

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Generative and Discriminative Models in Cognitive Science

Permalink

<https://escholarship.org/uc/item/0049w8mt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Love, Bradley C
Ramscar, Michael
Griffiths, Tom
[et al.](#)

Publication Date

2015

Peer reviewed

Generative and Discriminative Models in Cognitive Science

organizers and speakers

Bradley C. Love (b.love@ucl.ac.uk)

University College London

Michael Ramscar (michael.ramscar@uni-tuebingen.de)

University of Tübingen

invited speakers

Tom Griffiths (tom_griffiths@berkeley.edu)

University of California, Berkeley

Matt Jones (mcj@colorado.edu)

University of Colorado at Boulder

Generative and Discriminative Approaches

One popular distinction in machine learning is between discriminative and generative models (Ng & Jordan, 2001). Given the cross fertilization between research in human and machine learning, the time is ripe to ask whether the mind is a generative or discriminative learning device. This symposium tackles this question from a variety of perspectives. The aim is to explore the explanatory value of these two basic views of learning, which cut across existing distinctions in cognitive science (e.g., connectionist vs. Bayesian approaches).

In brief, generative and discriminative models characterize the task of the learner differently. Generative models attempt to learn an internal model of each class (i.e., category). In contrast, discriminative models attempt to find a boundary that separates classes. Generative models are typically Bayesian in form, whereas discriminative models include decision trees, SVMs, regression approaches, and some (but not all) connectionist models.

In generative models, the learning task is to estimate the joint probabilities between all variables. These models assume a hidden or latent variable (e.g., category label) generates observed features. In contrast, discriminative models perform a conditional estimation. For example, logistic regression only estimates the probability of a class (i.e., category) as a function of the predictive features. In this sense, discriminative models are more focused by the task, whereas generative models address a broader estimation problem, though models of all types have an inductive bias to make learning tractable.

Aims and Relevance

The aim of this symposium is to introduce these powerful ideas from machine learning to the broader cognitive science community. We will evaluate what these two views say about cognition and learning, and assess their utility in organizing findings in our science. At the broadest level, in what sense is the mind a generative or discriminative machine, and how can this understanding direct our future empirical and theoretical investigations of the mind?

The impact of sampling assumptions on learning from indirect negative evidence

Tom Griffiths (with Anne Hsu)

A classic debate in cognitive science revolves around understanding how children learn complex linguistic patterns, such as restrictions on verb alternations and contractions, without negative evidence. One factor that has been suggested as playing an important role in solving this problem is indirect negative evidence, in which the absence of a construction in the input provides evidence against its grammaticality. We consider two different sets of sampling assumptions that can operate in language learning, corresponding loosely to "generative" and "discriminative" approaches to learning. Only one set of assumptions licenses use of indirect negative evidence. We demonstrate in a series of experiments in which adults learn artificial languages that people can produce behavior consistent with the predictions of probabilistic models using both sets of sampling assumptions, depending on how the learning problem is presented. These results suggest that people use information about the way in which linguistic input is sampled to guide their learning, and show that adult learners make appropriate use of indirect negative evidence when the appropriate statistical assumptions are satisfied.

Language Learning From a Discriminative Perspective

Michael Ramscar

The development of morphological processing has been the focal topic in a debate over the nature of language, learning and the mind in cognitive science. Particular attention has been paid to the systematic nature of children's morphological errors (for example children tend to go through a phase of saying "mouses" as they learn the morphology of English nouns). Because these errors aren't explicitly corrected, it has been argued that the transition to adult language cannot be explained by learning, and that the acquisition of even relatively simple aspects of grammar must involve innate, language specific mechanisms. I'll describe the background to this debate, the generative models that have traditionally been proposed to explain

these behavioral patterns, and a model of morphological development based on discrimination learning that offers a very different perspective on morphological processing. This model also generates clear and surprising predictions, in particular that exposure to regular plurals (e.g. rats) can actually result in a decrease in children's tendency to overregularize irregular plurals (e.g. say "mouses"). I'll review some empirical results showing that testing memory for items with regular plural labels does result in a decrease in plural overregularization in six-year-olds, but also that it results in increases in four-year-olds. These models and results indicate that when the learning problem facing children is characterized discriminatively, overregularization can be seen to both arise and then resolve itself as a result of the distribution of evidence in the linguistic environment. I'll discuss the wider implications of these and some similar findings for our understanding of language and human communication.

Ramscar, M., Dye, M. & McCauley, S. (2013d) Error and expectation in language learning: The curious absence of 'mouses' in adult speech. *Language*, 89(4), 670-793

Sequential Effects as Signatures of Discriminative Learning

Matt Jones

An important class of psychological models of discriminative learning are those that learn incrementally from prediction error. One prediction of this iterated error correction is recency effects. In their simplest form, recency effects are simply a bias toward recent events, such that error rates and response times (RTs) are lower when the current trial matches recent feedback. Generative models can also predict these simple recency effects, by assuming nonstationarity in latent environmental parameters (Wilder, Jones, & Mozer, 2009; Yu & Cohen, 2008). This is because the nonstationarity assumption leads more recent events to be more informative about the current state of the environment. However, discriminative models also predict more complex sequential effects that generative models do not anticipate. This talk will focus on one set of such findings, in binary stimulus identification tasks (Jones, Curran, Mozer, & Wilder, 2013). In this paradigm, sequential effects in RT reveal learning of two statistics of the trial sequence: the base rate and the repetition rate. That is, RT is faster when the current response matches recent responses (a left response preceded by recent left responses, or a right response preceded by recent right responses), and RT is also faster on a repetition trial preceded by recent repetition trials or on an alternation trial preceded by recent alternation trials. This basic pattern is well fit by both a generative Bayesian model (Wilder et al., 2009) and a discriminative error-correction model (Jones et al., 2013). The two models diverge in their predictions for how the two learning mechanisms interact, with the error-correction model predicting cue-competition effects whereby the

expectancy derived from the base rate affects learning about the repetition rate and vice versa. This cue competition manifests in additional, subtle sequential effects that are confirmed in the data. These additional sequential effects thus appear to be signatures of discriminative learning.

Jones, M., Curran, T., Mozer, M. C., & Wilder, M. H. (2013). Sequential effects in response time reveal learning mechanisms and event representations. *Psychological Review*, 120, 628-666.

Getting Discriminative with a Generative Model

Bradley C. Love

Models, whether generative or discriminative, have an inductive bias that makes learning tractable. In this talk, I will present a generative model of learning and information sampling whose inductive bias follows from discriminative principles. The model, like people, is focused on properly estimating aspects of the environment that are goal relevant. This focus is consistent with conditional estimation in discriminative models. However, the model also benefits from the strengths of the generative approach, such as the ability to support planning and sampling processes critical in decision making (Giguère & Love, 2013). Like people, current goals and knowledge determine the information sampled in the world. Completing the cycle of mutual influence, the information sampled (i.e., attended) in the world updates the model's knowledge state. This cycle of influence depends on two model components. One model component determines the value of potential sources of information. The value of a piece of information depends on the decision maker's goals and assumptions about (i.e., knowledge of) the world. The second component of the model reflects the decision maker's knowledge of the world, which is used by the first component to direct information gathering. This learning component is updated by the information samples selected by the first component, completing the cycle of mutual influence. Human learning and eye tracking studies support the model. By introducing a notion of attention that focuses on goal-discriminating information, a generative model is imparted with discriminative characteristics and displays human-like behaviors.

Giguère, G. & Love, B.C. (2013). Limits in decision making arise from limits in memory retrieval. *Proceedings of the National Academy of Sciences of the United States of America* (PNAS), 110 (19), 7613-7618.

References

Ng, A.Y., & Jordan, M.I. (2001). On Discriminative vs. Generative classifiers: A comparison of logistic regression and naive Bayes. *Advances in neural information processing systems* 14, 841-849.