An Introduction to Causal Inference in Neuroimaging

Moritz Grosse-Wentrup

Max Planck Institute for Intelligent Systems
Department Empirical Inference
Tübingen, Germany

February 28, 2014

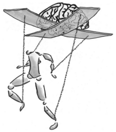

Why should we be interested in causal inference?

Why should we be interested in causal inference?

What should a scientific theory be able to do?

Why should we be interested in causal inference?

What should a scientific theory be able to do?

1. Explain already observed data

Why should we be interested in causal inference?

What should a scientific theory be able to do?

1. Explain already observed data
2. Correctly predict future observations

Why should we be interested in causal inference?

What should a scientific theory be able to do?

1. Explain already observed data
2. Correctly predict future observations
    - of systems that are passively observed.

Why should we be interested in causal inference?

What should a scientific theory be able to do?

1. Explain already observed data
2. Correctly predict future observations
   - of systems that are passively observed.
   - of systems that are actively intervened upon.

Why should we be interested in causal inference?

What should a scientific theory be able to do?

1. Explain already observed data
2. Correctly predict future observations
    - of systems that are passively observed.
    - of systems that are actively intervened upon.

$\rightarrow$The aim of causal inference is to predict how a system reacts to an intervention.

The potential outcomes framework

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

Example:

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$

Example:

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$

Example:

- $\mathcal{U}$: A set of patients

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$
- Individual *units* $u \in \mathcal{U}$

Example:

- $\mathcal{U}$: A set of patients

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$
- Individual *units* $u \in \mathcal{U}$

Example:

- $\mathcal{U}$: A set of patients
- u: An individual patient

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$
- Individual *units* $u \in \mathcal{U}$
- A *treatment variable* $S(u) : \mathcal{U} \rightarrow \{t, c\}$

Example:

- $\mathcal{U}$: A set of patients
- u: An individual patient

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$
- Individual *units* $u \in \mathcal{U}$
- A *treatment variable* $S(u) : \mathcal{U} \rightarrow \{t, c\}$

Example:

- $\mathcal{U}$: A set of patients
- u: An individual patient
- $S(u)$: Assignment of patient u to treatment- or control-group

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$
- Individual *units* $u \in \mathcal{U}$
- A *treatment variable* $S(u) : \mathcal{U} \to \{t, c\}$
- An *outcome variable* $Y(u, S(u)) : \mathcal{U} \times \{t, c\} \to \mathbb{R}$

Example:

- $\mathcal{U}$: A set of patients
- u: An individual patient
- $S(u)$: Assignment of patient u to treatment- or control-group

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Notation:

- A *population* $\mathcal{U}$
- Individual *units* $u \in \mathcal{U}$
- A *treatment variable* $S(u) : \mathcal{U} \rightarrow \{t, c\}$
- An *outcome variable* $Y(u, S(u)) : \mathcal{U} \times \{t, c\} \rightarrow \mathbb{R}$

Example:

- $\mathcal{U}$: A set of patients
- u: An individual patient
- $S(u)$: Assignment of patient u to treatment- or control-group
- $Y(u, S(u))$: The survival time of patient u under treatment $S(u)$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

1. Assume *unit homogeneity*: $Y_{t/c}(u_1) = Y_{t/c}(u_2)$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

1. Assume *unit homogeneity*: $Y_{t/c}(u_1) = Y_{t/c}(u_2)$
2. Assume *causal transience*: $Y_t(u) \perp\!\!\!\perp Y_c(u)$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

1. Assume *unit homogeneity*: $Y_{t/c}(u_1) = Y_{t/c}(u_2)$
2. Assume *causal transience*: $Y_t(u) \perp\!\!\!\perp Y_c(u)$
3. Compute the *average causal effect*: $T = E(Y_t) - E(Y_c)$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

1. Assume *unit homogeneity*: $Y_{t/c}(u_1) = Y_{t/c}(u_2)$
2. Assume *causal transience*: $Y_t(u) \perp\!\!\!\perp Y_c(u)$
3. Compute the *average causal effect*: $T = E(Y_t) - E(Y_c)$

   But we can only observe $E(Y_{t/c}|S = t/c)$.

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

1. Assume *unit homogeneity*: $Y_{t/c}(u_1) = Y_{t/c}(u_2)$
2. Assume *causal transience*: $Y_t(u) \perp\!\!\!\perp Y_c(u)$
3. Compute the *average causal effect*: $T = E(Y_t) - E(Y_c)$
   But we can only observe $E(Y_{t/c}|S = t/c)$.
   When does it hold that $E(Y_{t/c}|S = t/c) = E(Y_{t/c})$?

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

The fundamental problem of causal inference:

- The effect of cause t on u is measured as $Y_t(u) - Y_c(u)$.
- It is impossible to observe $Y_t(u)$ and $Y_c(u)$ *on the same unit*.

Solutions:

1. Assume *unit homogeneity*: $Y_{t/c}(u_1) = Y_{t/c}(u_2)$
2. Assume *causal transience*: $Y_t(u) \perp\!\!\!\perp Y_c(u)$
3. Compute the *average causal effect*: $T = E(Y_t) - E(Y_c)$

   But we can only observe $E(Y_{t/c}|S = t/c)$.

   When does it hold that $E(Y_{t/c}|S = t/c) = E(Y_{t/c})$?

   If $Y \perp\!\!\!\perp S$, i.e. if treatment assignments are done *randomly*.

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

Treatment

Survival

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

Treatment

Survival

- $E(\text{Survival}|\text{Treatment}) = E(\text{Survival}|\text{Control})$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

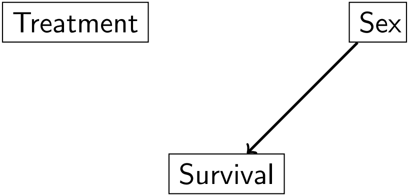

- $E(\text{Survival}|\text{Treatment}) = E(\text{Survival}|\text{Control})$
- Assume $E(\text{Survival}|\text{Sex} = \text{female}) > E(\text{Survival}|\text{Sex} = \text{male})$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

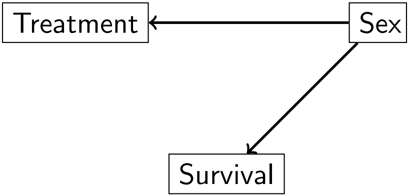

- $E(\text{Survival}|\text{Treatment}) = E(\text{Survival}|\text{Control})$
- Assume $E(\text{Survival}|\text{Sex} = \text{female}) > E(\text{Survival}|\text{Sex} = \text{male})$
- Assume $E(\text{Treatment}|\text{Sex} = \text{female}) > E(\text{Treatment}|\text{Sex} = \text{male})$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

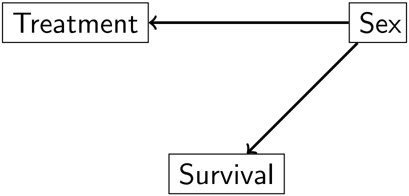

- $E(\text{Survival}|\text{Treatment}) = E(\text{Survival}|\text{Control})$
- Assume $E(\text{Survival}|\text{Sex} = \text{female}) > E(\text{Survival}|\text{Sex} = \text{male})$
- Assume $E(\text{Treatment}|\text{Sex} = \text{female}) > E(\text{Treatment}|\text{Sex} = \text{male})$
- Then $E(\text{Survival}|\text{Treatment}) > E(\text{Survival}|\text{Control})$

