Unit Selection Based on Counterfactual Logic

Causal Network: Learning and Inference - Final Exam

Marco Gigli
Greta Greco
Dagmawi Tegegn

Ang Li, Judea. Pearl. (2019). Unit Selection Based on Counterfactual Logic.

Outline

- What is unit selection
- Motivations
- State of art + Structural causal model
- Counterfactual expression
- Example churn
- Example advertisement
- Differences with statistical approach
- Conclusions

Motivating example

Phone company wants to

- identify customers likely to discontinue their services
- offer discount to **most promising** group

Response types:

- **Compliers** renew subscription if encouraged, otherwise they do not
- **Always takers** renew subscription anyway
- **Never takers** do not renew subscription anyway
- **Defiers** *do not* renew subscription if encouraged, otherwise they do
    - Reminded they pay for a service they no longer want
    - Feel that discount cheapens service
    - Are annoyed by the discount claim process

Expected benefit

- **Complier**: 140$ (profit from renewal) - 40$ (discount) = **100$**
- **Always taker**: - 40$ (discount) - 20$ (discount offering triggers need of additional discounts in the future) = **- 60$**
- **Never taker**: **0$** (no profit and no discount claimed)
- **Defier**: **-140$** (customer lost)

Expected benefit for characteristic c:

100 P(complier | c) - 60 P(always taker | c) - 0 P(never taker | c) - 140 P(defier | c)

Defined in **counterfactual terms**! Example: to distinguish an always taker from a complier, we would have to observe response both with and without discount.

Unit Selection Problem

Aim: identifying individuals most likely to **show a desired response** pattern if encouraged, and conversely if not.

More precisely: finding the characteristics c that **maximize** the percentage of compliers while **minimizing** the percentage of other classes.

Find *c* that maximizes:

$$f(\beta, \gamma, \theta, \delta) = \beta\, P(\text{complier} \mid c) + \gamma\, P(\text{always-taker} \mid c) + \theta\, P(\text{never-taker} \mid c) + \delta\, P(\text{defier} \mid c)$$

State of the art: Observational

P(Y=y|c, X = x)

Machine Learning models can be trained on past observational data:

- Customer churn models
- Click-through-rate models

This approach does not answer **causal** questions:

- Did the discount **cause** retention?
- Did the advertisement **cause** the click?

State of the art:
Randomized Controlled Trial

$$P(Y = y|c, do(X = x))$$

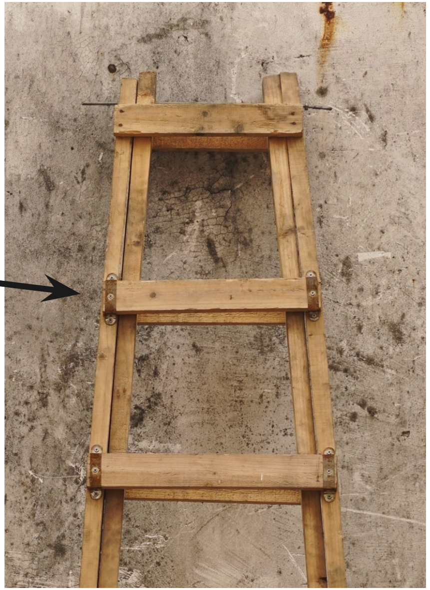

- Users randomly split in control and treatment
- Treatment group receives encouragement, control group doesn't

RCT can answer causal questions, but cannot answer **counterfactual** questions

Proposed approach: Counterfactual formulation

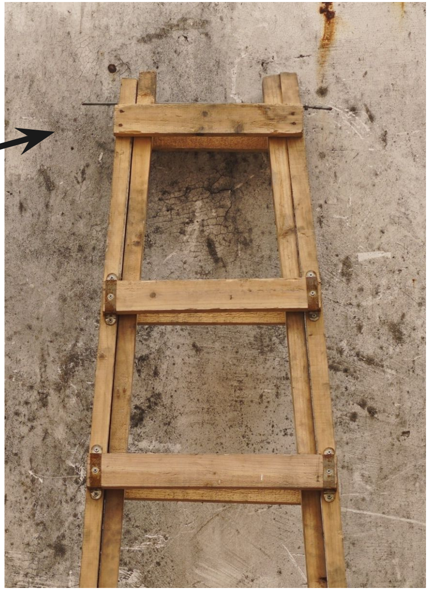

$$P(Y(x) = y \mid Y(x') = y', c)$$

The desired response pattern is **not observed directly** but rather is **defined counterfactually** in terms of what the individual would do under hypothetical unrealized conditions.

