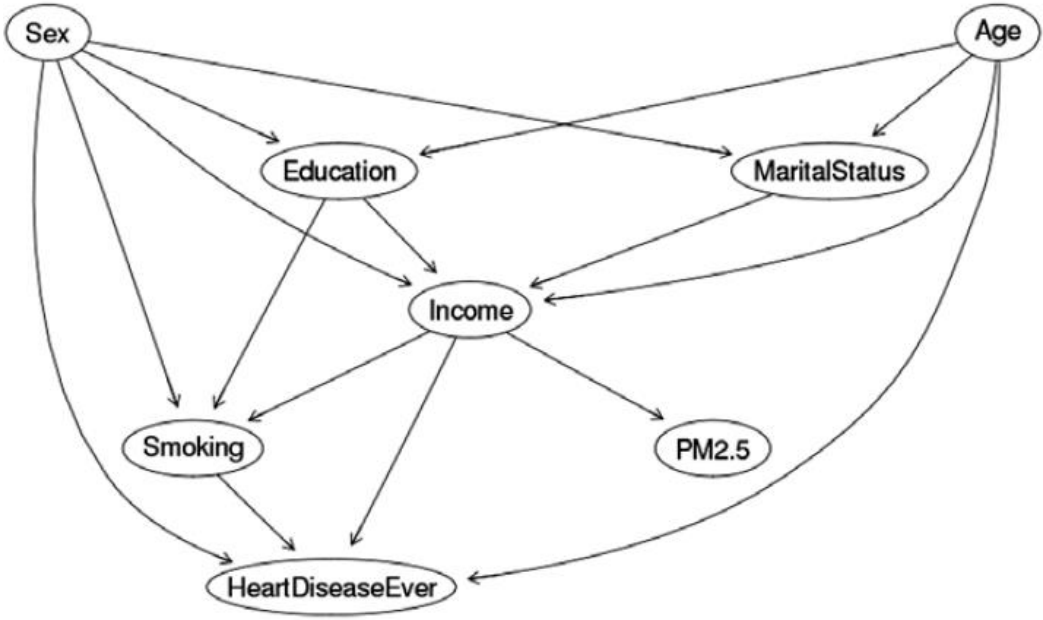

Arrows represent dependence (mutual information) between variables

Not necessarily manipulative causation

Some causal links lack clear directions (*MaritalStatus* ↔ *Income*)

Estimating a response CPT

- What does it mean?
  - Arrows into *Smoking*?
  - Non-causal (BN): factors joint pdf
  - **Causal: CPTs also represent invariant causal mechanisms**
- When are arrows "causal"?
  - Directed information flow
  - Homogeneous CPTs (latent vars.)
- How to learn DAG from data?
  - Test conditional independence implications, CPT invariance and homogeneity, compositionality
- Use in risk assessment?
  - Partial dependence plots for direct and total causal effects of interest
- How trustworthy are the results?
  - Non-parametric model ensembles
  - Robustness to different algorithms and principles

