ARTICLE IN PRESS

Journal of Mathematical Psychology ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Journal of Mathematical Psychology

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

Bayesian analysis of simulation-based models☆

Brandon M. Turner[a,*], Per B. Sederberg[a], James L. McClelland[b]

[a] Department of Psychology, The Ohio State University, United States

[b] Department of Psychology, Stanford University, United States

HIGHLIGHTS

- Neurologically plausible models of choice response time are compared.
- Several model comparison metrics are evaluated.
- Discrepancies between the metrics are observed.

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Bayes factor
Likelihood-free inference
Simulation models
Neurologically plausible cognitive models
Probability density approximation method
Leaky Competing Accumulator model
Feed Forward Inhibition model

ABSTRACT

Recent advancements in Bayesian modeling have allowed for likelihood-free posterior estimation. Such estimation techniques are crucial to the understanding of simulation-based models, whose likelihood functions may be difficult or even impossible to derive. One particular class of simulation-based models that have not yet benefited from the progression of Bayesian methods is the class of neurologically-plausible models of choice response time, in particular the Leaky, Competing Accumulator (LCA) model and the Feed-Forward Inhibition (FFI) model. These models are unique because their architecture was designed to embody actual neuronal properties such as inhibition, leakage, and competition. Currently, these models have not been formally compared by way of principled statistics such as the Bayes factor. Here, we use a recently developed algorithm – the probability density approximation method – to fit these models to empirical data consisting of a classic speed accuracy trade-off manipulation. Using this approach, we find some discrepancies between an assortment of model fit statistics. For some participants, one model appears to be superior when one fit statistic is used, while another appears superior when a different statistic is used. However, for 13 of the 20 participants, one model wins by all of the fit metrics considered. The FFI wins in 5 of these cases, while the LCA wins, often by a wide margin, for the others.

1. Introduction

The goals of cognitive modeling are to understand complex behaviors within a system of mathematically-specified mechanisms or processes, to assess the adequacy of the model in accounting for experimental data, and to obtain an estimate of the model parameters, which carry valuable information about how the model captures the observed behavior for both individuals and groups. Cognitive models are important because they provide a means with which cognitive theories can be explicitly tested and compared with one another.

Perhaps the greatest strength of many cognitive models is paradoxically the model's greatest weakness. Many cognitive models put forth sophisticated mechanisms meant to capture psychologically plausible processes. While these mechanisms are entirely plausible, they often render the cognitive model intractable, or at least difficult to fully analyze in a principled way such as with Bayesian statistics. The difficulties encountered in deriving the full likelihood function have prevented the application of fully Bayesian analyses for many cognitive models, especially those that attempt to capture neurally-plausible mechanisms.

Consider, for example, the Leaky Competing Accumulator (LCA; Usher & McClelland, 2001) model. The LCA model was proposed as a neurologically plausible model for choice response time in a c-alternative task. The model possesses mechanisms that extend other diffusion-type models (e.g., Ratcliff, 1978) by including leakage and competition by means of lateral inhibition. Because the evidence accumulation process used by the LCA model was

☆ This work was funded by NIH award (F32GM103288) to B.M.T.

* Corresponding author.
E-mail address: turner.826@gmail.com (B.M. Turner).

http://dx.doi.org/10.1016/j.jmp.2014.10.001
0022-2496/

designed to mimic actual neuronal activation patterns, one critical assumption is that the signal propagated from one accumulator to another can never be negative. This assumption can be implemented by specifying a floor on each accumulator's activation value, such that if the activation of an accumulator in the model becomes negative, it is reset to zero. The LCA model also assumes a competition among response alternatives that depends on the current state of each of the accumulators. Together, these features of the model sufficiently complicate the equations describing the joint distributions of choice and response time such that the likelihood function for the LCA model has not been derived. As a result, all model evaluations to this point have been performed using either a model simplification or least squares estimation (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Bogacz, Usher, Zhang, & McClelland, 2007; Gao, Tortell, & McClelland, 2011; Teodorescu & Usher, 2013; Tsetsos, Usher, & McClelland, 2011; Usher & McClelland, 2001; van Ravenzwaaij, van der Maas, & Wagenmakers, 2012), which have been shown to produce less accurate parameter estimates relative to techniques such as maximum likelihood or Bayesian estimation (e.g., Myung, 2003; Rouder, Sun, Speckman, Lu, & Zhou, 2003; Turner, Dennis, & Van Zandt, 2013; Van Zandt, 2000).

