# J.
# Conclusions



# Today

- Summary and final words, 'take-home message'
- Brief info about the upcoming exam
- Preliminary information about the project work
- When and how to provide course feedback



# Predictions, predictions, ...

- For predictions based on observations only, in an unchanging world, there is no need for causal concepts. We only need the concept of association! Examples:

An insurance firm decides on the cost of insurance for people with different kinds of cars, based on amount of past claims. For ecample: if red cars have frequent accidents, we can make their owners pay more without caring what the reason is. (Of course, predictions based on driving style might give better results...)

$P(\text{accident} \mid \text{red car}) = ?$

Forecasting the weather based on past and present measurement data (we don't care why it starts to rain, or what would happen if we made it rain, rather we only care about associations)

$P(\text{rain today} \mid \text{rain yesterday}) = ?$



- Predictions based on actions or novel events require understanding of causal effects! Examples:

I am trying to decide whether to buy a red car or a blue car. Will I have more accidents with a red car than with a blue car?

$P(\text{accident} \mid \text{do(red car)}) = ?$

We have to decide whether to limit emissions of greenhouse gases. Although in the past the $CO_2$ levels and the temperature have been highly correlated, is the relationship causal? Probably yes, as domain-specific models indicate, but if we had to base our argument on the past correlation only then we could not be certain.

$P(\text{temperature} \mid \text{do(greenhouse gases)}) = ?$

- So, in principle, a clear and crisp division of concepts. In practice, much more difficult and subtle! Examples:

  Due to recent fashion, red cars are becoming much more common than before. Should the insurance company update its pricing scheme? The answer depends on the causal connections between fashion, driver personality, car color, and accidents!

  Somebody else is making a decision. In his/her mind, the decision is clearly an exogeneous variable (having no parents, and giving no information about the past), but from your perspective his/her decision might provide some information about the past (and hence his/her decision is in your model endogenous).

⇒ We need notation which can express the various assumptions and particularities of any causal problem: probabilities, graphs & do- vs see-conditioning provide the necessary language



# Counterfactuals?

- If he had been president, we would now have...
- If wages had been lower, businesses would not have moved...
- What would have happened if...

  ...and...

- Drunk driving was the cause of the crash...
- A main factor behind the Second World War was the Treaty of Versailles ending the First World War...



- I hope to have convinced you that discussing such questions is in most cases irrelevant, since...
  - it is not clear when we can reasonably make the 'persistence' assumptions needed (i.e. which other variables would have attained the same values in the alternative scenario)
  - but more to the point: answering such questions is not relevant for either
    - evaluating credit/responsibility (decisions should be evaluated on the basis of expected outcomes at time of decision, not on what eventually happened!), nor
    - future decisions (future decisions depend on queries such as $P(\text{future} \mid \text{do}(\text{action})) = ?$ , not on counterfactuals!)



- ...but counterfactuals may seem useful in cases where the outcome provides clues about hidden variables on which we assign credit/responsibility
  - however, counterfactuals are not strictly needed here!

- ...and sometimes counterfactual questions are indirectly relevant:
  - decisions typically need to be made in a 'persistent' context, and computing counterfactuals in this case gives answers to predictive questions (the key is to understand the predictive implications of the counterfactual questions and hence be able to convert the counterfactuals into the relevant predictive questions, which are, at least in principle, answerable using controlled experiments and are not counterfactual!)

# Ingredients for causal analysis

- Background knowledge (expert domain knowledge)
  - We might know the presence/absence of various causal connections (i.e. we might know something about the generating graph)
  - Time order of the variables
  - Quantitative knowledge of strengths of connections?
  - In all cases, we might be uncertain about the knowledge, but not completely ignorant (Bayesian: we have a prior over various causal models)



- Data
  - Controlled, randomized experiments (We directly control the values of certain variables. Note that sometimes we are only indirectly controlling our target variable, leading to imperfect experiments...)

