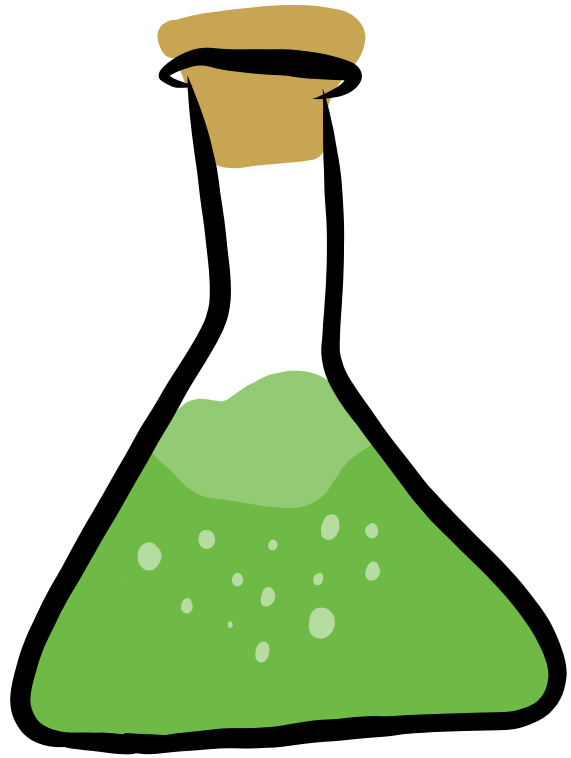

Wessel Bruinsma

29 October 2021

A Bayesian Truth Serum

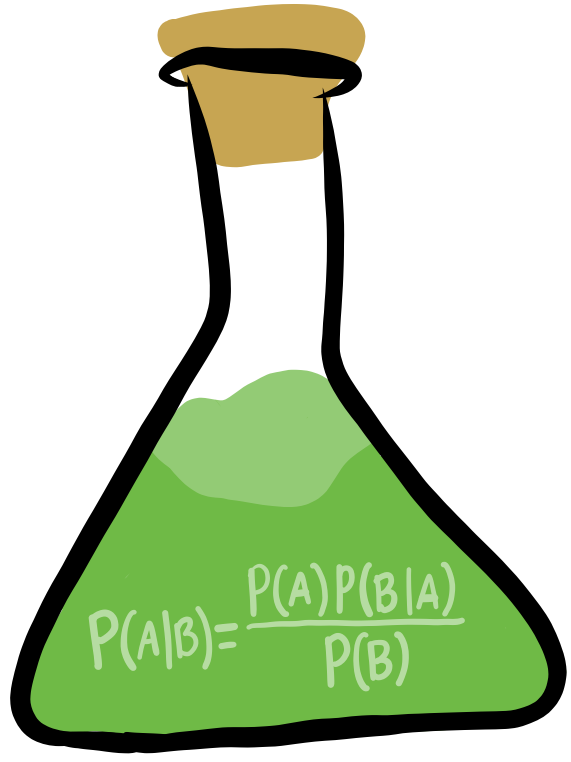

Wessel Bruinsma

29 October 2021

Prelec, D. (2004). A Bayesian Truth Serum for Subjective Data. *Science, 306*(5695), 462–466.

Motivation

"Do you like this painting?"

Motivation

"Do you like this painting?"

common prediction

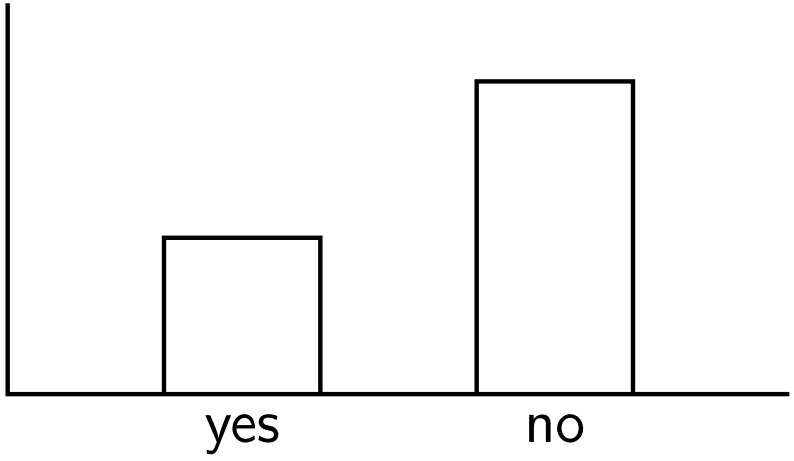

Motivation

"Do you like this painting?"