Recent advances in likelihood-free techniques have allowed for new insights to simulation-based cognitive models (Turner, Dennis et al., 2013; Turner & Sederberg, 2012, 2014; Turner & Van Zandt, 2012, 2014). In particular, the probability density approximation (PDA; Turner & Sederberg, 2014) method now allows for fully Bayesian analyses of computational models exclusively by way of simulation. In this article, we illustrate the importance of our method by comparing two neural network models of choice response time that have never been compared using Bayesian techniques due to their computational complexity: the LCA model (Usher & McClelland, 2001) and the Feed-Forward Inhibition (FFI; Shadlen & Newsome, 2001) model.[1] Both models embody neurologically plausible mechanisms such as "leakage", or the passive decay of evidence during a decision, and competition among alternatives through either lateral inhibition (in the LCA model) or feed-forward inhibition (in the FFI model). However, it remains unclear as to which dynamical system best accounts for empirical data, due to the limitations imposed by intractable likelihoods. Specifically, complexity measures that take into account posterior uncertainty and model complexity have yet to be applied. Here, we will compare the models on the basis of an approximation to the Bayes factor. We begin by describing in greater detail our method for fitting the models to data. We then describe how our posterior estimates are converted into a comparison between the models. Finally, we compare the relative merits of the two models by evaluating the models' fit to the data presented in Forstmann et al. (2011), which consisted of 20 subjects in three speed emphasis conditions.

2. Experiment

The data we will use to test the models were presented in Forstmann et al. (2011), and consist of 20 subjects. The experiment used a moving dots task where subjects were asked to decide whether a cloud of semi-randomly moving dots appeared to move to the left or to the right. Subjects indicated their response by pressing one of two spatially compatible buttons with either their left or right index finger. Before each decision trial, subjects were instructed whether to respond quickly (the speed condition), accurately (the accuracy condition), or at their own pace (the neutral condition). Following the trial, subjects were provided feedback about their performance. In the speed and neutral conditions, subjects were told that their responses were too slow whenever they exceeded a RT of 400 and 750 ms, respectively. In the accuracy condition, subjects were told when their responses were incorrect. Each subject completed 840 trials, equally distributed over the three conditions. These data serve as a benchmark for our metric comparison given that we have some experience in analyzing them in a variety of contexts (Turner et al., 2013; Turner & Sederberg, 2014; Turner, Sederberg, Brown, & Steyvers, 2013).

3. Likelihood-free inference

As the reader of this special issue is no doubt aware, there are many advantages of using Bayesian statistics in cognitive modeling. However, the widespread dissemination of Bayesian statistics can largely be attributed to advanced statistical techniques for approximating the posterior distribution (see, e.g., Gelman, Carlin, Stern, & Rubin, 2004; Gilks, Best, & Tan, 1995; Gilks & Wild, 1992; Robert & Casella, 2004; Ter Braak, 2006), rather than evaluating it precisely. Approximating any posterior distribution depends on efficient evaluation of two functions: (1) the prior distribution for the model parameters, and (2) the likelihood function relating the model parameters to the observed data. For purely statistical models, evaluating these functions is, generally speaking, straightforward. However, for cognitive models who attempt to provide mechanistic explanations for how data manifest, direct evaluation of the likelihood function can be difficult, if not impossible. We refer to these models as "simulation-based" to indicate that explicit equations for the likelihood function are either (1) intensely difficult to practically evaluate (e.g., Montenegro, Myung, & Pitt, 2011; Myung, Montenegro, & Pitt, 2007; Turner, Dennis et al., 2013), or (2) have not yet been derived (e.g., Shadlen & Newsome, 2001; Usher & McClelland, 2001). Recently, a suite of algorithms have been developed specifically for analyzing (simulation-based) cognitive models in a fully (hierarchical) Bayesian context (Turner & Sederberg, 2012, 2014; Turner & Van Zandt, 2014). While combinations of these algorithms can be used to effectively evaluate the joint posterior distribution, we require only one algorithm – the probability density approximation (PDA; Turner & Sederberg, 2014) method – to evaluate the models presented in this article.

3.1. The probability density approximation method

As discussed in Turner and Sederberg (2014), the PDA method is an alternative likelihood-free algorithm that does not require sufficient statistics for the parameters of interest. Turner and Sederberg demonstrated the utility of their algorithm by verifying that it could be used to accurately estimate the posterior distribution of the parameters of the Linear Ballistic Accumulator (LBA; Brown & Heathcote, 2008) model, which has a tractable likelihood function and is amenable to Bayesian estimation (Donkin, Averell, Brown, & Heathcote, 2009; Donkin, Heathcote, & Brown, 2009; Turner, Sederberg et al., 2013). In addition, Turner and Sederberg showed that the PDA method could be used to estimate the parameters of the LCA model in a fully hierarchical Bayesian context.

Although the details of how to apply the PDA method to various data types are explained in detail in Turner and Sederberg (2014), we will reproduce the relevant details for applying the method to data containing both discrete and continuous measures. For ease of exposition, we consider the common case of data consisting of one discrete measurement (e.g., choice) and one continuous measurement (e.g., response time). For the discrete measurements, suppose there are C options, and for the continuous measurements

[1] Although Ratcliff and Smith (2004) used the Bayesian information criteria to compare many simulation-based models, they did not obtain proper Bayesian posteriors, which is the endeavor of the current manuscript.