# Computational Cognitive Science

Lecture 13: Active Learning

Frank Keller

2 November 2026

School of Informatics
University of Edinburgh
keller@inf.ed.ac.uk

Slide credits: Chris Lucas, Benjamin Peters

# Overview

What is Active Learning?

Entropy and Information

Information Gain

Active Causal Learning

Recommended reading: Steyvers et al. (2003).

# What is Active Learning?

# Active learning

In many examples so far, we have treated people as passive observers. In reality, people often choose what information to obtain:

- Visual judgments: look at more informative parts of the scene
- Categorization: choose which examples to get labels for
- The number game: ask whether particular numbers belong to the concept
- Causal learning: intervene and design experiments

Active learning means choosing actions or information in order to learn.

# Active learning

To count as active learning, there must be:

1. selection of action or information that
2. serves some learning goal

# Active learning

Some active learning is unconscious, e.g., gaze during reading or viewing images:



# Entropy and Information

# Formalizing active learning

So active learning means acting or selecting information to learn more.

But what does it mean to learn more or less; how can we quantify learning?

Some intuitions:

- Learning reduces our uncertainty
- Learning changes our beliefs

# Reducing uncertainty

Let's try to express this intuition mathematically:

- If we're certain, we don't have much left to learn and don't expect to be surprised
- Let's define surprisal associated with event $x$ as $\log(1/P(x)) = -\log(P(x))$

For example:

| event $x$ | $P(x)$ | surprisal (bits) |
|---|---|---|
| drop object and it falls to the ground | $P(\text{object falls}) = 1$ | 0 |
| flip the switch and light goes on | $P(\text{light goes on}) = .99$ | .01 |
| flip a fair coin, it comes up heads | $P(\text{heads}) = .5$ | 1.0 |
| win the lottery (UK) | $P(\text{win}) = \frac{1}{45 \cdot 10^6}$ | 25 |
| drop object and it floats | $P(\text{object floats}) = 0$ | $\infty$ |

# Entropy

We can formalize uncertainty as expected surprisal:

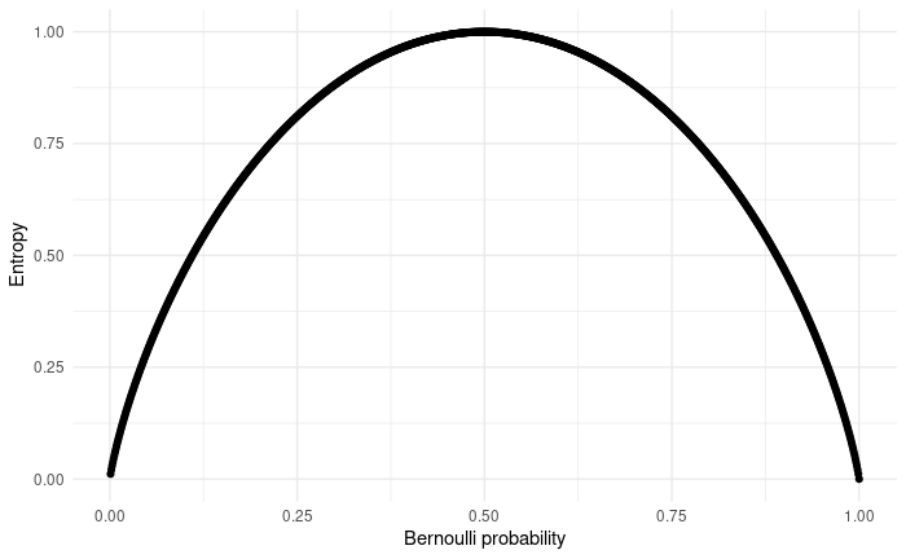
$$\begin{aligned} H(X) &= \mathbb{E}_{P(x)}[-\log P(x)] \\ &= -\sum_{x \in \mathcal{X}} P(x) \log P(x) \end{aligned}$$

This quantity is called entropy.

- High entropy: we are uncertain about which outcome will occur.
- Low entropy: one outcome is already highly predictable.

With log base 2, entropy is measured in bits.

