Learning as Hypothesis Testing

Learning Conditional and Probabilistic Information

Jonathan Vandenburgh

Northwestern University
Department of Philosophy

Progic 2021

Case: Conditional Learning

- Suppose your friend is rolling what you believe to be a fair die and you learn ‘If the die lands odd, it will land 3’
- How do you update your beliefs to incorporate the new information?
- You know $P(3|odd) = 1$. But what should posterior distribution be? i.e., what is $P(4)$?

Bayesian Learning

- This problem falls under Bayesian learning: given prior distribution P over a domain Ω and a constraint A, we want a posterior distribution P_A
- Bayesian learning says posterior is given by conditionalization: for $\omega \in \Omega$, $P_A(\omega) = \frac{P(\omega \wedge A)}{P(A)}$
- But this is only defined for constraints measurable w.r.t. the sample space, where $A \subseteq \Omega$: e.g., ‘The die lands odd’ corresponds to $A = \{1, 3, 5\}$
- Conditional constraints are not measurable: there is no subset of $\{1, ..., 6\}$ corresponding to ‘If the die lands odd, it lands 3’

Relative Information Minimization

- Common proposal: extend Bayesian learning by choosing posterior distribution which incorporates constraint but changes as little as possible about prior beliefs
- Change can be measured by relative entropy or Kullback-Leibler divergence, $\sum_{\omega\in\Omega} P_A(\omega)\log(\frac{P_A(\omega)}{P(\omega)})$
- Similar intuition to AGM model: change beliefs as little as necessary

Relative Information Minimization: Problems

- Example: when learning that $P(3|odd) = 1$, set $P(1) = P(5) = 0$ and keep other probabilities proportional, so for $\omega = 2, 3, 4, 6$, $P(\omega) = \frac{1}{4}$
- This is equivalent to updating by the material conditional $odd \supset 3$
- But there is no way to realize this posterior distribution with a six-sided die
- This proposal runs into other problems: sundowners problem (Douven & Romeijn 2011) and Judy Benjamin problem (Van Fraassen 1981)

Theory-level Constraints

- We can think of a prior distribution as a theory or hypothesis: (Ω, P) tells us what outcomes are possible and how likely each outcome is
- Bayesian learning applies to constraints on outcomes: learning $A \subseteq \Omega$ tells us the outcome is in A
- Conditional constraints are constraints on the entire theory (Ω, P): $A \rightarrow B$ true in (Ω, P) iff all A-worlds in Ω are B-worlds
- Ex: 'If the die lands odd, it lands 3' is false on theory that the die is fair, but true on the theory that all faces of the die are 3

Theory-level Constraints

- Other theory-level constraints include: $P(A) = p$; $P(B|A) = p$; '*A* might/must happen'
- Such constraints are sometimes called 'tests' of an information state (Veltman 1996)
- When a theory-level constraint is false in (Ω, π), it is *incompatible* with the theory
- This differs from outcome-level constraints, which are always compatible with a theory

Hypothesis Testing: Intuitions

- Learning information incompatible with a hypothesis requires (sometimes radical) theory revision
- Children test and radically revise concepts (i.e., learning clocks are not alive) (Carey & Spelke 1994, Gopnik 1996)
- Evidence that adults test and radically change theories of opponent behavior during strategic interaction (Salmon 2004, Young 2004)
- Kuhn's theory of paradigm shift: epistemic communities abandon and replace theories when sufficient incompatible evidence is discovered

Learning by Hypothesis Testing

- Constraint C incompatible with hypothesis $h = (\Omega, P)$ triggers 'theory change,' or choice of new theory from hypotheses consistent with C, $\mathcal{H}_C = \{(\Omega_i, P_i)\}$
- New hypotheses can differ substantially from prior beliefs, and there can be multiple elements in $\mathcal{H}_C$
- When multiple options, can use 'prior over priors,' or distribution π over $\mathcal{H}_C$

Determining Hypothesis Space

- To make reasonable predictions, need way to narrow down hypothesis space
- Method 1: parameterize hypotheses, i.e., learning height distribution from finite sample by parameterizing with normal distributions (μ, σ)
- Method 2: associate theories with causal models (Rehder 2003, Griffiths & Tenenbaum 2009)