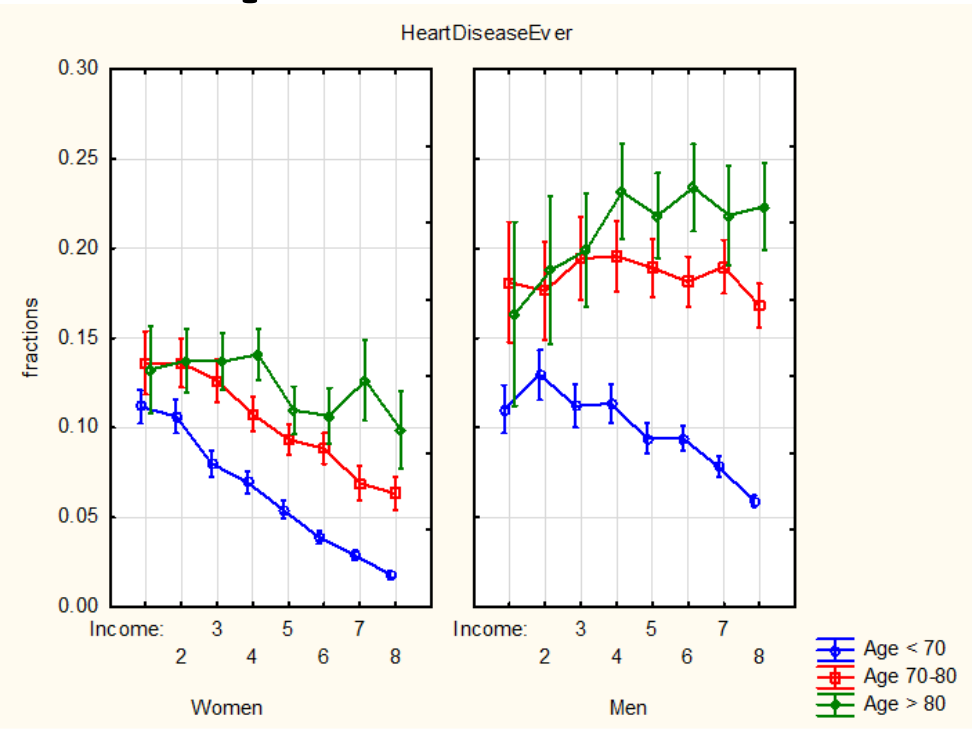

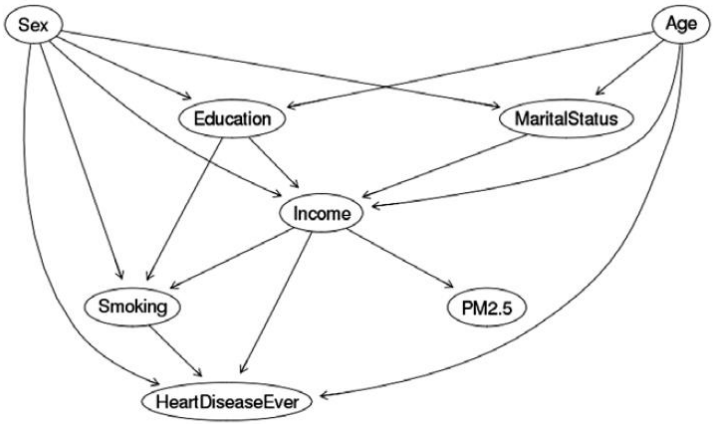

Model = dependencies + causal CPTs

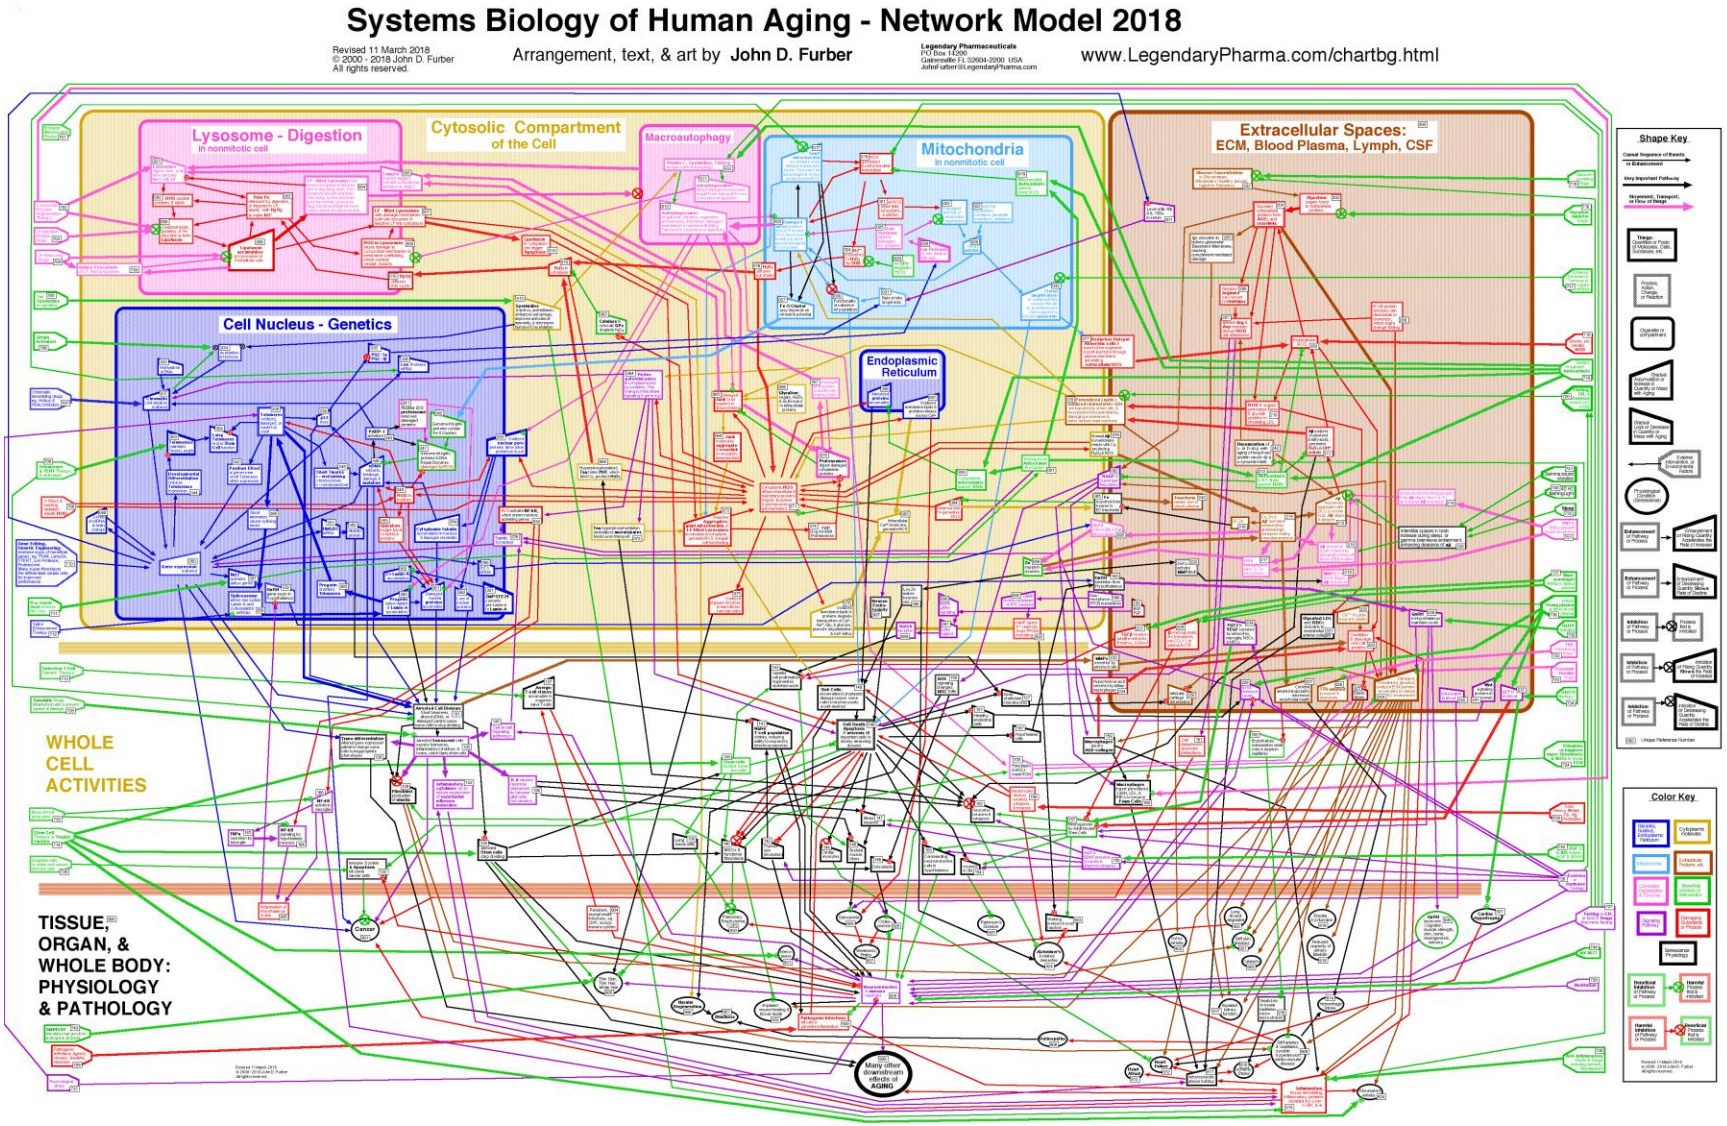

Partial dependence plot for natural direct causal effect in DAG model

Risk of heart disease

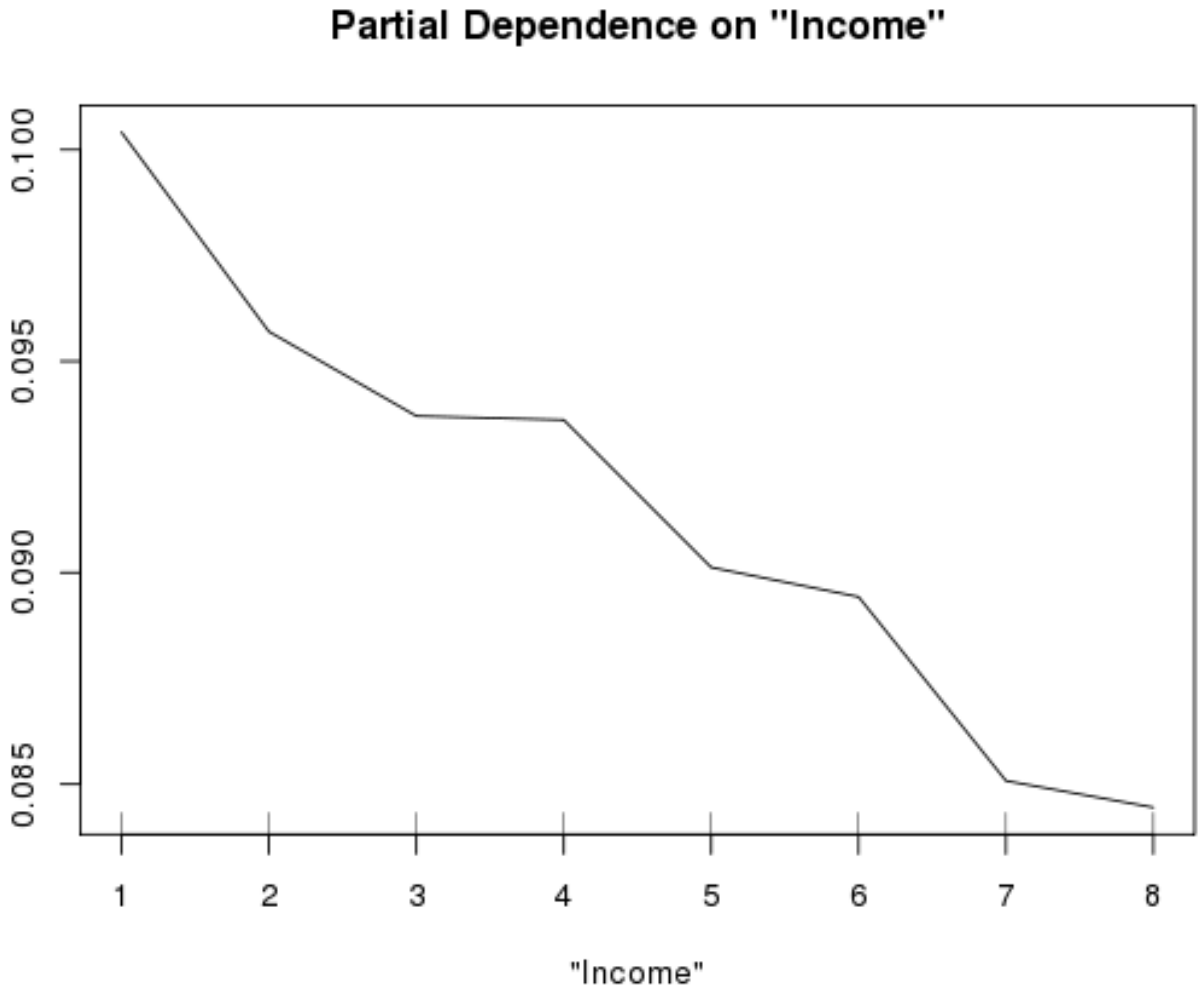

Updating Hill for manipulative causation: Mutual information

Bradford-Hill considerations	Modern causal discovery and inference principles
Strength of association: Stronger associations are more likely to be causal	• *Mutual information principle:* Causes are informative about their direct effects and help to predict their values. • *Conditional independence version:* Effects are not conditionally independent of their direct causes. • Direct causes contain at least as much information about their effects as do more remote indirect causes

Hill's strength and consistency considerations often fail (mislead) in practice (Ioannidis, 2016; Pearl and Mackenzie, 2018). Strong association usually indicates strong biases or confounding. Consistency of effects estimates (e.g., regression coefficients) in different populations may indicate common omitted confounders, p-hacking.