# Entropy


1.00
0.75
0.50
0.25
0.00
Entropy
0.00
0.25
0.50
0.75
1.00
Bernoulli probability

# Entropy and communication: Example

Alice randomly chooses a candy from a bag:

$$P(A) = \tfrac{1}{2}, \qquad P(B) = P(C) = P(D) = P(E) = \tfrac{1}{8}$$

She wants to tell us which candy she has by blinking (left $= 0$, right $= 1$).

A sensible code gives the common outcome a short message:

| Candy | A | B | C | D | E |
|---|---|---|---|---|---|
| Code | 0 | 111 | 110 | 101 | 100 |

How many bits will she need on average?

# Entropy and communication: Solution

The surprisals are

$$-\log_2(1/2, 1/8, 1/8, 1/8, 1/8) = (1, 3, 3, 3, 3)$$

So the expected surprisal is

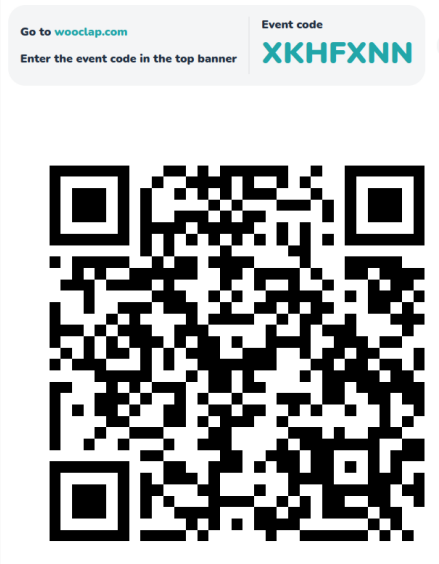
$$H(X) = \tfrac{1}{2}(1) + 4\left(\tfrac{1}{8}\right)(3) = 2 \text{ bits}$$

This matches the code:

$$\tfrac{1}{2} \cdot 1 + \tfrac{1}{2} \cdot 3 = 2 \text{ bits}$$

Entropy is the average amount of information needed to identify the outcome.

# Time for a short quiz on WooClap!



# Information Gain

# Entropy and active learning

An informative action is one that we expect to reduce uncertainty about what we want to learn.

This can be expressed as information gain:

$$H(Y) - \mathbb{E}_{P(d|a)}\left[H(Y|a,d)\right]$$

information gain = uncertainty before – expected uncertainty after

- $Y$ is what we care about
- $a$ is our action
- $d$ is the unknown result of our action

# Entropy and active learning

How do we maximize our information gain?

$$H(Y) - \mathbb{E}_{P(d|a)}\left[H(Y|a,d)\right]$$

$H(Y)$ is fixed when we choose the action; want to minimize the expected posterior entropy $H(Y|a,d)H(Y|a,d)$:

$$\mathbb{E}_{P(d|a)}\left[H(Y|a,d)\right] = \sum_{d\in\mathcal{D}} P(d|a)\, H(Y|a,d)$$

That is, we want to pick actions that are likely to leave us with little uncertainty.

# Focusing on uncertain events

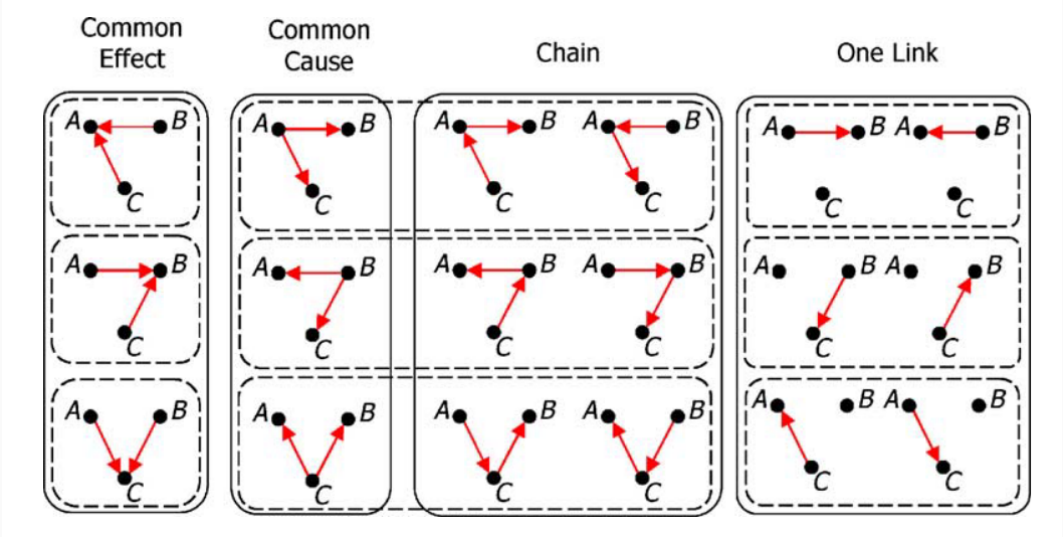
We learn little from experiments where we know what the outcome will be.
What if we focus on experiments where we don't know what will happen?