Two Coin Flips

- Friend is flipping coin twice: $\Omega = \{HH, HT, TH, TT\}$, each event has $P(\omega) = \frac{1}{4}$
- Learn: 'Probability first toss lands heads is 0.7'; $P(HH \vee HT) = 0.7$
- This is probabilistic constraint incompatible with initial theory

Two Coin Flips II

- Assume: coin has constant bias, so $\mathcal{H}$ parameterized by $\theta \in [0, 1]$
- Only hypothesis consistent with constraint is $\theta = 0.7$, so posterior has $P_C(HH) = .49$, $P_C(HT) = P_C(TH) = .21$, and $P_C(TT) = .09$
- Posterior which minimizes new information only changes expectation of first coin toss: $P_J(HH) = P_J(HT) = .35$ and $P_J(TH) = P_J(TT) = .15$

Sundowners

- Learn 'If it rains, sundowners will be canceled'
- Prior beliefs: 50% chance of rain (R) and cancellation of sundowners ($\neg S$); independent
- Information minimization: only need to ensure $P(R \wedge S) = 0$ to make conditional true, so $R \wedge \neg S, \neg R \wedge S$, and $\neg R \wedge \neg S$ all have probability $\frac{1}{3}$
- But then $P(R) = \frac{1}{3}$; learn it is less likely to rain!

Sundowners II

- Causal theories relating rain and sundowners in $\mathcal{H}_C$ (Vandenburgh 2020):

$$
\begin{array}{ccc}
R & \qquad\qquad & \neg S \\
\downarrow & & \downarrow \\
\neg S & & R
\end{array}
$$

- First is most plausible: R is independent variable, so probability stays the same, and probability that sundowners is canceled is now .75 (100% if it rains, 50% if it doesn't)

Die Example: Underdetermination

- Recall: ‘If die lands odd, it lands 3’. Is there most plausible new hypothesis?
- Perhaps not: could have all faces 3, other odd faces replaced with 3, etc.
- Learning problems can be underdetermined: not enough contextual information to settle on new hypothesis

Judy Benjamin

- Judy Benjamin dropped onto grid, trying to figure out location

R2	R1
B2	B1

- Initially, each quadrant equally likely. Then learn $P(R1|R) = \frac{3}{4}$, where $R = R1 \vee R2$
- Information minimization: $P(R)$ decreases, so Benjamin less likely to be in enemy territory (Van Fraassen 1981)

Judy Benjamin II

Learning outcome depends on the hypothesis explaining new information

- Perhaps there is a small area off-limits where Benjamin couldn't be dropped. Given information, this would most likely be in $R2$, so $P(R)$ decreases
- Perhaps there was a target for Benjamin's drop. Targets of $R1$, $B1$, and $R1 \vee B1$ are consistent with new information. $P(R)$ could increase, decrease, or remain same.

Summary

- Learning a constraint incompatible with a theory requires finding a new (sometimes very different) theory
- This applies to theory-level constraints like conditionals and probabilistic constraints
- Hypothesis testing can offer compelling predictions through tools like parameterization and causal modeling
- Hypothesis testing is compatible with some learning problems being underdetermined

Bibliography I

Carey, S. & Spelke, E. (1994), ‘Domain-specific knowledge and conceptual change’, *Mapping the mind: Domain specificity in cognition and culture* **169**, 200.

Douven, I. & Romeijn, J.-W. (2011), ‘A new resolution of the Judy Benjamin problem’, *Mind* **120**(479), 637–670.

Gopnik, A. (1996), ‘The scientist as child’, *Philosophy of science* **63**(4), 485–514.

Griffiths, T. L. & Tenenbaum, J. B. (2009), ‘Theory-based causal induction’, *Psychological review* **116**(4), 661–716.

Rehder, B. (2003), ‘A causal-model theory of conceptual representation and categorization.’, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29**(6), 1141.

Salmon, T. C. (2004), ‘Evidence for learning to learn behavior in normal form games’, *Theory and decision* **56**(4), 367–404.

Bibliography II

Van Fraassen, B. C. (1981), ‘A problem for relative information minimizers in probability kinematics’, *The British Journal for the Philosophy of Science* **32**(4), 375–379.

Vandenburgh, J. (2020), ‘Conditional learning through causal models’, *Synthese* . https://doi.org/10.1007/s11229-020-02891-x.

Veltman, F. (1996), ‘Defaults in update semantics’, *Journal of philosophical logic* **25**(3), 221–261.

Young, H. P. (2004), *Strategic learning and its limits*, OUP Oxford.