Cognition 120 (2011) 302–321

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

A tutorial introduction to Bayesian models of cognitive development

Amy Perfors [a,*], Joshua B. Tenenbaum [b], Thomas L. Griffiths [c], Fei Xu [c]

[a] School of Psychology, Level 4, Hughes Building, University of Adelaide, Adelaide, SA 5050, Australia
[b] Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, United States
[c] Department of Psychology, 3210 Tolman Hall, University of California, Berkeley, Berkeley, CA 94720-1650, United States

ARTICLE INFO

Article history:
Received 10 April 2010
Revised 11 October 2010
Accepted 25 November 2010
Available online 26 January 2011

Keywords:
Bayesian models
Cognitive development

ABSTRACT

We present an introduction to Bayesian inference as it is used in probabilistic models of cognitive development. Our goal is to provide an intuitive and accessible guide to the *what*, the *how*, and the *why* of the Bayesian approach: what sorts of problems and data the framework is most relevant for, and how and why it may be useful for developmentalists. We emphasize a qualitative understanding of Bayesian inference, but also include information about additional resources for those interested in the cognitive science applications, mathematical foundations, or machine learning details in more depth. In addition, we discuss some important interpretation issues that often arise when evaluating Bayesian models in cognitive science.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

One of the central questions of cognitive development is how we learn so much from such apparently limited evidence. In learning about causal relations, reasoning about object categories or their properties, acquiring language, or constructing intuitive theories, children routinely draw inferences that go beyond the data they observe. Probabilistic models provide a general-purpose computational framework for exploring how a learner might make these inductive leaps, explaining them as forms of Bayesian inference.

This paper presents a tutorial overview of the Bayesian framework for studying cognitive development. Our goal is to provide an intuitive and accessible guide to the *what*, the *how*, and the *why* of the Bayesian approach: what sorts of problems and data the framework is most relevant for, and how and why it may be useful for developmentalists. We consider three general inductive problems that learners face, each grounded in specific developmental challenges:

1. Inductive generalization from examples, with a focus on learning the referents of words for object categories.
2. Acquiring inductive constraints, tuning and shaping prior knowledge from experience, with a focus on learning to learn categories.
3. Learning inductive frameworks, constructing or selecting appropriate hypothesis spaces for inductive generalization, with applications to acquiring intuitive theories of mind and inferring hierarchical phrase structure in language.

We also discuss several general issues as they bear on the use of Bayesian models: assumptions about optimality, biological plausibility, and what idealized models can tell us about actual human minds. The paper ends with an appendix containing a glossary and a collection of useful resources for those interested in learning more.

2. Bayesian basics: inductive generalization from examples

The most basic question the Bayesian framework addresses is how to update beliefs and make inferences in light of observed data. In the spirit of Marr's (1982)

* Corresponding author. Tel.: +61 8 8303 5744.
E-mail address: amy.perfors@adelaide.edu.au (A. Perfors).

0010-0277/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2010.11.015

computational level of analysis, it begins with understanding the logic of the inference made when generalizing from examples, rather than the algorithmic steps or specific cognitive processes involved. A central assumption is that degrees of belief can be represented as probabilities: that our conviction in some hypothesis h can be expressed as a real number ranging from 0 to 1, where 0 means something like "h is completely false" and 1 that "h is completely true." The framework also assumes that learners represent probability distributions and that they use these probabilities to represent uncertainty in inference. These assumptions turn the mathematics of probability theory into an engine of inference, a means of weighing each of a set of mutually exclusive and exhaustive hypotheses $\mathcal{H}$ to determine which best explain the observed data. Probability theory tells us how to compute the degree of belief in some hypothesis h_i, given some data d.

Computing degrees of belief as probabilities depends on two components. One, called the *prior probability* and denoted $P(h_i)$, captures how much we believe in h_i prior to observing the data d. The other, called the *likelihood* and denoted $P(d|h_i)$, captures the probability with which we would expect to observe the data d if h_i were true. These combine to yield the *posterior probability* of h_i, given via Bayes' Rule:

$$P(h_i|d) = \frac{P(d|h_i)P(h_i)}{\sum_{h_j \in \mathcal{H}} P(d|h_j)P(h_j)}. \quad (1)$$

As we will see, the product of priors and likelihoods often has an intuitive interpretation. It balances between a sense of plausibility based on background knowledge on one hand and the data-driven sense of a "suspicious coincidence" on the other. In the spirit of Ockham's Razor, it expresses the tradeoff between the intrinsic complexity of an explanation and how well it fits the observed data.