An outcome can be unpredictable because it is informative, or because it is noisy.

- "maximum entropy sampling" – sometimes very useful, but misleading if some observations are inherently noisy
- sometimes a pitfall for new scientists!

# Active Causal Learning

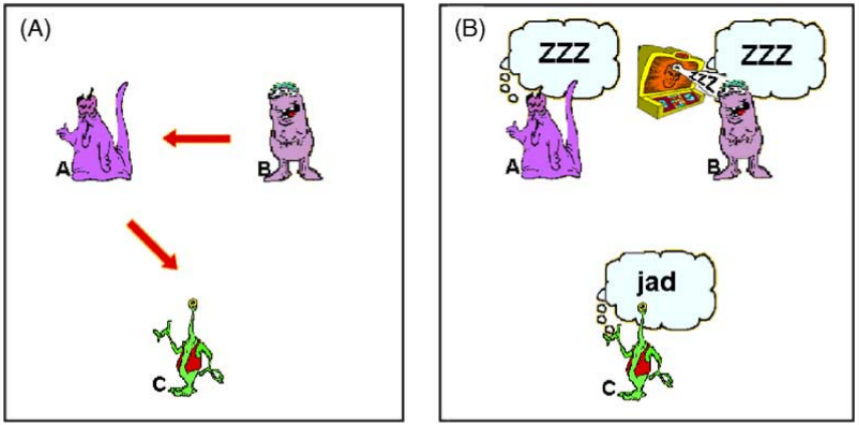
# Reminder: causal structures




Different causal graphs may induce the same joint distribution (dashed lines).

# Example: Alien mind reading

Steyvers et al. (2003) asked whether people choose causal interventions to learn efficiently, using an information-gain approach.



# Example: Alien mind reading

Participants saw 18 kinds of acyclic causal graphs and made causal structure judgments based on:

- Observations
- Interventions with a brain zapper

The causal relationships were stochastic – aliens could fail to read other minds.

# Example: Alien mind reading

Steyvers et al. compared active-learning models that considered different sets of competing hypotheses:

- Rational identification: Which of all possible graphs is correct?
- Rational test 1: Does a working hypothesis beat a null hypothesis (independence)?
- Rational test 2: Does a working hypothesis beat a simpler model with one fewer edge?

Identification: Which of all possible graphs is correct?

Testing: Does my current hypothesis beat a specific alternative?

## Steyvers et al.: Main result

What do people's interventions look like?

- People were not best described as globally identifying the correct graph from all possible structures.
- Their choices were better captured by rational test models: select an intervention that discriminates a current hypothesis from a focused alternative.
- There was individual variation: some participants behaved more like rational identification.

Human active causal learning appears to be goal-directed and informative, but often more like local hypothesis testing than exhaustive optimal search.

Caveat: the task involved only three variables and 18 structures, and these models optimize the information gain from the next observation (a myopic policy).

## Summary

- Active learning means selecting actions or information in order to achieve a learning goal.
- Surprisal measures how unexpected an outcome is; entropy is expected surprisal and quantifies uncertainty.
- A rational learner can choose actions to maximize expected information gain: the expected reduction in entropy after observing the result.
- Simply targeting uncertain outcomes is not always optimal: uncertainty can reflect noise rather than useful information.
- In causal learning, interventions can be selected to distinguish competing causal structures.
- Steyvers et al. (2003) found that people's interventions were better described by focused rational tests than by global rational identification, with substantial individual variation.

# References

Steyvers, Mark, Joshua B. Tenenbaum, Eric-Jan Wagenmakers, and Ben Blum. 2003. Inferring causal networks from observations and interventions. *Cognitive Science* 27(3):453–489.