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

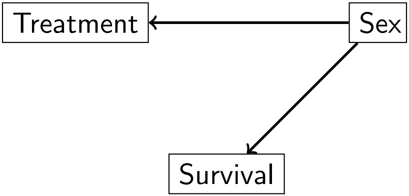

- $E(\text{Survival}|\text{Treatment}) = E(\text{Survival}|\text{Control})$
- Assume $E(\text{Survival}|\text{Sex} = \text{female}) > E(\text{Survival}|\text{Sex} = \text{male})$
- Assume $E(\text{Treatment}|\text{Sex} = \text{female}) > E(\text{Treatment}|\text{Sex} = \text{male})$
- Then $E(\text{Survival}|\text{Treatment}) > E(\text{Survival}|\text{Control})$
- Incorrect conclusion: Treatment $\rightarrow$ Survival

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

The potential outcomes framework

Example:

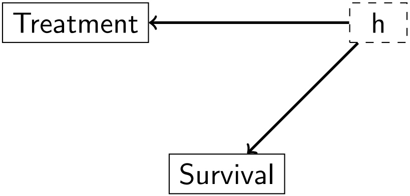

- $E(\text{Survival}|\text{Treatment}) = E(\text{Survival}|\text{Control})$
- Assume $E(\text{Survival}|\text{Sex} = \text{female}) > E(\text{Survival}|\text{Sex} = \text{male})$
- Assume $E(\text{Treatment}|\text{Sex} = \text{female}) > E(\text{Treatment}|\text{Sex} = \text{male})$
- Then $E(\text{Survival}|\text{Treatment}) > E(\text{Survival}|\text{Control})$
- Incorrect conclusion: Treatment $\rightarrow$ Survival

(Holland PW, Statistics and Causal Inference. *Journal of the American Statistical Association*, 1986)

Causal inference in neuroimaging

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
    - In general, no.

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
   - In general, no.
   - Under (usually untestable) assumptions, yes.

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
   - In general, no.
   - Under (usually untestable) assumptions, yes.
2. So what makes for a good causal inference algorithm?

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
   - In general, no.
   - Under (usually untestable) assumptions, yes.
2. So what makes for a good causal inference algorithm?
   - It is provably correct under reasonable assumptions.

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
   - In general, no.
   - Under (usually untestable) assumptions, yes.
2. So what makes for a good causal inference algorithm?
   - It is provably correct under reasonable assumptions.
   - It makes testable predictions on the effect of interventions.

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
   - In general, no.
   - Under (usually untestable) assumptions, yes.
2. So what makes for a good causal inference algorithm?
   - It is provably correct under reasonable assumptions.
   - It makes testable predictions on the effect of interventions.
   - It is able to deal with hidden confounders.

Causal inference in neuroimaging

1. Can we do causal inference based on *observational* data only?
   - In general, no.
   - Under (usually untestable) assumptions, yes.
2. So what makes for a good causal inference algorithm?
   - It is provably correct under reasonable assumptions.
   - It makes testable predictions on the effect of interventions.
   - It is able to deal with hidden confounders.
   - It performs well on finite data.

Outline

Outline

Granger causality

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Example: First-order autoregressive (AR) process

$$x[t] = ax[t-1] + \epsilon_x[t]$$
$$y[t] = by[t-1] + cx[t-1] + \epsilon_y[t]$$

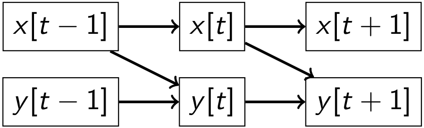

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Example: First-order autoregressive (AR) process

$$x[t] = ax[t-1] + \epsilon_x[t]$$
$$y[t] = by[t-1] + cx[t-1] + \epsilon_y[t]$$

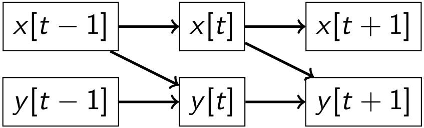

Inference procedure: $x \to y$?

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Example: First-order autoregressive (AR) process

$$x[t] = ax[t-1] + \epsilon_x[t]$$
$$y[t] = by[t-1] + cx[t-1] + \epsilon_y[t]$$

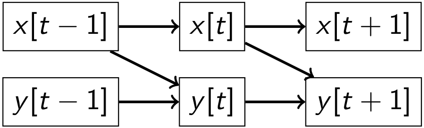

Inference procedure: $x \rightarrow y$?

- Predict $y[t]$ from its past values $\Rightarrow \sigma^2(y[t]|y[t-1])$.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Example: First-order autoregressive (AR) process

$$x[t] = ax[t-1] + \epsilon_x[t]$$
$$y[t] = by[t-1] + cx[t-1] + \epsilon_y[t]$$

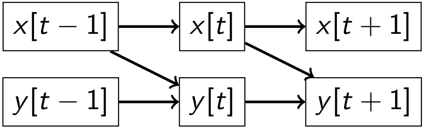

Inference procedure: $x \rightarrow y$?

- Predict $y[t]$ from its past values $\Rightarrow \sigma^2(y[t] | y[t-1])$.
- Predict $y[t]$ from the past of y and $x \Rightarrow \sigma^2(y[t] | y[t-1], x[t-1])$.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Example: First-order autoregressive (AR) process

$$x[t] = ax[t-1] + \epsilon_x[t]$$
$$y[t] = by[t-1] + cx[t-1] + \epsilon_y[t]$$

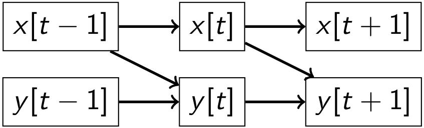

Inference procedure: $x \rightarrow y$?

- Predict $y[t]$ from its past values $\Rightarrow \sigma^2(y[t]|y[t-1])$.
- Predict $y[t]$ from the past of y and $x \Rightarrow \sigma^2(y[t]|y[t-1], x[t-1])$.
- If $\sigma^2(y[t]|y[t-1], x[t-1]) < \sigma^2(y[t]|y[t-1])$ conclude that $x \rightarrow y$.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Confounding

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Confounding

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Definition B (by Granger): If $\sigma^2(Y|U) < \sigma^2(Y|\overline{U - X})$, we say that X is causing Y, denoted by $X_t \Rightarrow Y_t$.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Confounding

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Definition B (by Granger): If $\sigma^2(Y|U) < \sigma^2(Y|\overline{U - X})$, we say that X is causing Y, denoted by $X_t \Rightarrow Y_t$.

Confounding:

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Confounding

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Definition B (by Granger): If $\sigma^2(Y|U) < \sigma^2(Y|\overline{U - X})$, we say that X is causing Y, denoted by $X_t \Rightarrow Y_t$.

Confounding:

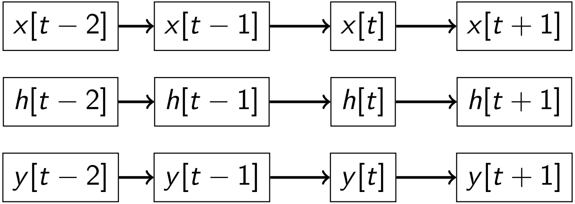

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Confounding

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Definition B (by Granger): If $\sigma^2(Y|U) < \sigma^2(Y|\overline{U - X})$, we say that X is causing Y, denoted by $X_t \Rightarrow Y_t$.

Confounding:

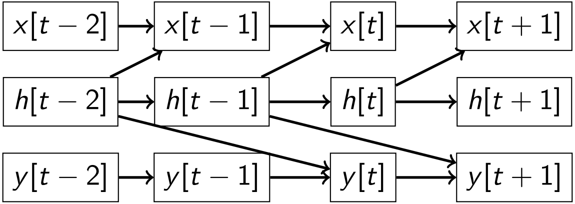

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Confounding

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Definition B (by Granger): If $\sigma^2(Y|U) < \sigma^2(Y|\overline{U - X})$, we say that X is causing Y, denoted by $X_t \Rightarrow Y_t$.

Confounding:

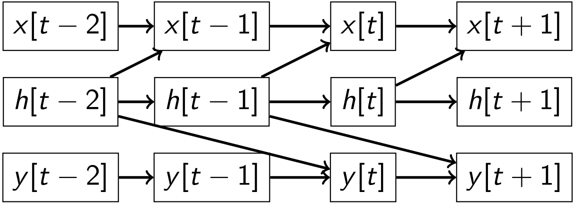

$\rightarrow$ Control for $\overline{h}$: $\sigma^2(y|\overline{y}, \overline{x}, \overline{h}) < \sigma^2(y|\overline{y}, \overline{h})$!

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Definition A (not by Granger): One stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first.

Definition B (by Granger): If $\sigma^2(Y|U) < \sigma^2(Y|\overline{U - X})$, we say that X is causing Y, denoted by $X_t \Rightarrow Y_t$.

Confounding:

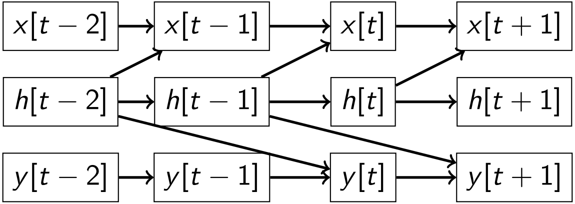

$\rightarrow$ Control for $\overline{h}$: $\sigma^2(y|\overline{y}, \overline{x}, \overline{h}) < \sigma^2(y|\overline{y}, \overline{h})$!Impossible for latent variables.

(Granger C.W.J., Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 1969)

Granger causality: Directed transfer function (DTF)

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).
- Pick an order p of the AR-process.

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ ($=$ number of signals).
- Pick an order p of the AR-process.
- Learn parameters $A[i]$ of AR-process $\mathbf{x}[t] = \sum_{i=1}^{p} A[i]\mathbf{x}[t-i] + \epsilon[t]$.

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).
- Pick an order p of the AR-process.
- Learn parameters $A[i]$ of AR-process $\mathbf{x}[t] = \sum_{i=1}^{p} A[i]\mathbf{x}[t-i] + \epsilon[t]$.
- Let $A[0] = -I$:

$$-\sum_{i=0}^{p} A[i]\mathbf{x}[t-i] = \epsilon[t]$$

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).
- Pick an order p of the AR-process.
- Learn parameters $A[i]$ of AR-process $\mathbf{x}[t] = \sum_{i=1}^{p} A[i]\mathbf{x}[t-i] + \epsilon[t]$.
- Let $A[0] = -I$:

$$
\begin{aligned}
& -\sum_{i=0}^{p} A[i]\mathbf{x}[t-i] = \epsilon[t] \\
\Leftrightarrow \quad & -A[t] * \mathbf{x}[t] = \epsilon[t]
\end{aligned}
$$

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).
- Pick an order p of the AR-process.
- Learn parameters $A[i]$ of AR-process $\mathbf{x}[t] = \sum_{i=1}^{p} A[i]\mathbf{x}[t-i] + \boldsymbol{\epsilon}[t]$.
- Let $A[0] = -I$:

$$\begin{aligned} & -\sum_{i=0}^{p} A[i]\mathbf{x}[t-i] = \boldsymbol{\epsilon}[t] \\ \Leftrightarrow \quad & -A[t] * \mathbf{x}[t] = \boldsymbol{\epsilon}[t] \\ \overset{\mathcal{F}}{\leftrightarrow} \quad & -A(f)\mathbf{x}(f) = \boldsymbol{\epsilon}(f) \end{aligned}$$

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).
- Pick an order p of the AR-process.
- Learn parameters $A[i]$ of AR-process $\mathbf{x}[t] = \sum_{i=1}^{p} A[i]\mathbf{x}[t-i] + \epsilon[t]$.
- Let $A[0] = -I$:

$$\begin{aligned} & -\sum_{i=0}^{p} A[i]\mathbf{x}[t-i] = \epsilon[t] \\ \Leftrightarrow \quad & -A[t] * \mathbf{x}[t] = \epsilon[t] \\ \overset{\mathcal{F}}{\leftrightarrow} \quad & -A(f)\mathbf{x}(f) = \epsilon(f) \\ \Leftrightarrow \quad & H(f) = \frac{\mathbf{x}(f)}{\mathbf{e}(f)} = -A^{-1}(f). \end{aligned}$$

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Directed transfer function (DTF)

- Observe T samples of $\mathbf{x}[t] \in R^N$ (= number of signals).
- Pick an order p of the AR-process.
- Learn parameters $A[i]$ of AR-process $\mathbf{x}[t] = \sum_{i=1}^{p} A[i]\mathbf{x}[t-i] + \boldsymbol{\epsilon}[t]$.
- Let $A[0] = -I$:

$$\begin{aligned} & -\sum_{i=0}^{p} A[i]\mathbf{x}[t-i] = \boldsymbol{\epsilon}[t] \\ \Leftrightarrow \quad & -A[t] * \mathbf{x}[t] = \boldsymbol{\epsilon}[t] \\ \overset{\mathcal{F}}{\Leftrightarrow} \quad & -A(f)\mathbf{x}(f) = \boldsymbol{\epsilon}(f) \\ \Leftrightarrow \quad & H(f) = \frac{\mathbf{x}(f)}{\mathbf{e}(f)} = -A^{-1}(f). \end{aligned}$$

- $h_{ij}(f)$ describes the frequency-specific effect of $x_j[t]$ on $x_i[t]$.

(Kaminski et al., Evaluating causal relations in neural systems. *Biological Cybernetics*, 2001)

Granger causality: Case study

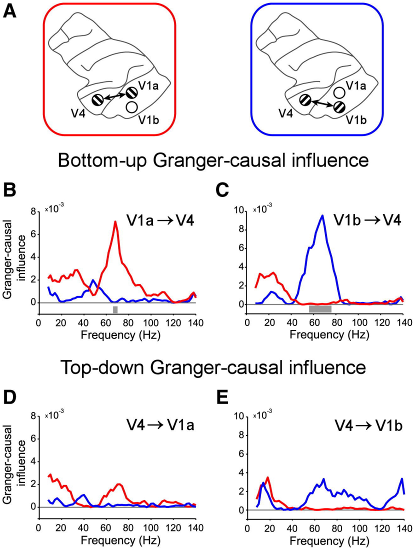

(Bosman et al., Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron*, 2012)

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions				
Testable interventions				
Hidden confounders				
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗			
Testable interventions				
Hidden confounders				
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗			
Testable interventions	✓			
Hidden confounders				
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗			
Testable interventions	✓			
Hidden confounders	✗			
Empirical performance				

Outline

Causal structure
(unkown)

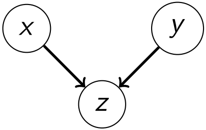

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

Causal structure
(unkown)

Empirical data
(observable)

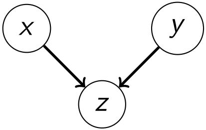

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

Causal structure
(unkown)

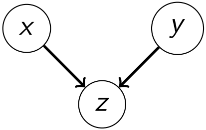

Empirical data
(observable)

$$\{x_1, y_1, z_1\}$$
$$\{x_2, y_2, z_2\}$$
$$\vdots$$
$$\{x_N, y_N, z_N\}$$

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

Causal structure (unkown)

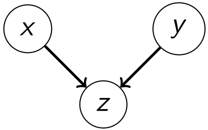

Empirical data (observable)

$$\{x_1, y_1, z_1\}$$
$$\{x_2, y_2, z_2\}$$
$$\vdots$$
$$\{x_N, y_N, z_N\}$$
$$\rightarrow p(x, y, z)$$

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

Causal structure (unkown)

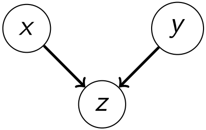

Empirical data (observable)

$$\{x_1, y_1, z_1\}$$
$$\{x_2, y_2, z_2\}$$
$$\vdots$$
$$\{x_N, y_N, z_N\}$$
$$\rightarrow p(x, y, z)$$

Dependency structure (inferable)

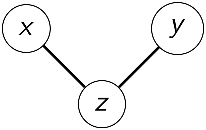

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

Causal structure

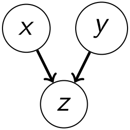

Dependency structure

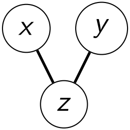

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

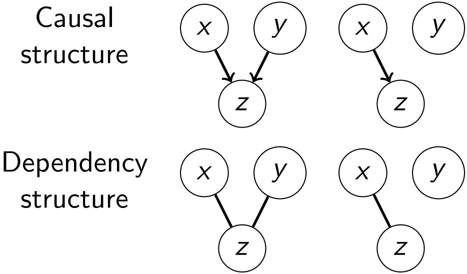

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

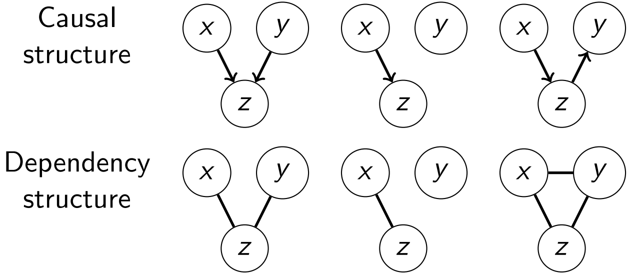

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Introductory example

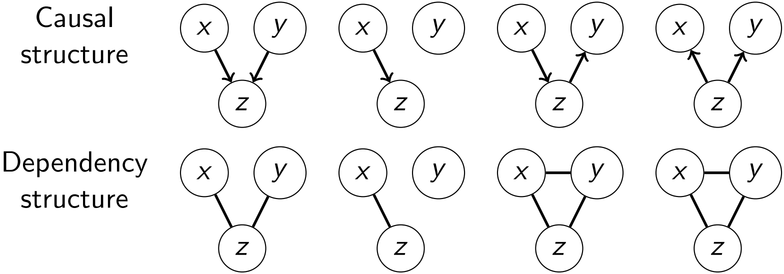

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Under the assumptions of *faithfulness* and *causal sufficiency*, the following conditions are sufficient for x to be a cause of y:

- $x \not\perp\!\!\!\perp z$
- $y \not\perp\!\!\!\perp z$
- $x \perp\!\!\!\perp y$

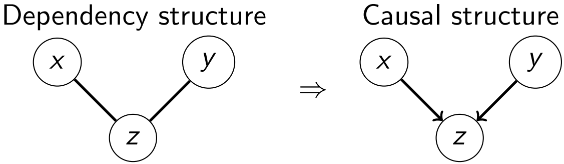

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Under the assumptions of *faithfulness* and *causal sufficiency*, the following conditions are sufficient for x to be a cause of y:

- $x \not\perp\!\!\!\perp z$
- $y \not\perp\!\!\!\perp z$
- $x \perp\!\!\!\perp y$

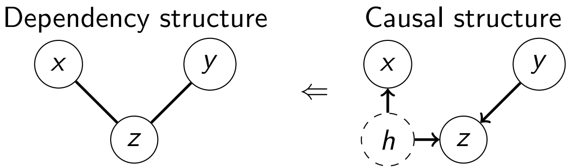

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Spurious association

Under the assumption of *faithfulness*, the following conditions are sufficient for x and y to be spuriously related:

- $x \not\perp\!\!\!\perp y$
- $x \not\perp\!\!\!\perp z$
- $y \not\perp\!\!\!\perp w$
- $z \perp\!\!\!\perp w$
- $x \perp\!\!\!\perp w$
- $z \perp\!\!\!\perp y$

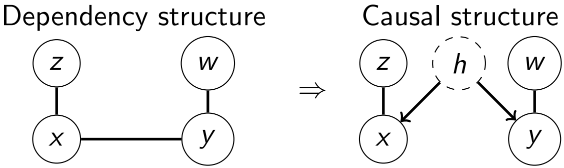

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Spurious association

Under the assumption of *faithfulness*, the following conditions are sufficient for x and y to be spuriously related:

- $x \not\perp\!\!\!\perp y$
- $x \not\perp\!\!\!\perp z$
- $y \not\perp\!\!\!\perp w$
- $z \perp\!\!\!\perp w$
- $x \perp\!\!\!\perp w$
- $z \perp\!\!\!\perp y$

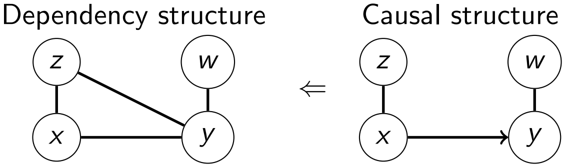

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Spurious association

Under the assumption of *faithfulness*, the following conditions are sufficient for x and y to be spuriously related:

- $x \not\perp\!\!\!\perp y$
- $x \not\perp\!\!\!\perp z$
- $y \not\perp\!\!\!\perp w$
- $z \perp\!\!\!\perp w$
- $x \perp\!\!\!\perp w$
- $z \perp\!\!\!\perp y$

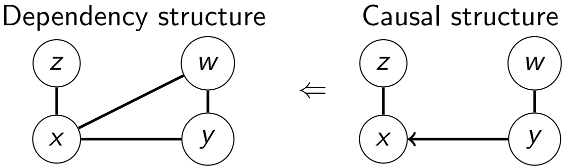

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Genuine causation

Under the assumption of *faithfulness*, the following conditions are sufficient for x to be a genuine cause of y:

- z is a potential cause of x
- $z \not\perp\!\!\!\perp y$
- $z \perp\!\!\!\perp y | x$