| | | | |
|---|---|---|---|
| 271 | -0.2359217915 | 0.726533516 | -1.2132623 |
| 276 | 0.6480936623 | 0.543305086 | 1.1861087 |
| 360 | 1.6462430548 | -1.388650662 | -0.3220947 |
| 403 | -0.0281080387 | -0.431375402 | -1.1740618 |
| 507 | 1.0760912272 | -0.999268658 | -0.0012266 |
| 699 | -1.1366690801 | 1.810894451 | -0.0912933 |
| 262 | -1.5577371006 | 0.624488149 | -0.0117425 |
| 181 | 0.2229495015 | 0.542347397 | 0.5502321 |
| 140 | -0.4896986175 | -0.317770032 | -1.1222234 |
| 243 | 0.6028803517 | 0.961736329 | -0.5125305 |
| 393 | -0.7918150918 | 1.801467030 | 0.7815183 |
| 347 | -2.0381386734 | -1.193032887 | -0.3801571 |
| 034 | 2.4977753407 | 1.757324321 | -0.9982583 |
| 094 | -0.1974430277 | -0.220404514 | -0.1618838 |
| 345 | 0.3497108092 | 0.880653052 | 0.4979430 |
| 712 | -1.5667439085 | -0.265948553 | 1.4953565 |

  - Observational data (No control over the values of the variables. One might argue that most data today is of this form)
  - Selection bias? (In either case, we must be careful to think about the possibility of selection bias. Do we truly have a random sample, or do some variables in the model directly or indirectly make it more or less likely for the data vector to be included in our data set)



- Assumptions
  - DAG assumption? (Are the data the result of a feedback process? Or can we model the process using a directed acyclic model?)
  - Faithfulness assumption? (Are all independencies due to the causal structure or could some perhaps arise from some 'cancelling of effects'?)
  - No hidden variables assumption? (Usually not justified to assume, since it is often hard to verify the non-existence of a variable we don't know anything about. Fortunately some advanced methods can handle these.)

Bayesian view: assumptions are a form of priors. Not all generative processes are equally likely. We will never be able to be 100% sure about the validity of any of our assumptions, but they might guide us to probable models!



# Model learning

- Given only observational data, what can be done?
  - DAG assumption
  - Linearity assumption?
  - Faithfulness assumption?
  - No hidden variables and no selection ⇒ IC (Pearl) / PC (Spirtes et al) algorithm
  - Linear non-gaussian case: ⇒ LiNGAM
  - Algorithms exist for handling hidden variables & selection: FCI (Fast Causal Inference, not described in this course)

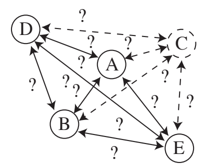


$p(a,b,d,e) = \ldots$

And the causal model is...?

# Estimation of do-probabilities

- When we know the generating DAG from background knowledge, or can infer it from the data, we can combine it with data to answer do-queries (predictions of the results of actions)
- Back-door adjustment
- Front-door adjustment
- Do-algebra (general case)
- Linear regression
- Instrumental variable method

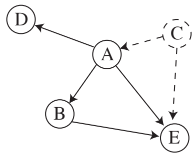


$p(a,b,d,e) = \ldots$

$p(e \mid do(b)) = ?$



# Confidence intervals!

- Resampling for handling limited data
- Allow for multiple possible models (different assumptions, observational equivalence, etc)
- Average over models if needed
- Future decisions: Give predictions (including confidence intervals) for all the different scenarios ('business as usual' vs alternative scenarios)



# 'Take-home message'

- Possible to define causality using (hypothetical) interventions
- The combination of graphs and probabilities provides a natural and versatile language for
  - expressing/communicating causal assumptions
  - formalizing the derivation of causal implications resulting from the combination of assumptions and data
- We have discussed some early and basic results, but...
  - much more is already known
  - there is still much to be discovered (active research field!)