$$f(\beta, \gamma, \theta, \delta) = \beta\, P(\text{complier} \mid c) + \gamma\, P(\text{always-taker} \mid c) + \theta\, P(\text{never-taker} \mid c) + \delta\, P(\text{defier} \mid c)$$

$$= \beta\, P(R(a) = r, R(a') = r' \mid c) + \gamma\, P(R(a) = r, R(a') = r \mid c) + \theta\, P(R(a) = r', R(a') = r' \mid c) + \delta\, P(R(a) = r', R(a') = r \mid c)$$

Main results

- Upper and lower **bounds** are given for the objective function $f(\beta, \gamma, \theta, \delta)$, depending only on **experimental** and **observational** data. The bounds do not require specifying a Structural Causal Model.

- In general, without a SCM, $f(\beta, \gamma, \theta, \delta)$ is not identifiable. It is **identifiable** under additional assumptions:

  - **Monotonicity**: no defiers
  - **Gain equality**, $\beta + \delta = \gamma + \theta$: benefit(complier) + benefit(defier) = benefit(always taker) + benefit(never taker)

- Under monotonicity or gain equality, $f(\beta, \gamma, \theta, \delta)$ takes the **same form**:

$$f(\beta, \gamma, \theta, \delta) = (\beta - \theta)P(y \mid c, do(x)) + (\gamma - \beta) P(y \mid c, do(x'))$$

- Experiments show that the **bound midpoint** can be effectively used when $f(\beta, \gamma, \theta, \delta)$ is not identifiable

Bounds

The objective function **f(β, γ, θ, δ)** is bounded as follows:

$$\max\{p_1, p_2, p_3, p_4\} \leq f \leq \min\{p_5, p_6, p_7, p_8\}\ if\ \alpha < 0.5$$

$$\max\{p_5, p_6, p_7, p_8\} \leq f \leq \min\{p_1, p_2, p_3, p_4\}\ if\ \alpha > 0.5$$

where

$$\alpha = \frac{\beta - \gamma - \theta}{\beta - \gamma - \theta - \delta}$$

$$p_1 = (\beta - \theta)P(y_x|z) + \delta P(y_{x'}|z) + \theta P(y'_{x'}|z)$$

$$p_2 = \gamma P(y_x|z) + \delta P(y'_x|z) + (\beta - \gamma)P(y'_{x'}|z)$$

$$p_3 = (\gamma - \delta)P(y_x|z) + \delta P(y_{x'}|z) + \theta P(y'_{x'}|z) + (\beta - \gamma - \theta + \delta)P(y,x|z)[P + (y',x'|z)]$$

$$p_4 = (\beta - \theta)P(y_x|z) - (\beta - \gamma - \theta)P(y_{x'}|z) + \theta P(y'_{x'}|z) + (\beta - \gamma - \theta + \delta)[P(y,x'|z) + P(y,x|z)]$$

$$p_5 = (\gamma - \delta)P(y_x|z) + \delta P(y'_x|z) + \theta P(y''_x|z)$$

$$p_6 = (\beta - \theta)P(y_x|z) - (\beta - \gamma - \theta)P(y'_x|z) + \theta P(y'_x|z)$$

$$p_7 = (\gamma - \delta)P(y_x|z) - |(\beta - \gamma - \theta)P(y'_x|z) + \theta P(y''_x|z) + (\beta - \gamma - \theta + \delta)P(y|z)$$

$$p_8 = (\beta - \theta)P(y_x|z) + \delta P(y'_x|z) + \theta P(y''_x|z) - (\beta - \gamma - \theta + \delta)P(y|z)$$

Churn Management Experiment

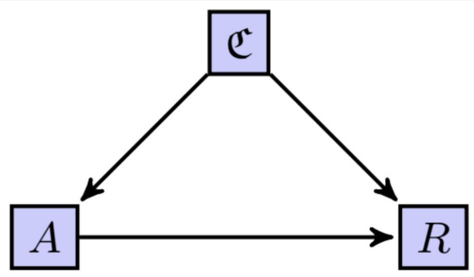

Causal graph for the customer selection model.

Problem

Predicting which customers are about to churn, but are likely to change their minds if enticed toward retention.

(Unknown) distribution of response types

	Complier	Always-taker	Never-taker	Defier	Benefit
Group 1	60%	28%	2%	10%	0.22
Group 2	50%	3%	27%	20%	0.27

Gains	1	-1	0	-1

Randomized Controlled Trial

Randomly select 700 customers from each group and offer the special renewal deal to 350 customers in each group.