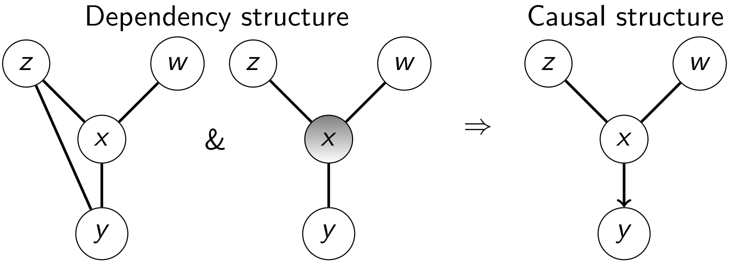

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

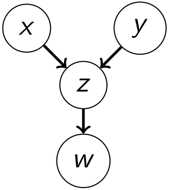

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

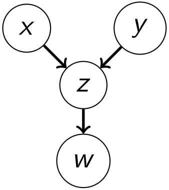

1. Factorize $p(x, y, z, w)$ accoding to it's DAG:

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

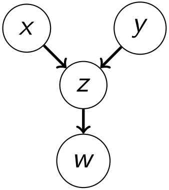

1. Factorize $p(x, y, z, w)$ accoding to it's DAG:
$p(x, y, z, w) = p(w|z)p(z|x, y)p(x)p(y)$

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

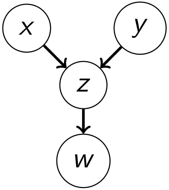

1. Factorize $p(x, y, z, w)$ accoding to it's DAG:
   $p(x, y, z, w) = p(w|z)p(z|x, y)p(x)p(y)$
2. Compute the interventional distribution:

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

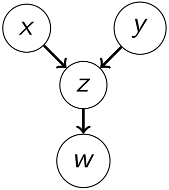

1. Factorize $p(x, y, z, w)$ accoding to it's DAG:
   $p(x, y, z, w) = p(w|z)p(z|x, y)p(x)p(y)$
2. Compute the interventional distribution:
   $p(x, y, z, w|\text{do}(x = \tilde{x})) = p(w|z)p(z|\tilde{x}, y)p(y)$

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

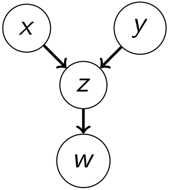

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

1. Factorize $p(x, y, z, w)$ accoding to it's DAG:
   $p(x, y, z, w) = p(w|z)p(z|x, y)p(x)p(y)$
2. Compute the interventional distribution:
   $p(x, y, z, w|\text{do}(x = \tilde{x})) = p(w|z)p(z|\tilde{x}, y)p(y)$
3. Compute marginal distribution of variable of interest:

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Predicting interventions

Given: Causal structure (DAG) & $p(x, y, z, w)$
Goal: Predict the effect of experimentally controlling x

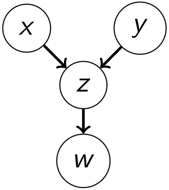

1. Factorize $p(x, y, z, w)$ accoding to it's DAG:
   $p(x, y, z, w) = p(w|z)p(z|x, y)p(x)p(y)$
2. Compute the interventional distribution:
   $p(x, y, z, w|\mathrm{do}(x = \tilde{x})) = p(w|z)p(z|\tilde{x}, y)p(y)$
3. Compute marginal distribution of variable of interest:
   $p(w|\mathrm{do}(x = \tilde{x})) = \int_z \int_y p(w|z)p(z|\tilde{x}, y)p(y)\mathrm{d}y\mathrm{d}z$

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Faithfulness

Faithfulness: All observed (conditional) independences are *structural*.

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Faithfulness: All observed (conditional) independences are *structural*.

Causal structure	Functional model	Dependency structure
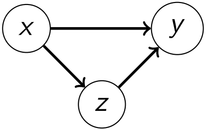	$x = u_x$ $y = ax + bz + u_y$ $z = cx + u_z$	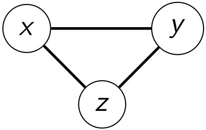

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Faithfulness

Faithfulness: All observed (conditional) independences are *structural*.

Causal structure | Functional model | Dependency structure

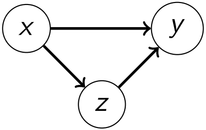

$$x = u_x$$
$$y = ax + bz + u_y$$
$$(\text{if } a = -bc)$$
$$= bu_z + u_y$$
$$z = cx + u_z$$

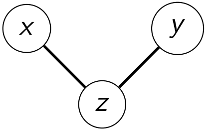

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Faithfulness

Faithfulness: All observed (conditional) independences are *structural*.

Causal structure | Functional model | Dependency structure

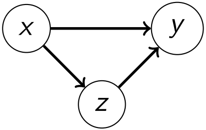

$$x = u_x$$
$$y = ax + bz + u_y$$
$$\text{(if } a = -bc\text{)}$$
$$= bu_z + u_y$$
$$z = cx + u_z$$

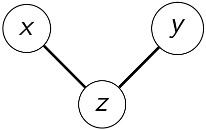

For a given causal structure (DAG), the unfaithful distributions have measure zero in the space of all distributions that can be generated by the DAG (Meek. *UAI*, 1995).

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Conditional independence tests

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Conditional independence tests

- Independence tests have to be carried out on finite data.

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Conditional independence tests

- Independence tests have to be carried out on finite data.
- Uncorrelatedness does not imply independence.

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Conditional independence tests

- Independence tests have to be carried out on finite data.
- Uncorrelatedness does not imply independence.
- Nonlinear independence tests are hard (Gretton et al., *JMLR*, 2012).

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Conditional independence tests

- Independence tests have to be carried out on finite data.
- Uncorrelatedness does not imply independence.
- Nonlinear independence tests are hard (Gretton et al., *JMLR*, 2012).
- Conditional non-linear independence tests are very hard (Zhang et al. *UAI*, 2012).

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

Causal Bayesian Networks: Conditional independence tests

- Independence tests have to be carried out on finite data.
- Uncorrelatedness does not imply independence.
- Nonlinear independence tests are hard (Gretton et al., *JMLR*, 2012).
- Conditional non-linear independence tests are very hard (Zhang et al. *UAI*, 2012).
- Not finding a dependence is not evidence for independence.

(Pearl J., *Causality: Models, reasoning, and inference*, 2000)

What are the causes of inter-subject variations in performance when operating a sensorimotor brain-computer interface (BCI)?

(Grosse-Wentrup et al. Causal influence of gamma-oscillations on the sensorimotor-rhythm. *NeuroImage*, 2011)

Causal Bayesian Networks: Case study

What are the causes of inter-subject variations in performance when operating a sensorimotor brain-computer interface (BCI)?

Causal hypothesis

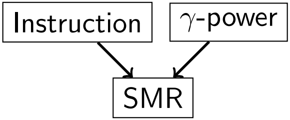

(Grosse-Wentrup et al. Causal influence of gamma-oscillations on the sensorimotor-rhythm. *NeuroImage*, 2011)

What are the causes of inter-subject variations in performance when operating a sensorimotor brain-computer interface (BCI)?