Mutual information (arrow in DAG) provides a useful non-parametric alternative

Mutual information may be positive even if correlation is zero ($y = x^2$) or zero even if correlation is positive (spurious regression)

Direct causes are adjacent to their effects in valid causal graphs

Bradford-Hill considerations	Modern causal discovery and inference principles
Strength of association: Stronger associations are more likely to be causal	• *Mutual information principle:* Causes are informative about their direct effects and help to predict their values. • *Conditional independence version:* Effects are not conditionally independent of their direct causes. • Direct causes contain at least as much information about their effects as do more remote indirect causes

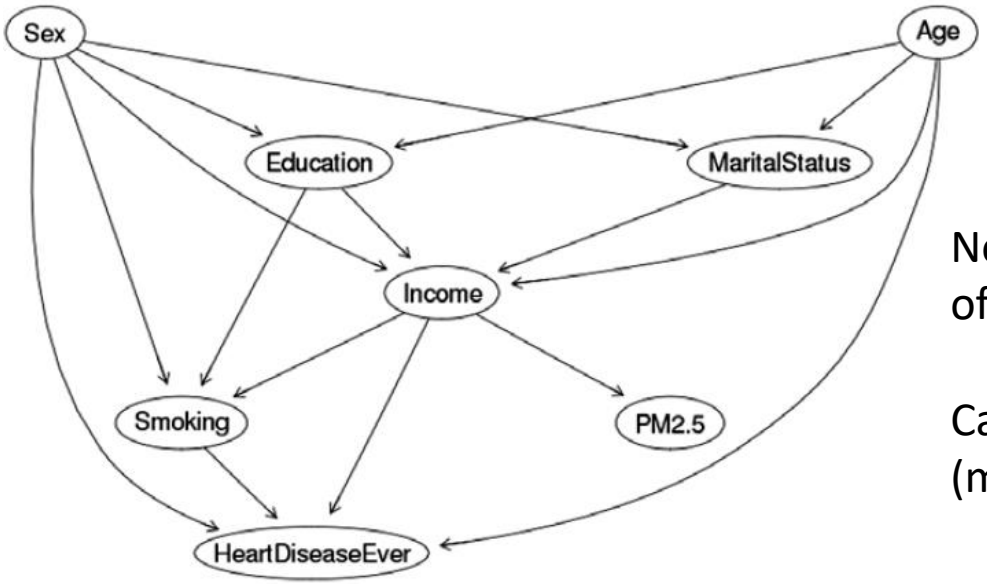

No evidence in DAG that PM2.5 is a direct cause of increased heart disease risk (adjacency)

Can bound maximum size of undetected effect (missing arrow)

Updating Hill for manipulative causation: Consistency

Bradford-Hill considerations	Modern causal discovery and inference principles
Strength of association: Stronger associations are more likely to be causal	• *Mutual information principle:* Causes are informative about their direct effects and help to predict their values. • *Conditional independence version:* Effects are not conditionally independent of their direct causes. • Direct causes contain at least as much information about their effects as do more remote indirect causes
Consistency of findings across populations, study designs, times, locations, investigators, etc.	• *Internal consistency:* Similar effects are estimated using different adjustment sets • *External consistency:* Causal laws, expressed as conditional probability tables (CPTs), are invariant, homogeneous, and transportable across study settings

External consistency: *Invariant causal prediction* (ICP) property of CPTs across studies provides a testable foundation for making unambiguous counterfactual predictions.

Homogeneity of CPTs within and across studies provides a testable basis for managing latent variables. (If homogeneity fails, use finite mixture distributions, HMMs, etc.)

Invariant causal prediction (ICP)

- Testable property of response CPTs
  - Study ID does not help to predict response
- Uses of ICP:
  - *Generalize* from individual study results
    - Invariant causal CPTs can be “transported” across study settings, allowing predictions in novel settings
  - *Synthesize* studies with overlapping variables
  - *Detect* omitted/unobserved causes (latent variables) via homogeneity tests of response CPTs
- Provides sound basis for counterfactual causality

ICP addresses challenge of ambiguous (model-dependent) counterfactual predictions