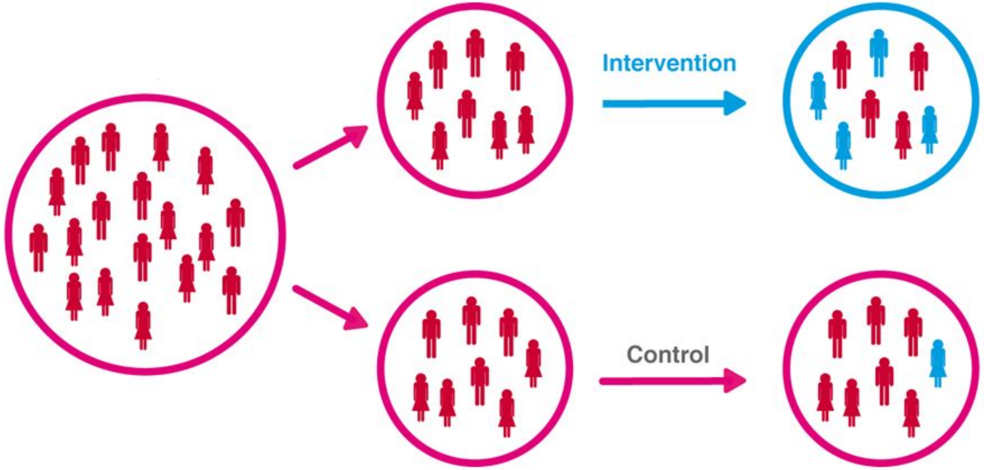

RCT data

		$do(a)$	$do(a')$
Group 1	r	308	133
	r'	42	217
Group 2	r	186	81
	r'	164	269

Observational data

$$P(r|c_1) = 0.3$$
$$P(r|c_2) = 0.1$$

Objective functions

		Group 1	Group 2
HEURISTIC	Average Treatment Effect: $P(r \mid c, do(a)) - P(r \mid c, do(a'))$ →	**0.5**	0.3
	Assuming monotonicity: $P(r \mid c, do(a)) - 2 \cdot P(r \mid c, do(a'))$ →	**0.12**	0.07
COUNTERFACTUAL	$f(1, -1, 0, 1) = 1 \cdot P(r_a, r'_{a'} \mid c) + (-1) \cdot P(r_a, r_{a'} \mid c) + 0 \cdot P(r'_{a'}, r'_{a'} \mid c) + (-1) \cdot P(r'_a, r_{a'} \mid c)$ → Bound midpoint	0.22	**0.25** ✓

The **heuristic** approach comprises of two objective functions, respectively based on controlled experiment and weighted controlled experiment.

The **counterfactual** objective function is the proposed approach: midpoint of bounds is used.

Advertisement Recommendations Experiment

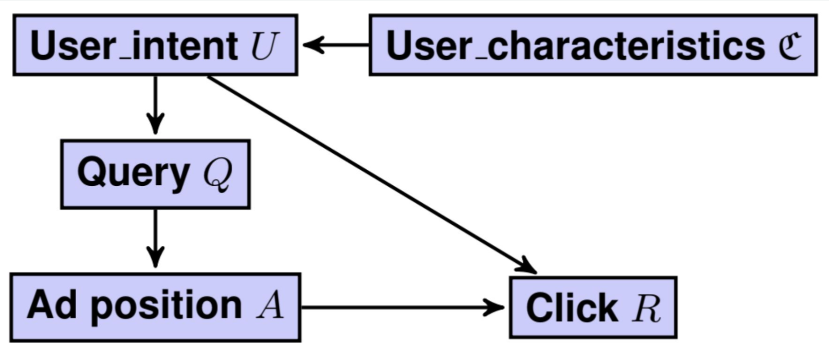

Causal graph for the user selection model.

Problem

Identifying users who are likely to click on a given advertisement if (and only if) the advertisement is placed in top position

(Unknown) distribution of response types

	Complier	Always-taker	Never-taker	Defier	Benefit
Group 1	70%	6%	4%	20%	0.20
Group 2	60%	20%	15%	5%	0.15
Benefit	1	-1	-1	-2	

		RCT		Observational	
		$do(a)$	$do(a')$	a	a'
Group 1	r	266	91	20	67
	r'	84	259	143	470
Group 2	r	280	87	226	30
	r'	70	263	10	434

Objective functions

	Objective function		Group 1	Group 2
HEURISTIC	Average Treatment Effect: $P(r \mid c, do(a)) - P(r \mid c, do(a'))$	→	0.5	**0.55**
	Naive observational: $P(r \mid c, a) - P(r \mid c, a')$	→	0	**0.89**
COUNTERFACTUAL	$f(1, -1, 0, 1) = 1 \cdot P(r_a, r'_{a'} \mid c) + (-1) \cdot P(r_a, r_{a'} \mid c) + (-1) \cdot P(r'_{a'}, r'_{a'} \mid c) + (-2) \cdot P(r'_a, r_{a'} \mid c)$	Bound midpoint →	**0.17** ✓	0.13

The **heuristic** approach comprises of two objective functions, respectively based on controlled experiment and observational data.

