Abstraction and generalization in causal learning

Christopher Lucas

Causal learning

Causal learning

Prediction:

- What happens if I push this button?

Explanation and diagnosis:

- What caused the noise downstairs?
- Why do some objects float in water?

Planning and achieving goals:

- How can I stop killing my plants?

Causal learning

Scientific discovery:

- What causes disease?
- Discovering new planets
- Genetic inheritance

Causal learning

It's difficult!

- Many variables
- Coincidence/spurious covariation
- Causes and effects can relate in many ways:

  A causes B,

  B causes A,

  common causes,

  intereractions, …

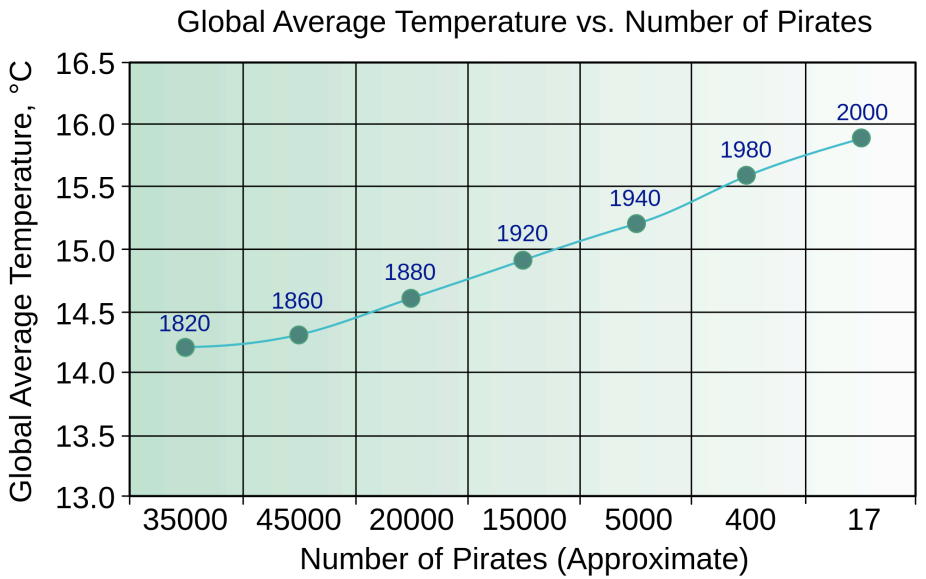

(wikimedia commons: RedAndr)

Inductive biases

"Squad Helps Dog Bite Victim"

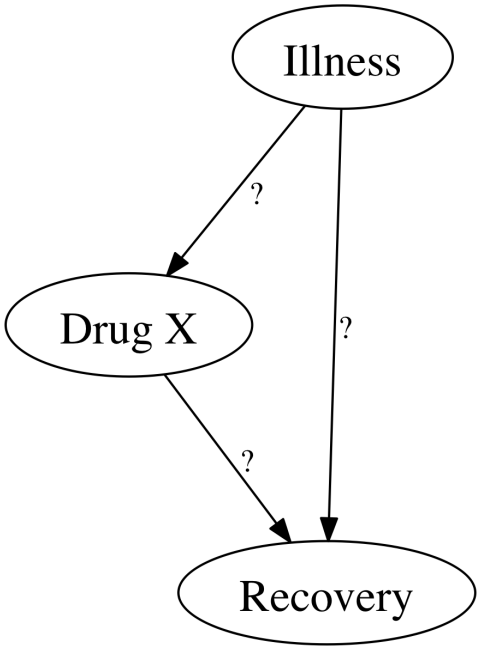

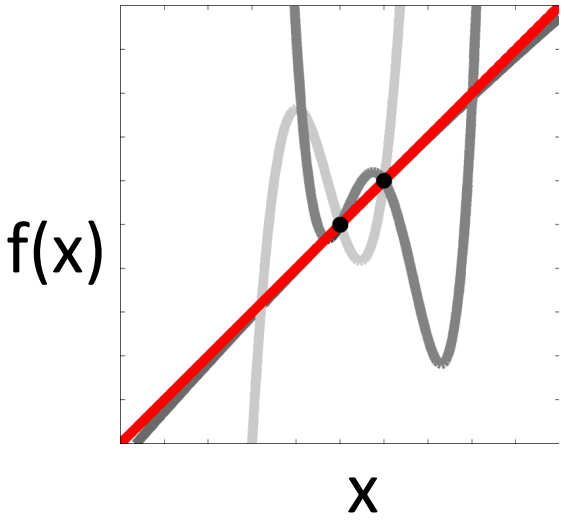

(Cooper, 1980; Geman Bienenstock & Doursat, 1992)

Questions

What representations and inductive biases allow us to generalize?