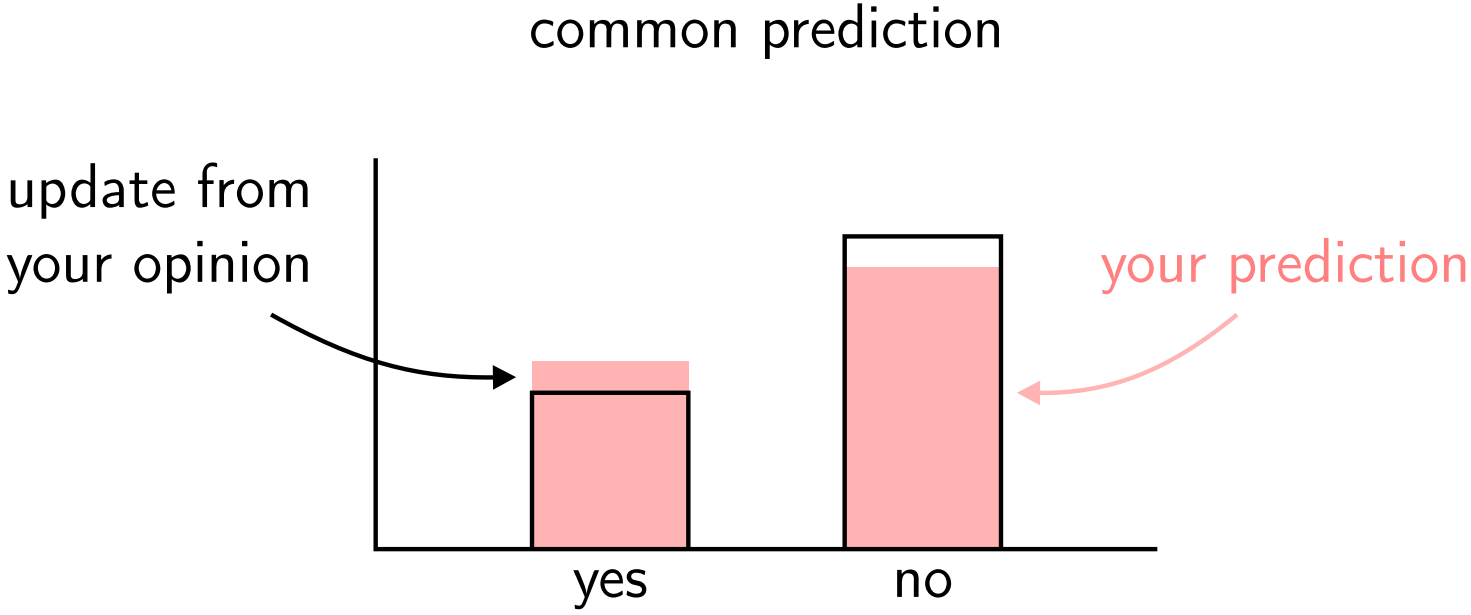

Motivation

"Do you like this painting?"

Your opinion is the opinion that you believe has the highest probability of being *more common than commonly predicted*.