The **counterfactual** objective function is the proposed approach.

Benefit comparison

The benefit of group 1 letting **δ vary**, and keeping β, γ, and θ fixed, as calculated from different objective functions.

The proposed objective is the closest to the real one.

The real benefit lies between the bounds.

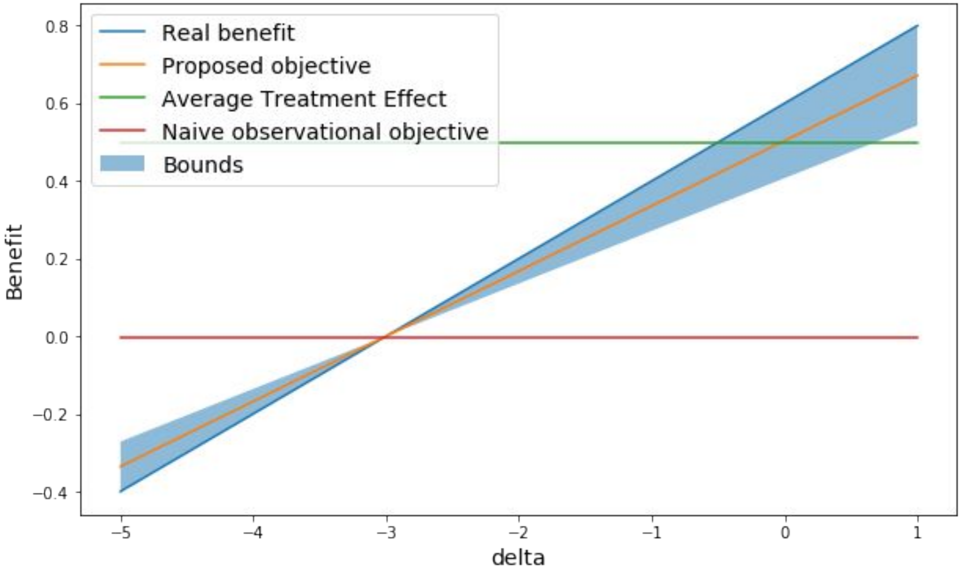

Exploiting Causal Graphs

Derived **bounds** require

- Observational data:
  - P(r, a | c), P(r', a | c), …
  - P(r | c)
- Experimental data:
  - P(r | c, do(a))
  - P(r | c, do(a'))

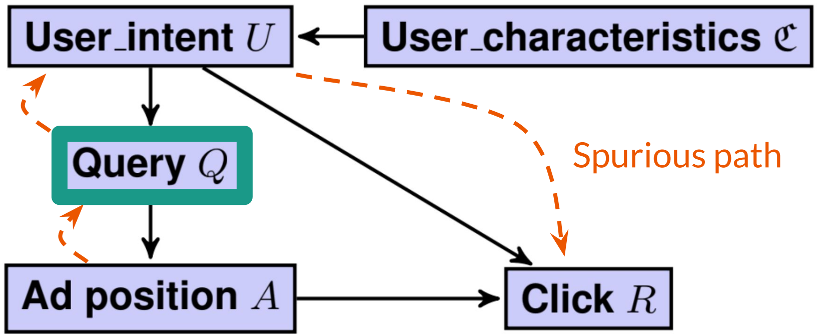

If we have a causal graph and observational data P(R, A, Q) and a set of variables satisfying the **backdoor criterion**, we don't need an experiment.

Backdoor adjustment formula

$$P\left[R=r|do(A=a)\right]=\sum_{q}P\left[R=r|A=a,Q=q\right]P(Q=q)$$

Conclusions

- **Unit Selection Problem** properly treated in counterfactual setting
- Counterfactual **objective** f(β, γ, θ, δ)
- Identifiable **bounds** for f(β, γ, θ, δ): need only observational and experimental data
- Identifiable objective with additional **assumptions**
- Bound **midpoint** effective in practical settings
- Bounds do not need a **causal graph**
- If a causal graph is known and **backdoor criterion** can be applied, observational data suffice

Questions?