Where do they come from?

Questions

How do we discover new concepts or abstractions, and refine the ones we have?

- Discovering and reasoning about hidden variables
- What *kinds* of causal relationships are likely

Methods

Predictions from hierarchical Bayesian models, e.g.:

- functional causal models;
- mixtures of Gaussian process experts;
- probabilistic grammars/programs

(Pearl, 2000; Rasmussen & Gharamani, 2002; Meeds & Osindero, 2006; Goodman et al. 2008)

Methods

Experiments: surveys, demonstrations, classification and prediction tasks, and games

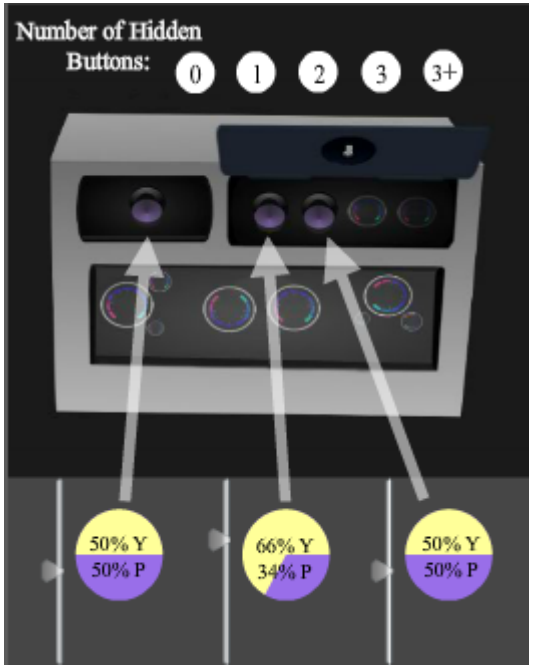

Example experiment

Can people learn about forms of relationships using one set of variables and use that knowledge when reasoning about new variables?

Example experiment

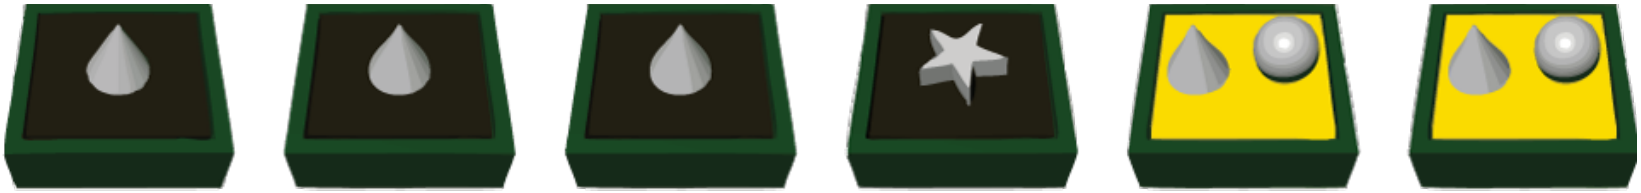

D D D E D,F D,F

Which objects are blickets?

D? E? F?

Example experiment

Training: See unlabelled event data, identify causal structure behind events.

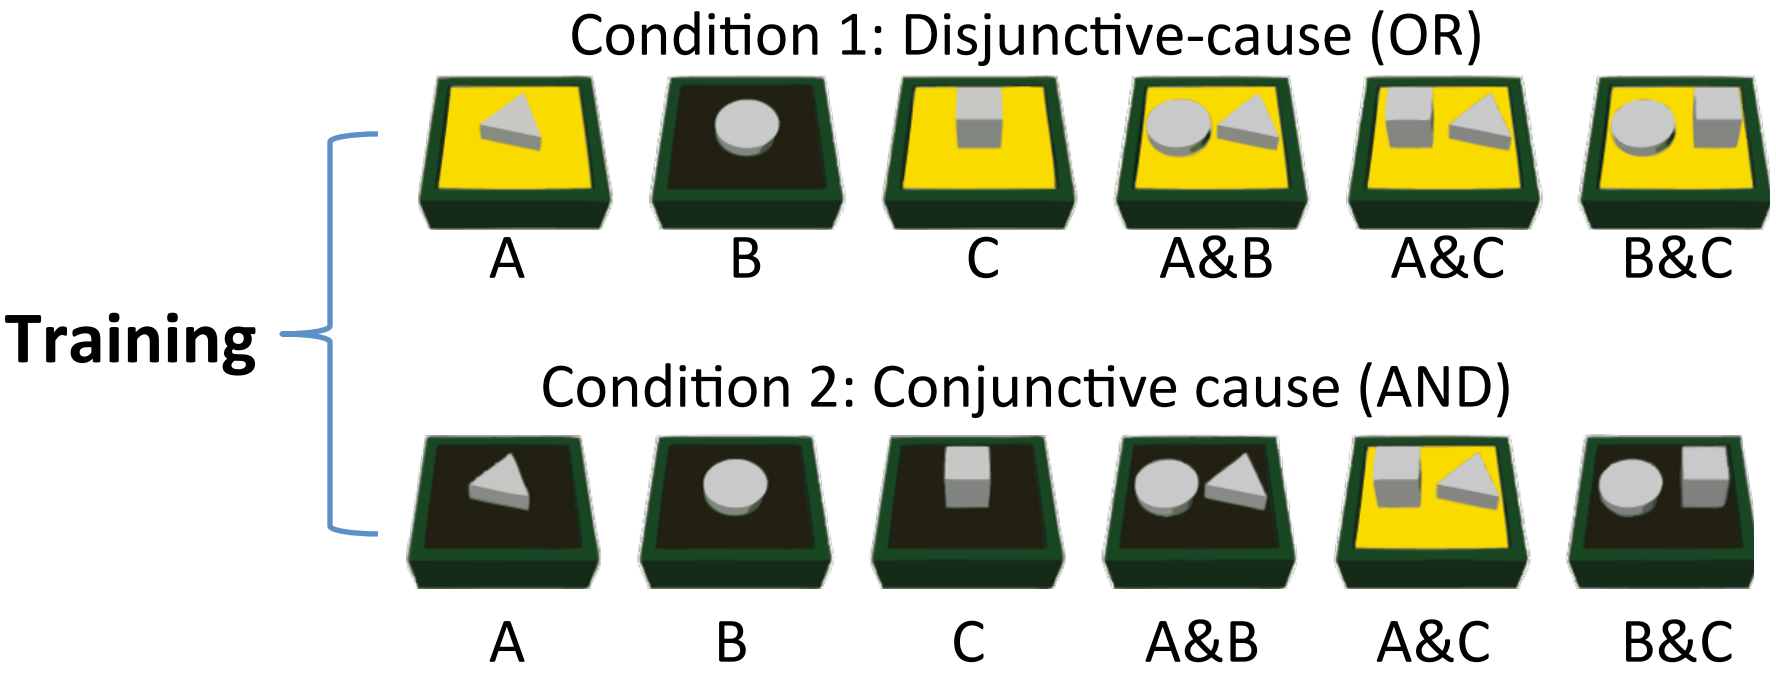

(Lucas et al., , 2014)

Example experiment

Test: See ambiguous data involving new objects. Infer causal structure by generalizing from training.

Test

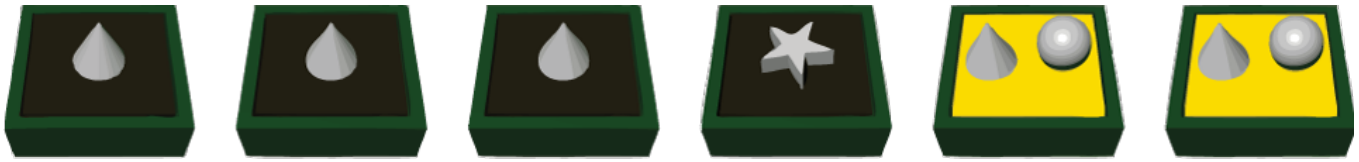

D D D E D&F D&F

Developmental differences

If our experiences today shape our interpretation of tomorrow's data:

- Generalizations should change over development
- Children will be more sensitive to atypical data than adults

(Lucas et al., 2014)

Results

- Both adults and children generalize from training.
- Children show greater flexibility.
- Adults have stronger *a priori* commitments.

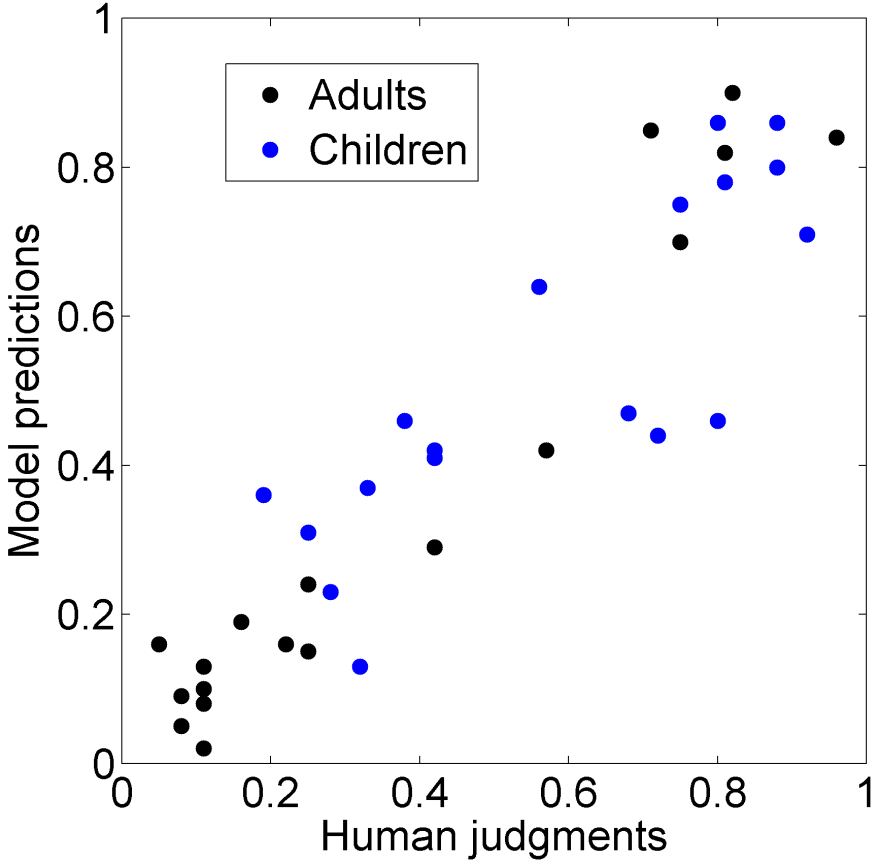

Example: continuous relationships

E.g., causes and effects of climate change

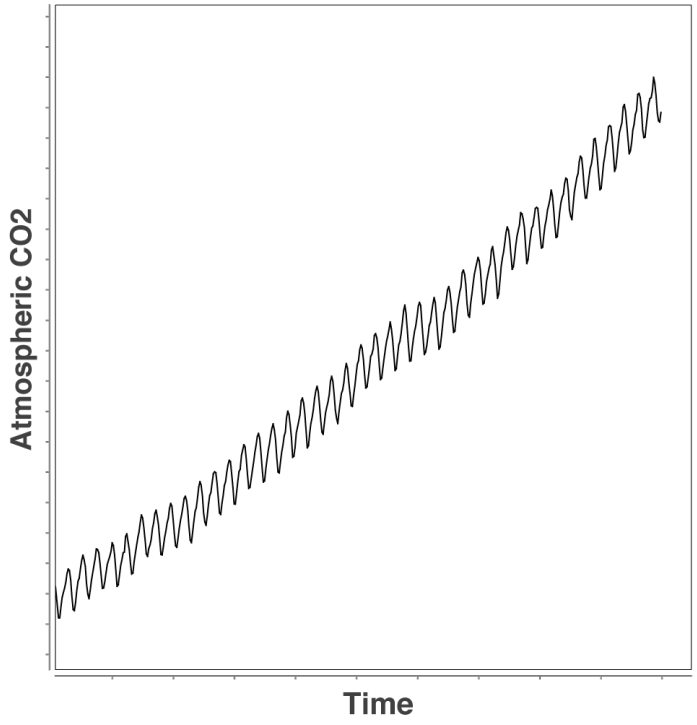

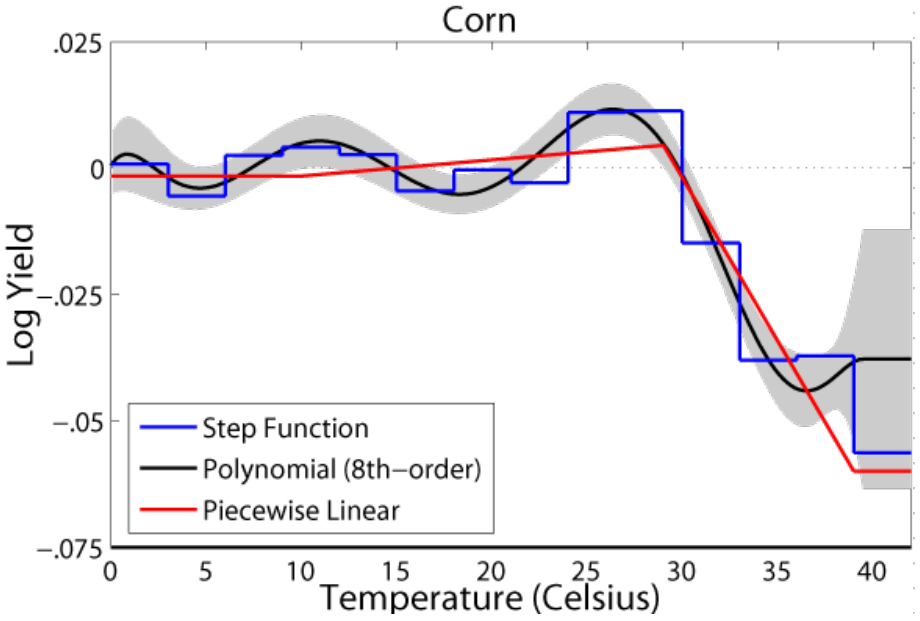

(Schenker & Roberts, 2009)

A Bayesian model explains diverse phenomena

Extrapolation

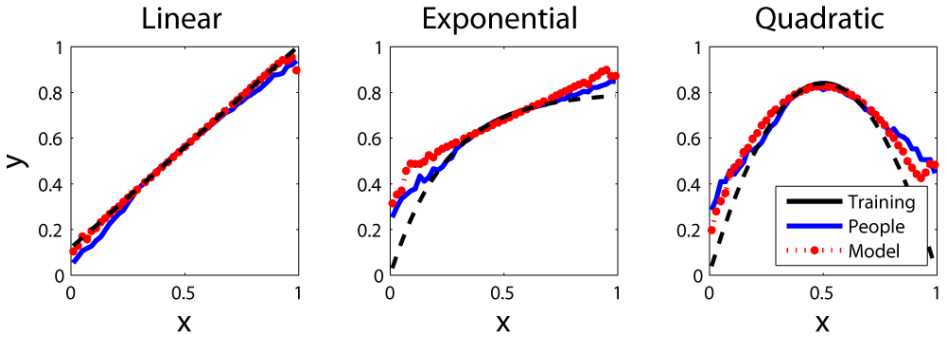