The Serum

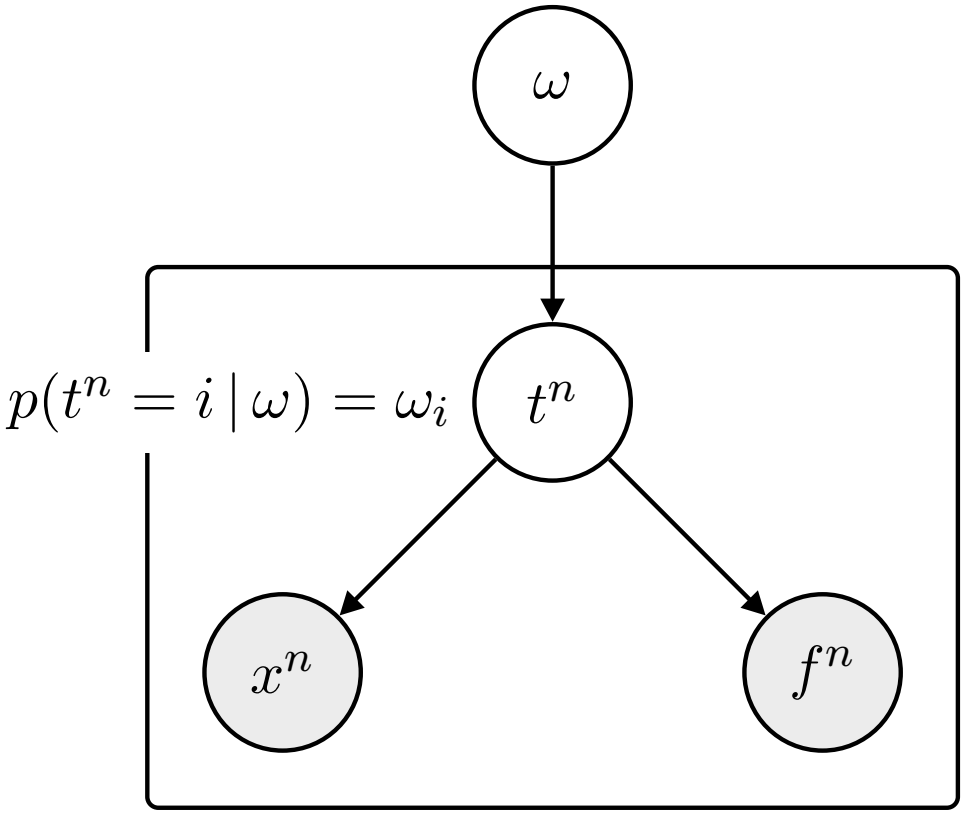

$\omega =$ distribution of opinions

$t^n =$ opinion

$x^n =$ answer

$f^n =$ prediction of
frequencies of answers

The Serum

$$\text{ }(x = i, f) = \overbrace{\log \frac{\langle x \rangle_i}{\langle f \rangle_i}}^{\text{information score}} - \overbrace{\sum_j \langle x \rangle_j \log \frac{f_j}{\langle x \rangle_j}}^{\text{prediction penalty}}$$

$\langle x \rangle_i =$ average of $(\mathbb{1}(x^n = i))$

$\langle f \rangle =$ geometric average of (f^n)

Truth Telling is a Bayesian Nash Equilibrium

Truth Telling is a Bayesian Nash Equilibrium

Proposition

Suppose that a respondent holds opinion t, answers x, and predicts f; and everyone else answers and predicts honestly. Then the respondent does best also by answering and predicting honestly:

$$\max_{(x,f)} \mathbb{E}\left[\text{ }(x, f) \,\middle|\, t \right] = (t, p(t' \mid t)).$$

Recapitulation

Prelec, D. (2004). A Bayesian Truth Serum for Subjective Data. *Science, 306*(5695), 462–466.

- Your opinion is the opinion that you believe has the highest probability of being *more common than commonly predicted*.

$$\text{ }(x_i, f) = \overbrace{\log \frac{\langle x \rangle_i}{\langle f \rangle_i}}^{\text{information score}} - \overbrace{\sum_j \langle x \rangle_j \log \frac{f_j}{\langle x \rangle_j}}^{\text{prediction penalty}}$$

- Truth telling is a Bayesian Nash equilibrium.

These slides: `https://wesselb.github.io/pdf/serum`.
Write-up: `https://wesselb.github.io/pdf/serum-write-up`.