# What to study for the exam?

- Material: lecture slides, required readings, exercises and their solutions
- Focus on the central concepts:
  - DAG models: functional causal model / causal bayes net
  - d-separation and relationship to independencies:
    - d-separations in DAG $\Rightarrow$ independencies in joint distribution
    - faithfulness: no additional independencies in distribution
  - semantics of interventions (do-conditioning)
  - discrete models (with probability tables) vs linear continuous models (with covariances)

- identifiability of causal effects (post-intervention distributions)
- estimating post-intervention distributions with discrete vars
  - by randomized controlled trials
  - by back-door adjustment
  - by front-door adjustment
  - using the 'do-calculus'
- estimating causal effects in linear models for continuous vars
  - using multivariate linear regression
  - using the instrumental variable method
- counterfactuals: what they are, arguments for and against, computing answers in functional causal models assuming persistence of disturbances



- impossibility of model learning without assumptions
- assumptions for model learning: DAG?, faithfulness?, no-hidden-variables?
- d-separation equivalence
- IC algorithm
  - simulating the algorithm
  - conditions for correctness
  - practical aspects



# Exam

- Friday 17.10 at 9.00-12.00 in Exactum B123, no need to enroll for this exam. Exam questions given in both English and Finnish
- Future exam possibilities: 18.11.2008, 6.2.2009, 27.3.2009, and 16.6.2009 (must enroll using department system)
- Bring along: writing utensils and calculator. Nothing else allowed.
- 5 questions/problems, 6 points each, so max total = 30 points
- Some explanations or an essay, but mostly mathematical, similar to exercise problems. The goal is to measure understanding
- Do the best you can on each problem, try not to get stuck on any one of them



# Project work

- Basic idea: Allows you to try to apply the methods of the course in practice using simulated data (where we know the true answer)
- The project is individual: Each student gets his/her own networks and data. The overall problems and structure of the project is the same though
- You are free to use any ready-made software such as Matlab / R / Tetrad / B-course / Lingam / anything else that you choose.
- Alternative projects: Students who are involved in research may discuss with the lecturer regarding applying causal analysis methods on their own data (as an alternative project topic)

- Practical issues:
  - Doing the voluntary project work (in addition to the exam) makes the course worth a total of 6 credits. Doing only the exam gives you 4 cr (2 cu)
  - When doing the project work, the final grade is determined by the combination of exam (30 points) and project work (15 points), both must be acceptably passed. However, there is no risk in attempting the project work as you can always choose to obtain the credits+grade from just the exam if you wish
  - Let me know, by email or on your exam, whether you intend to participate in the project work (so that I know whether to generate data for you), preferrably by 20.10.2008
  - Detailed information, project description, and data sets will be available on 24.10.2008 at the latest (possibly earlier)
  - The DL will be 30.11.2008



  - Help will be available from the lecturer during the project (until 11th of November in person, after that only by email because of travels).
  - Feedback/grading explained individually after the submission deadline



# Project work problems...

The full details are not yet decided, but the following types of problems are likely:

- Estimation of causal effects: I provide you with some data and a number of possible different DAG graphs and your task is to:
  - Evaluate the different DAGs in terms of fitting the data
  - Estimate some specified causal effects using the combination of do-calculus and the observed data
  - Make some decisions so as to try to maximize a given utility function



- Counterfactuals: I provide you with a parametrized functional causal model and your task is to:
  - Generate some data according to the model
  - Calculate some counterfactual probabilities
  - Estimate the counterfactuals from the observed data

- Model learning: I have generated several datasets using various models and your task is to:
  - Investigate the data and infer the generating causal model (using some of the methods presented in the course)
  - Report your remaining uncertainty regarding the correct model

- Combining the results of a number of experiments: I have generated some data using a single model but in a number of different experimental conditions, and your task is to:
  - join all the available information and infer the full causal model as well as possible
  - suggest future experiments to perform which optimally distinguish between the remaining model candidates



# Feedback

- Official channel for course feedback: https://ilmo.cs.helsinki.fi/kurssit/servlet/Valinta?kieli=en
- If you want to you can wait until after the exam and final grading, and for those of you who are doing the project work you may want to wait until it is finished (I will email you and remind you to fill out the form, don't worry!)
- Please also feel free to send comments directly by email.