Scientific Method

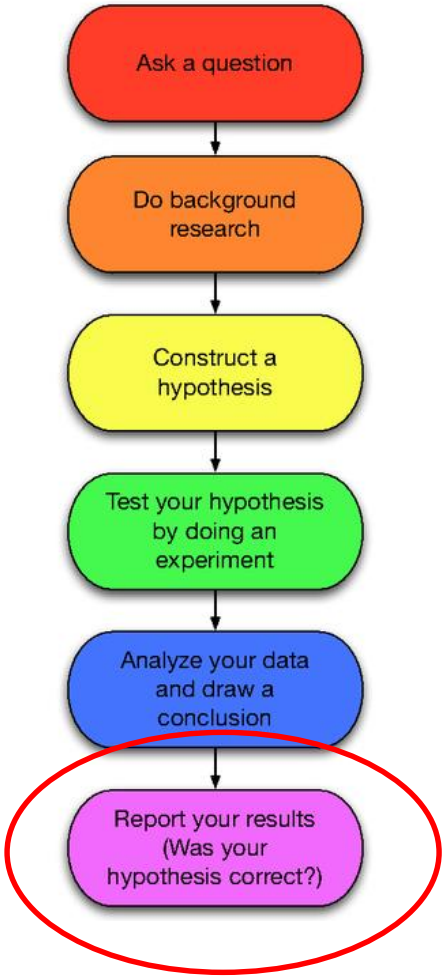

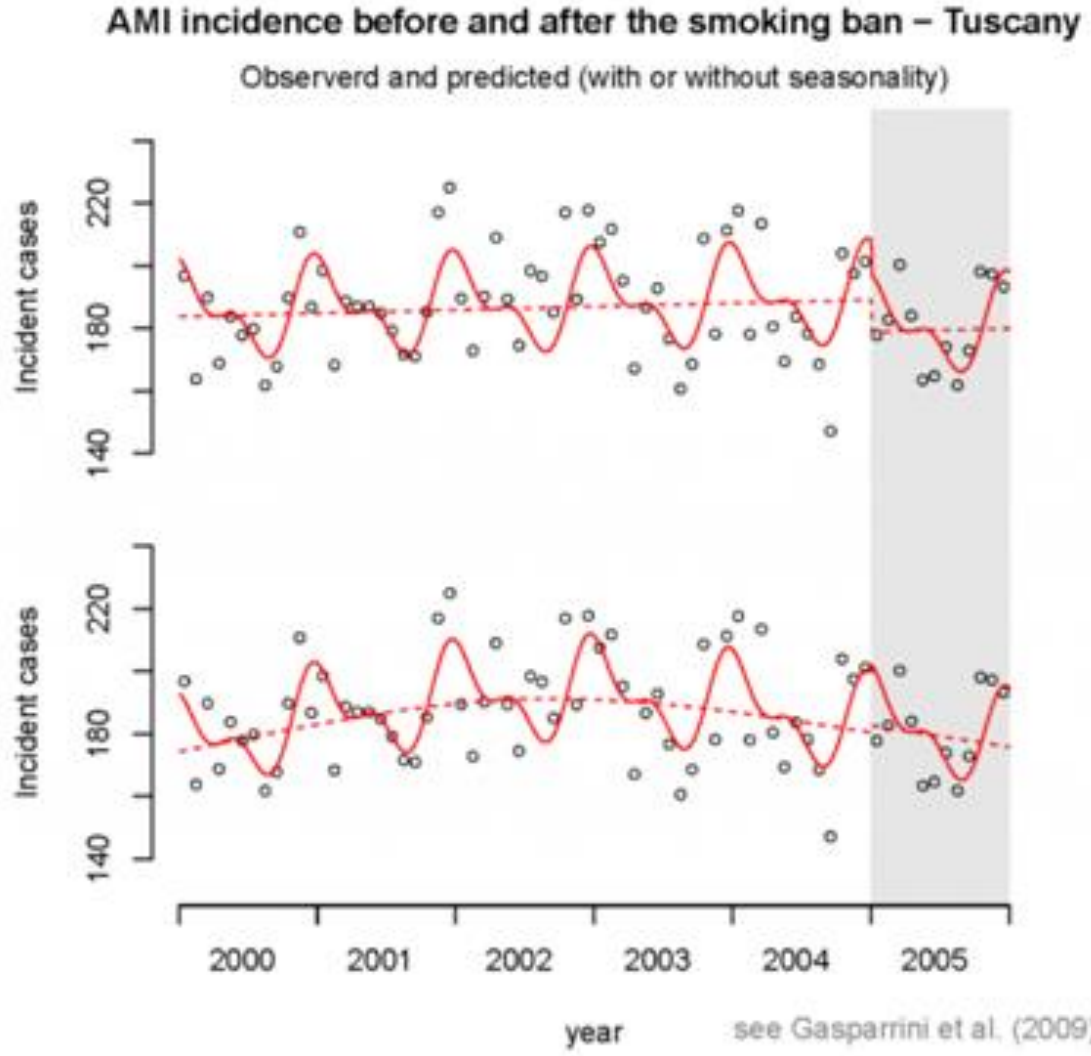

Results depend on modeling choices

Updating Hill for manipulative causation: Orienting arrows with data constraints

Temporality: Causes precede their effects	• *Directed dependence principle:* Information flows from causes to their effects over time. • *Predictive causation principle:* Changes in causes help to predict changes in their effects • Techniques: Granger causality tests, transfer entropy, directed information graphs, dynamic Bayesian networks • LiNGAM principle for linear non-Gaussian models: Prediction error distributions vary simply and predictably with predicted values

There are many ways to orient arrows from data (Granger, DBN, DIG, homoscedasticity, exogeneity, compositionality …)

But directions of some arrows may not be uniquely determined by data, or may not be clearly defined/interpretable conceptually → Data constrain possible causal graphs

Updating Hill for manipulative causation: Coherent explanation

Biological plausibility: Exposure-response association has a plausible biological mechanism

Coherence: Agrees with knowledge of disease biology

- *Coherent structural causation:* Effects are derived from their causes. Directed paths lead from causes (e.g., doses) to their effects (e.g., responses) in a causal graph
- *Coherent causal explanation:* A valid causal graph explains the observed data
- Exposures affect response probabilities via pathways (sequences of mechanisms in a causal biological graph)

Coherent causal explanation/biological plausibility: Identify paths (explanations) from exposure to response through causal biological network, consistent with data

QRA: exposure $\rightarrow$ [PBPK] $\rightarrow$ internal dose $\rightarrow$ [PD] $\rightarrow$ response

Causal mediation confirmation:

Compositionality: If $X \rightarrow Z \rightarrow Y$ is valid, then $dY/dX = (dY/dZ)(dZ/dX)$

Chapman-Kolmogorov: $P(y \mid x) = \Sigma_z P(z \mid x)P(y \mid z)$

Special cases in Hill considerations (No new criteria needed)

Experiment: Reducing exposure reduces effect	• Exogenous changes in causes produce predictable changes in the probability distributions of their effects that can be calculated via CPTs.
Specificity of effects: A specific cause produces a specific effect	• Connectivity: One or more directed paths in a causal graph lead from causes to their effects. • Direct effects of a cause are its children in a causal graph
Biological gradient: Larger responses at higher exposures	• Variations in direct causes help to predict and explain variations in their (possibly joint) effects via a CPT

Final step: Refute alternative (non-causal) explanations