Individual differences

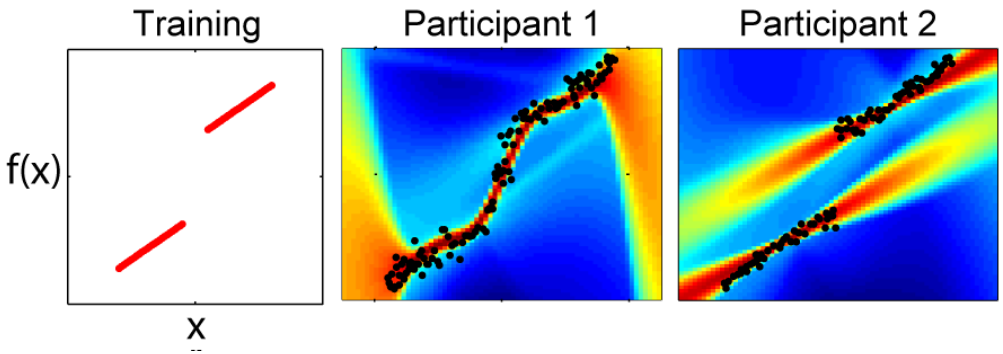

Interpolation performance (13 experiments)

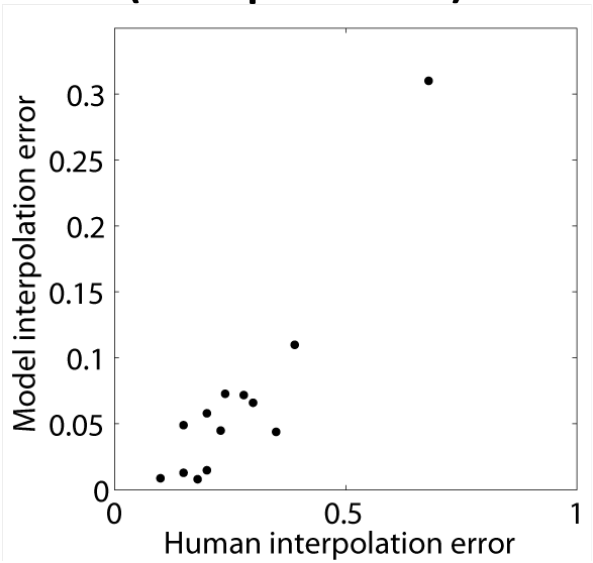

...and more:
Iterated learning;
Knowledge partitioning

(Delosh, Busemeyer, & McDaniel, 1997; Kalish, Lewandowsky, & Kruschke, 2004; Byun, 1995; Kalish, Griffiths, & Lewandowsky, 2007)