Causal hypothesis

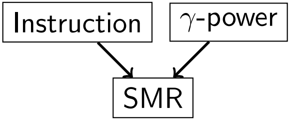

Empirical independence stucture

$\perp\!\!\!\perp$ /p	Instruction	SMR	γ-power
Instruction	-	1e-4	0.44
SMR	$\not\perp\!\!\!\perp$	-	2e-4
γ-power	$\perp\!\!\!\perp$	$\not\perp\!\!\!\perp$	-

(Grosse-Wentrup et al. Causal influence of gamma-oscillations on the sensorimotor-rhythm. *NeuroImage*, 2011)

Causal Bayesian Networks: Case study

What are the causes of inter-subject variations in performance when operating a sensorimotor brain-computer interface (BCI)?

Causal hypothesis

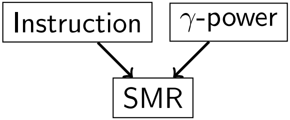

Empirical independence stucture

$\perp\!\!\!\perp$ /p	Instruction	SMR	γ-power
Instruction	-	1e-4	0.44
SMR	$\not\perp\!\!\!\perp$	-	2e-4
γ-power	$\perp\!\!\!\perp$	$\not\perp\!\!\!\perp$	-

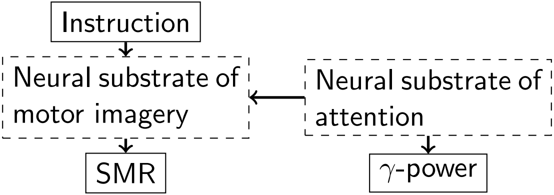

(Grosse-Wentrup et al. Causal influence of gamma-oscillations on the sensorimotor-rhythm. *NeuroImage*, 2011)

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗			
Testable interventions	✓			
Hidden confounders	✗			
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓		
Testable interventions	✓			
Hidden confounders	✗			
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓		
Testable interventions	✓	✓		
Hidden confounders	✗			
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓		
Testable interventions	✓	✓		
Hidden confounders	✗	✓		
Empirical performance				

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓		
Testable interventions	✓	✓		
Hidden confounders	✗	✓		
Empirical performance		✗		

Outline

Dynamic causal modelling (DCM)

Causality in DCM is used in a control theory sense and means that, under the model, activity in one brain area causes dynamics in another, and that these dynamics cause the observations. (Friston, 2009)

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling (DCM)

Causality in DCM is used in a control theory sense and means that, under the model, activity in one brain area causes dynamics in another, and that these dynamics cause the observations. (Friston, 2009)

Inference procedure:

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling (DCM)

Causality in DCM is used in a control theory sense and means that, under the model, activity in one brain area causes dynamics in another, and that these dynamics cause the observations. (Friston, 2009)

Inference procedure:

1. Observe an N-dimensional time-series $\mathbf{x}(t) \in \mathbb{R}^N$ for $t = 1, \ldots, T$ (e.g., BOLD signals).

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling (DCM)

Causality in DCM is used in a control theory sense and means that, under the model, activity in one brain area causes dynamics in another, and that these dynamics cause the observations. (Friston, 2009)

Inference procedure:

1. Observe an N-dimensional time-series $\mathbf{x}(t) \in \mathbb{R}^N$ for $t = 1, \ldots, T$ (e.g., BOLD signals).
2. Define a set of M models $\mathcal{M} = \{m_1, \ldots, m_M\}$, where each model consists of a set of differential equations with a different connectivity structure.

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling (DCM)

Causality in DCM is used in a control theory sense and means that, under the model, activity in one brain area causes dynamics in another, and that these dynamics cause the observations. (Friston, 2009)

Inference procedure:

1. Observe an N-dimensional time-series $\mathbf{x}(t) \in \mathbb{R}^N$ for $t = 1, \ldots, T$ (e.g., BOLD signals).
2. Define a set of M models $\mathcal{M} = \{m_1, \ldots, m_M\}$, where each model consists of a set of differential equations with a different connectivity structure.
3. Fit each model to the data (which is a tough problem).

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling (DCM)

Causality in DCM is used in a control theory sense and means that, under the model, activity in one brain area causes dynamics in another, and that these dynamics cause the observations. (Friston, 2009)

Inference procedure:

1. Observe an N-dimensional time-series $\mathbf{x}(t) \in \mathbb{R}^N$ for $t = 1, \ldots, T$ (e.g., BOLD signals).
2. Define a set of M models $\mathcal{M} = \{m_1, \ldots, m_M\}$, where each model consists of a set of differential equations with a different connectivity structure.
3. Fit each model to the data (which is a tough problem).
4. Take the connectivity of the model with the best data fit as the true causal structure.

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling: The hemodynamic model

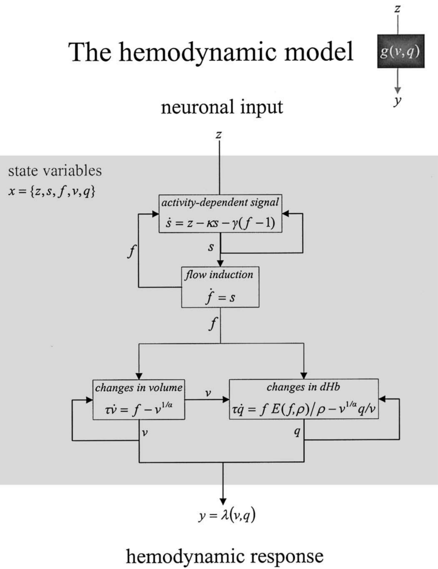

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling: The bilinear model

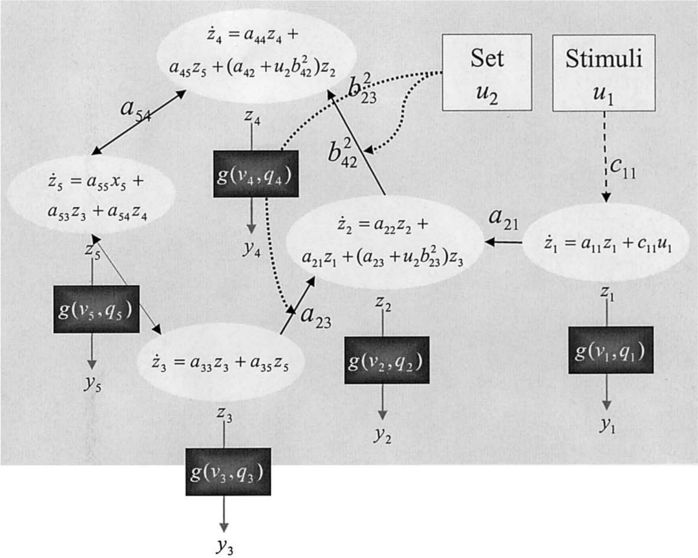

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling: Model comparison

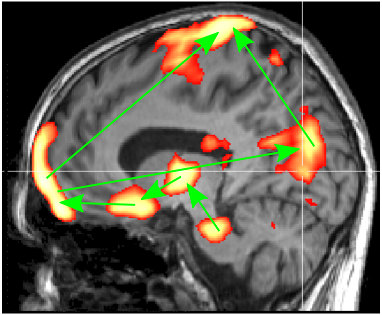

vs.

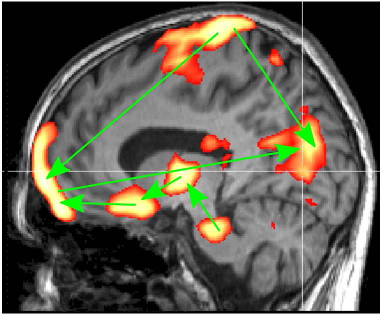

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling

Inference procedure:

1. Observe an N-dimensional time-series $\mathbf{x}(t) \in \mathbb{R}^N$ for $t = 1, \ldots, T$ (e.g., BOLD signals).
2. Define a set of M models $\mathcal{M} = \{m_1, \ldots, m_M\}$, where each model consists of a set of differential equations with a different connectivity structure.
3. Fit each model to the data (which is a tough problem).
4. Take the connectivity of the model with the best data fit as the true causal structure.

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling

Inference procedure:

1. Observe an N-dimensional time-series $\mathbf{x}(t) \in \mathbb{R}^N$ for $t = 1, \ldots, T$ (e.g., BOLD signals).
2. Define a set of M models $\mathcal{M} = \{m_1, \ldots, m_M\}$, where each model consists of a set of differential equations with a different connectivity structure.
3. Fit each model to the data (which is a tough problem).
4. Take the connectivity of the model with the best data fit as the true causal structure.

If $\mathcal{M}$ does not contain the true model, is the best-fitting model in $\mathcal{M}$ similar to the true one in terms of its connectivity structure?

(Friston K.J. et al., Dynamic causal modelling. *NeuroImage*, 2003)

Dynamic causal modelling: Model fit & structure similarity

(Lohmann et al., Critical comments on dynamic causal modelling. *NeuroImage*, 2012)

Dynamic causal modelling: Model fit & structure similarity

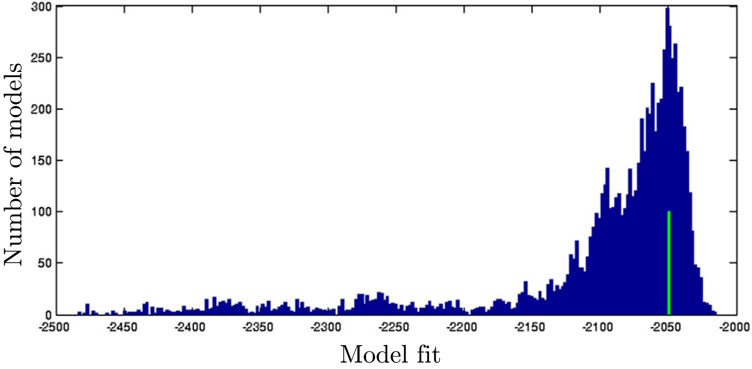

(Lohmann et al., Critical comments on dynamic causal modelling. *NeuroImage*, 2012)

Dynamic causal modelling: Model fit & structure similarity

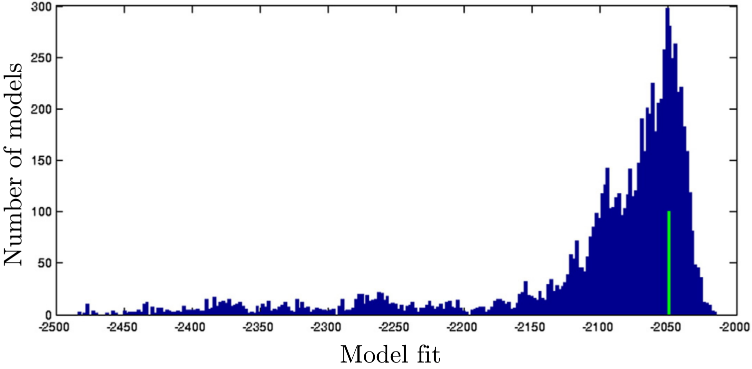

→ Similarity in terms of model fit does not translate into similarity in terms of connectivity structure.

(Lohmann et al., Critical comments on dynamic causal modelling. *NeuroImage*, 2012)

Dynamic causal modelling: Model fit & structure similarity

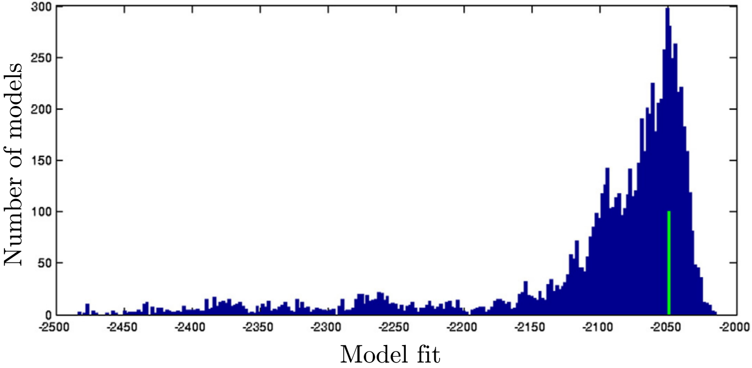

→ Similarity in terms of model fit does not translate into similarity in terms of connectivity structure.

→ There is no reason to believe that DCM selects a causal structure that is structurally similar to the true one.

(Lohmann et al., Critical comments on dynamic causal modelling. *NeuroImage*, 2012)

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓		
Testable interventions	✓	✓		
Hidden confounders	✗	✓		
Empirical performance		✗		

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	
Testable interventions	✓	✓		
Hidden confounders	✗	✓		
Empirical performance		✗		

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	
Testable interventions	✓	✓	✗	
Hidden confounders	✗	✓		
Empirical performance		✗		

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	
Testable interventions	✓	✓	✗	
Hidden confounders	✗	✓	✗	
Empirical performance		✗		

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	
Testable interventions	✓	✓	✗	
Hidden confounders	✗	✓	✗	
Empirical performance		✗	✗	

Outline

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

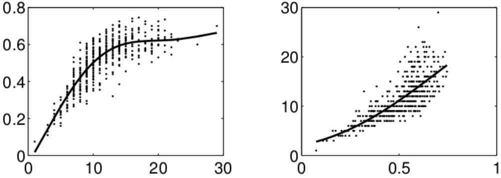

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

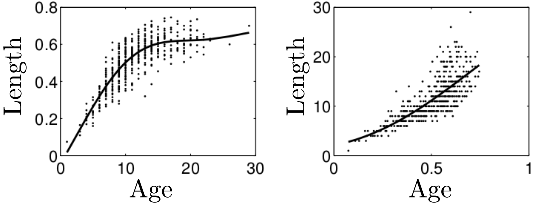

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

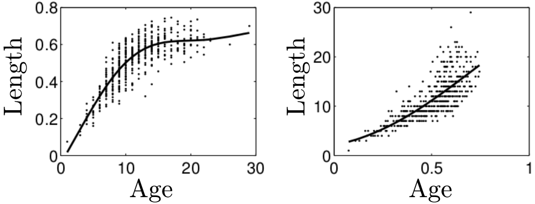

Let

$$y = f(x) + n$$

for some arbitrary non-linear function f and $p(x, n) = p(x)p(n)$ $(x \perp\!\!\!\perp n)$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

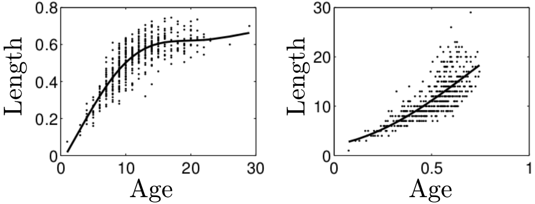

Let

$$y = f(x) + n$$

for some arbitrary non-linear function f and $p(x, n) = p(x)p(n)$ $(x \perp\!\!\!\perp n)$.
Is it possible to invert this model to obtain

$$x = g(y) + \tilde{n}$$

with some arbitrary non-linear function g and $p(y, \tilde{n}) = p(y)p(\tilde{n})$ $(y \perp\!\!\!\perp \tilde{n})$?

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

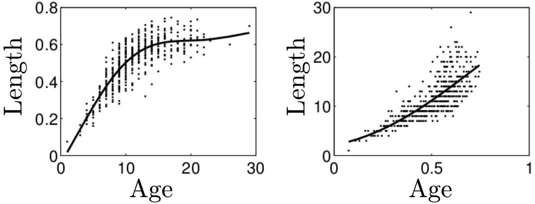

Let

$$y = f(x) + n$$

for some arbitrary non-linear function f and $p(x, n) = p(x)p(n)$ $(x \perp\!\!\!\perp n)$.
Is it possible to invert this model to obtain

$$x = g(y) + \tilde{n}$$

with some arbitrary non-linear function g and $p(y, \tilde{n}) = p(y)p(\tilde{n})$ $(y \perp\!\!\!\perp \tilde{n})$?
$\rightarrow$ In general, no!

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

1. Observe N samples of $\{x_i, y_i\}$ with $i = 1, \ldots, N$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

1. Observe N samples of $\{x_i, y_i\}$ with $i = 1, \dots, N$.
2. Perform a non-linear regression from x to y and test whether for the residuals e_y it holds that $e_y \perp\!\!\!\perp x$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

1. Observe N samples of $\{x_i, y_i\}$ with $i = 1, \ldots, N$.
2. Perform a non-linear regression from x to y and test whether for the residuals e_y it holds that $e_y \perp\!\!\!\perp x$.
3. Perform a non-linear regression from y to x and test whether for the residuals e_x it holds that $e_x \perp\!\!\!\perp y$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

1. Observe N samples of $\{x_i, y_i\}$ with $i = 1, \ldots, N$.
2. Perform a non-linear regression from x to y and test whether for the residuals e_y it holds that $e_y \perp\!\!\!\perp x$.
3. Perform a non-linear regression from y to x and test whether for the residuals e_x it holds that $e_x \perp\!\!\!\perp y$.
4. If $e_y \perp\!\!\!\perp x$ decide that $x \rightarrow y$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

1. Observe N samples of $\{x_i, y_i\}$ with $i = 1, \dots, N$.
2. Perform a non-linear regression from x to y and test whether for the residuals e_y it holds that $e_y \perp\!\!\!\perp x$.
3. Perform a non-linear regression from y to x and test whether for the residuals e_x it holds that $e_x \perp\!\!\!\perp y$.
4. If $e_y \perp\!\!\!\perp x$ decide that $x \rightarrow y$.
5. If $e_x \perp\!\!\!\perp y$ decide that $y \rightarrow x$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Non-Linear Non-Gaussian Acyclic Models (Non-LiNGAM)

Inference procedure:

1. Observe N samples of $\{x_i, y_i\}$ with $i = 1, \ldots, N$.
2. Perform a non-linear regression from x to y and test whether for the residuals e_y it holds that $e_y \perp\!\!\!\perp x$.
3. Perform a non-linear regression from y to x and test whether for the residuals e_x it holds that $e_x \perp\!\!\!\perp y$.
4. If $e_y \perp\!\!\!\perp x$ decide that $x \rightarrow y$.
5. If $e_x \perp\!\!\!\perp y$ decide that $y \rightarrow x$.
6. Do not decide on causal direction if neither $e_y \perp\!\!\!\perp x$ nor $e_x \perp\!\!\!\perp y$.

(Hoyer et al., Nonlinear causal discovery with additive noise models. *NIPS*, 2008)

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	
Testable interventions	✓	✓	✗	
Hidden confounders	✗	✓	✗	
Empirical performance		✗	✗	

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	✓
Testable interventions	✓	✓	✗	
Hidden confounders	✗	✓	✗	
Empirical performance		✗	✗	

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	✓
Testable interventions	✓	✓	✗	✓
Hidden confounders	✗	✓	✗	
Empirical performance		✗	✗	

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	✓
Testable interventions	✓	✓	✗	✓
Hidden confounders	✗	✓	✗	✓
Empirical performance		✗	✗	

Outline

Empirical Performance

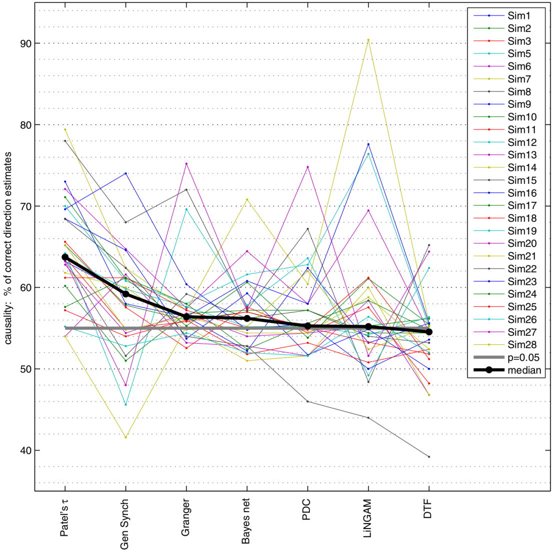

(Smith et al., Network modelling methods for fMRI. *NeuroImage*, 2011)

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	✓
Testable interventions	✓	✓	✗	✓
Hidden confounders	✗	✓	✗	✓
Empirical performance			✗	

Causal inference in neuroimaging

	GC	CBN	DCM	Non-LiNGAM
Provably correct under reasonable assumptions	✗	✓	✗	✓
Testable interventions	✓	✓	✗	✓
Hidden confounders	✗	✓	✗	✓
Empirical performance	✗	✗	✗	✗

Conclusions

Conclusions

- Every causal inference algorithms rests on *untestable* assumptions.

Conclusions

- Every causal inference algorithms rests on *untestable* assumptions.
- Several causal inference algorithms appear to perform *above chance-level.*

Conclusions

- Every causal inference algorithms rests on *untestable* assumptions.
- Several causal inference algorithms appear to perform *above chance-level.*
- Causal inference may be useful

Conclusions

- Every causal inference algorithms rests on *untestable* assumptions.
- Several causal inference algorithms appear to perform *above chance-level.*
- Causal inference may be useful
  - to guide the design of interventional studies

Conclusions

- Every causal inference algorithms rests on *untestable* assumptions.
- Several causal inference algorithms appear to perform *above chance-level.*
- Causal inference may be useful
  - to guide the design of interventional studies
  - when qualitative conclusions do not depend on individual results.

Conclusions

- Every causal inference algorithms rests on *untestable* assumptions.
- Several causal inference algorithms appear to perform *above chance-level.*
- Causal inference may be useful
  - to guide the design of interventional studies
  - when qualitative conclusions do not depend on individual results.
- Causal inference is (at present) not useful, when qualitative conclusions depend on one individual inference result.

4th Int. Workshop on Pattern Recognition in Neuroimaging (PRNI 2014)

June 4-6, 2014, Tübingen
http://prni.org