Non-causal explanation	Methods for addessing non-causal associations
Unobserved (latent) confounders (Pearl and Mackenzie, 2018)	These can be tested for and their effects modeled and controlled for using the *Tetrad*, *Invariant Causal Prediction,* and *BACKSHIFT* algorithms, among others.
Spurious regression in time series or spatial observations with trends (Yule, 1926)	Spurious regression arising from coincident trends can be detected and avoided by using conditional independence tests and predictive causation (e.g., Granger causality) instead of regression models.
Collider bias; stratification or selection bias (Cole et al., 2010; Pearl and Mackenzie, 2018)	A study that stratifies or matches individuals on certain variables, such as membership in an occupation, or an analysis that conditions on certain variables by including them on the right-hand side of a regression model, can induce exposure-response associations if the variables conditioned, matched, or stratified on are common descendents of the exposure and response variables. The association does not indicate causality between exposure and response, but that they provide alternative explanations of an observed value. Such biases can be avoided by using *dagitty* to compute adjustment sets and conditioning only on variables in an adjustment set.

Refute non-causal explanations for exposure-response association

Other threats to internal validity (Campbell and Stanley, 1963)	Threats to internal validity (e.g., regression to the mean, coincident historical trends, sample selection or attrition biases, reporting biases, etc.) were enumerated by Campbell and Stanley (1963), who also discuss ways to refute them as plausible explanations, when possible, using observational data.
Model specification errors (Lenis et al., 2018; Linden et al., 2017; Pirracchio et al., 2015)	Model specification errors arise when an analysis assumes a particular parametric modeling form that does not accurately describe the data-generating process. Assuming a linear regression model when there are nonlinear effects present is one example; omitting high-order interactions terms is another. Model specification errors can be avoided by using non-parametric model ensemble methods such as PDPs.
P-hacking, i.e., adjusting modeling assumptions to produce an association (e.g., a statistically significantly positive regression coefficient). (Fraser et al., 2018)	Automated modeling using CAT or packages such as *randomForest* and *bnlearn* to automate modeling choices such as which predictors to select, how to code them (i.e., aggregate their values into ranges), and which high-order interactions to include can help to avoid p-hacking biases.

Refute non-causal explanations for exposure-response association

Omitted errors in explanatory variables. (Rhomberg et al., 2011)	Using job exposure matrices, remote-sensing and satellite imagery for pollutant concentration estimation, or other error-prone techniques for estimating exposures, creates exposure estimates for individuals that can differ substantially from their true exposures. In simple regression models, omitting errors from the estimated values of explanatory variables tends to bias regression coefficients toward the null (i.e., 0), but the bias can be in either direction in multivariate models, and failing to carefully model errors in explanatory variables can create false-positive associations. These errors and biases can be avoided by modeling errors in explanatory variables.
Omitted interdependencies among explanatory variables. (Pearl and Mackenzie, 2018; Textor et al., 2016)	Regression models that ignore dependencies among right-hand side variables can create non-causal exposure-response associations. This can be avoided by using *dagitty* to compute adjustment sets for the causal effect of exposure on response and then conditioning on variables in an adjustment set to estimate that effect.

Proposed criteria can be made operational via statistical tests

Criterion	Test	Methods
Mutual information	Reject null hypothesis that *Y* is conditionally independent of *X*	Reject null hypothesis if *X* and *Y* are linked in DAG models learned from data by causal discovery algorithms (e.g., those in the *bnlearn* package). Other tests for independence (e.g., chi-squared tests) can also be used.
Directed dependence	For longitudinal data: Reject null hypothesis that future values of *Y* are conditionally independent of past values of *X,* even after conditioning on past values of *Y* and other variables.	Reject null hypothesis if *X* and *Y* are linked in DBNs or DIGs learned from data (e.g., via the *bnstruct* package). Granger tests can also be used for time series data.
	For cross sectional data: Reject null hypothesis that direction of dependence is undetermined by data.	Reject null hypothesis if constraints determining the direction of an arrow can be identified from data (e.g., using LiNGAM, BACKSHIFT, or Simon-Iwasaki causal ordering)

Operational criteria

Internal consistency	Do not reject null hypothesis that effects estimated from different adjustment sets are the same	Reject null hypothesis if confidence bands for effects estimated from different adjustment sets do not overlap.
External consistency	Do not reject null hypothesis that response CPTs estimated from different studies or data sets are the same	Reject null hypothesis if study ID is a parent of response in DAG models learned from data.
Causal mediation confirmation	Do not reject the null hypothesis that variations in *Y* caused by variations in *X* are explained by resulting variations in mediating variables (e.g., as described by the Chapman-Kolmogorov identities implied by an explanation in the form of a probabilistic causal graph model)	Reject null hypothesis if variations in mediating variables do not explain variations in *Y* for different values of *X* (e.g., if a chi-squared test rejects the conditional independence and Chapman-Kolmogorov implications of the explanation).

Operational criteria

Causal coherence & biological plausibility	Reject the null hypothesis that identified biologically plausible pathway(s) directed from *X* to *Y* cannot explain the observed dependence of *Y* on *X*	Reject the null hypothesis if one or more biologically plausible coherent causal explanations (pathways in a causal biological network) are identified that can explain the dependence of *Y* on *X*.
Refutation of non-causal explanations	Reject the null hypothesis that the observed statistical dependence of *Y* on *X* has a non-causal explanation	Reject the null hypothesis if threats to validity are refuted (Campbell and Stanley, 1963).

Summary: Information-based causal perspectives

- Focus specifically on *manipulative* causation
- *Quantify* causal exposure-response dependence
  - Multiple paths/explanations for association
    - Qualitative determination of causality is not well-defined
  - Partial dependence plots for effects of interest
- *Connect* exposure to response probability via paths in causal graphs and causal biological networks
- *Test* implications of hypothesized manipulative causal exposure-response explanations using data
  - Mutual information, directed information flow, internal and external consistency (ICP), coherent explanation, CMC
- *Refute* non-causal explanations using data
- Result: Assess consistency of evidence with manipulative causal interpretation of exposure-response dependence using information in data sets

Thanks!