Extrapolating in higher-dimensional spaces

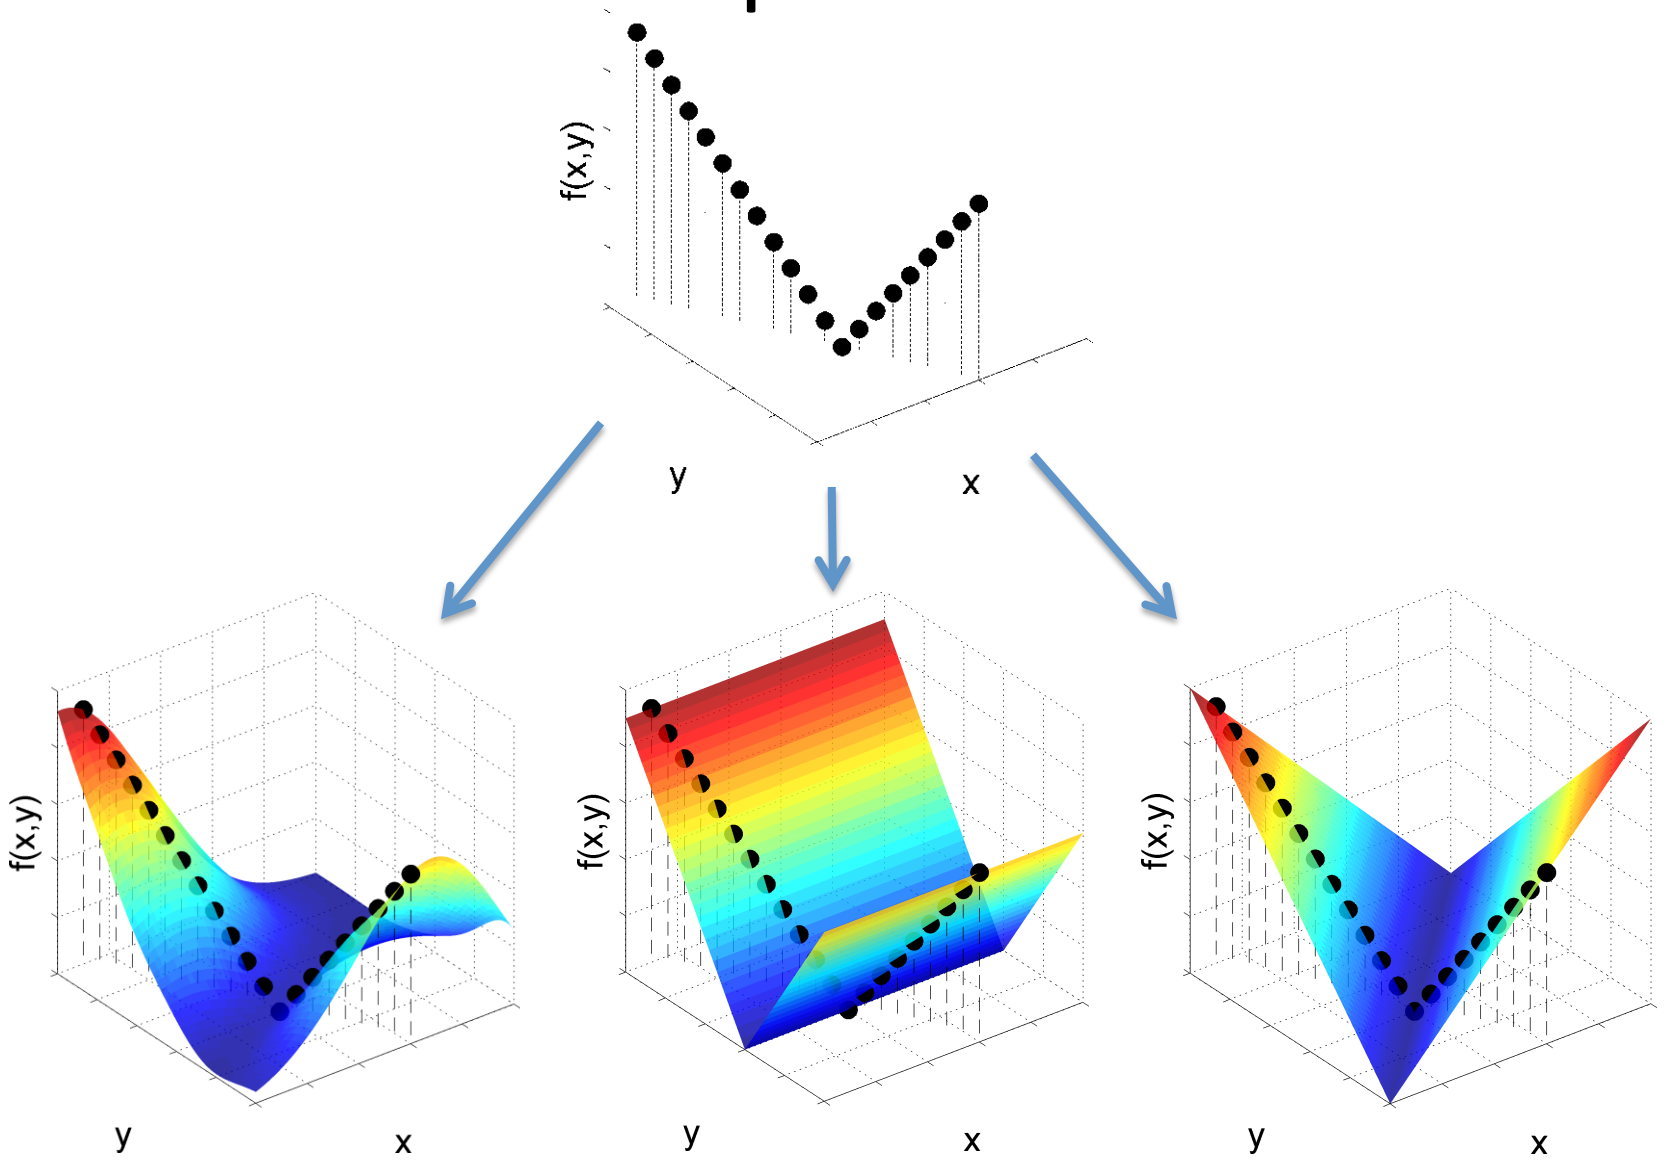

Example

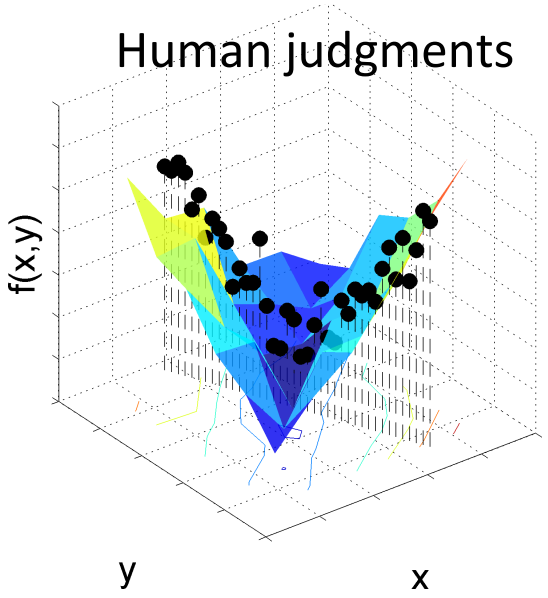

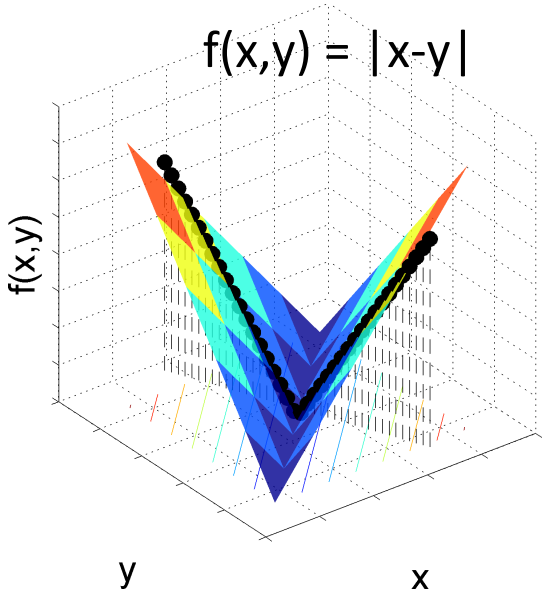

Example

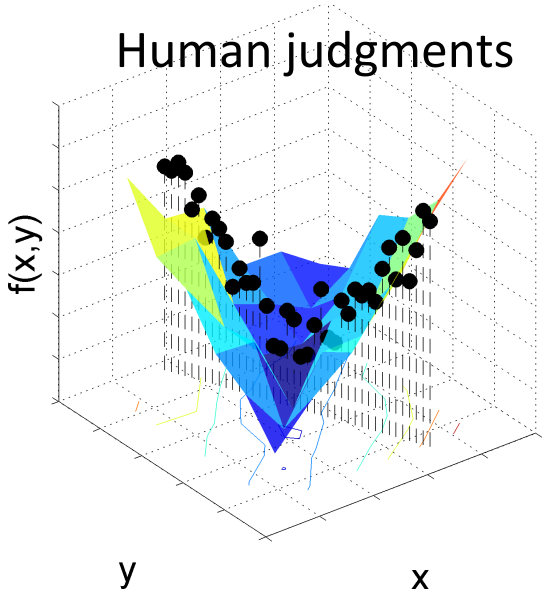

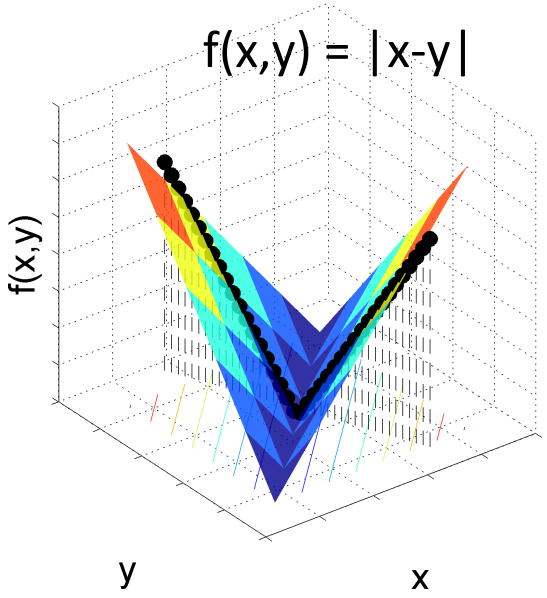

Rotated

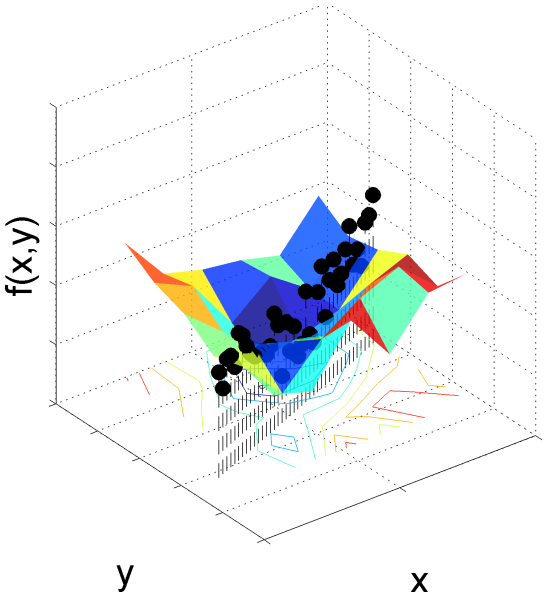

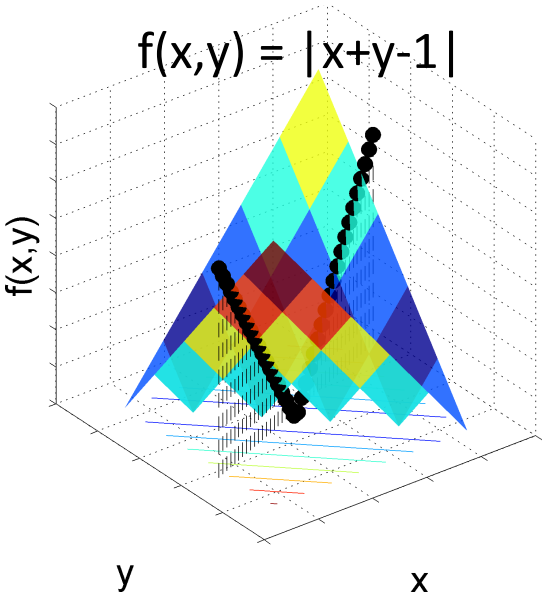

Other projects

- Compositions of programs as priors
- Preference understanding
- Using statistical “concidences” to discover hidden causes
- Temporal information and causal inference
- Counterfactual reasoning