The denominator in Eq. (1) provides a normalizing term which is the sum of the probability of each of the possible hypotheses under consideration; this ensures that Bayes' Rule will reflect the proportion of all of the probability that is assigned to any single hypothesis h_i, and (relatedly) that the posterior probabilities of all hypotheses sum to one. This captures what we might call the "law of conservation of belief": a rational learner has a fixed "mass" of belief to allocate over different hypotheses, and the act of observing data just pushes this mass around to different regions of the hypothesis space. If the data lead us to strongly believe one hypothesis, we must decrease our degree of belief in all other hypotheses. By contrast, if the data strongly disfavor all but one hypothesis, then (to paraphrase Sherlock Holmes) whichever remains, however implausible *a priori*, is very likely to be the truth.

To illustrate how Bayes' Rule works in practice, let us consider a simple application with three hypotheses. Imagine you see your friend Sally coughing. What could explain this? One possibility (call it h_{cold}) is that Sally has a cold; another (call it h_{cancer} is that she has lung cancer; and yet another (call it $h_{\text{heartburn}}$) is that she has heartburn. Intuitively, in most contexts, h_{cold} seems by far the most probable, and may even be the only one that comes to mind consciously. Why? The likelihood favors h_{cold} and h_{cancer} over $h_{\text{heartburn}}$, since colds and lung cancer cause coughing, while heartburn does not. The prior, however, favors h_{cold} and $h_{\text{heartburn}}$ over h_{cancer}: lung cancer is thankfully rare, while colds and heartburn are common. Thus the posterior probability – the product of these two terms – is high only for h_{cold}.

The intuitions here should be fairly clear, but to illustrate precisely how Bayes' Rule can be used to back them up, it can be helpful to assign numbers.[1] Let us set the priors as follows: $P(h_{\text{cold}}) = 0.5$, $P(h_{\text{heartburn}}) = 0.4$, and $P(h_{\text{cancer}}) = 0.1$. This captures the intuition that colds are slightly more common than heartburn, but both are significantly more common than cancer. We can set our likelihoods to be the following: $P(d|h_{\text{cold}}) = 0.8$, $P(d|h_{\text{cancer}}) = 0.9$, and $P(d|h_{\text{heartburn}}) = 0.1$. This captures the intuition that both colds and cancer tend to lead to coughing, and heartburn generally does not. Plugging this into Bayes' Rule gives:

$$P(h_{\text{cold}}|d) = \frac{P(d|h_{\text{cold}})P(h_{\text{cold}})}{P(d|h_{\text{cold}})P(h_{\text{cold}}) + P(d|h_{\text{cancer}})P(h_{\text{cancer}}) + P(d|h_{\text{heartburn}})P(h_{\text{heartburn}})}$$
$$= \frac{(0.8)(0.5)}{(0.8)(0.5) + (0.9)(0.1) + (0.1)(0.4)} = \frac{0.4}{0.4 + 0.09 + 0.04} = 0.7547.$$

Thus, the probability that Sally is coughing because she has a cold is much higher than the probability of either of the other two hypotheses we considered. Of course, these inferences could change with different data or in a different context. For instance, if the data also included coughing up blood, chest pain, and shortness of breath, you might start to consider lung cancer as a real possibility: the likelihood now explains that data better than a cold would, which begins to balance the low prior probability of cancer in the first place. On the other hand, if you had other information about Sally – e.g., that she had been smoking two packs of cigarettes per day for 40 years – then it might raise the prior probability of lung cancer in her case. Bayes' Rule will respond to these changes in the likelihood or the prior in a way that accords with our intuitive reasoning.

The Bayesian framework is generative, meaning that observed data are assumed to be generated by some underlying process or mechanism responsible for creating the data. In the example above, data (symptoms) are generated by an underlying illness. More cognitively, words in a language may be generated by a grammar of some sort, in combination with social and pragmatic factors. In a physical system, observed events may be generated by some underlying network of causal relations. The job of the learner is to evaluate different hypotheses about the underlying nature of the generative process, and to make predictions based on the most likely ones. A probabilistic model is simply a specification of the generative processes at work, identifying the steps (and associated probabilities) involved in generating data. Both priors and likelihoods are typically describable in generative terms.

To illustrate how the nature of the generative process can affect a learner's inference, consider another example, also involving illness. Suppose you observe that 80% of the

[1] Note that we have assumed that these are the only possible hypotheses, and that exactly one applies. That is why the priors are much higher than the base rates of these diseases. In a real setting, there would be many more diseases under consideration, and each would have much lower prior probability. They would also not be mutually exclusive. Adding such details would make the math more complex but not change anything else, so for clarity of exposition we consider only the simplified version.

دریافت فوری ← متن کامل مقاله

ISIArticles
مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات