

Look-Ahead Monte Carlo with People

Charles Blundell (c.blundell@gatsby.ucl.ac.uk)

Gatsby Computational Neuroscience Unit, University College London, London, UK

Adam Sanborn (a.n.sanborn@warwick.ac.uk)

Department of Psychology, University of Warwick, Coventry, UK

Thomas L. Griffiths (tom_griffiths@berkeley.edu)

Department of Psychology, University of California, Berkeley, Berkeley, CA USA

Abstract

Investigating people’s representations of categories of complicated objects is a difficult challenge, not least because of the large number of ways in which such objects can vary. To make progress we need to take advantage of the structure of object categories – one compelling regularity is that object categories can be described by a small number of dimensions. We present Look-Ahead Monte Carlo with People, a method for exploring people’s representations of a category where there are many irrelevant dimensions. This method combines ideas from Markov chain Monte Carlo with People, an experimental paradigm derived from an algorithm for sampling complicated distributions, with hybrid Monte Carlo, a technique that uses directional information to construct efficient statistical sampling algorithms. We show that even in a simple example, our approach takes advantage of the structure of object categories to make experiments shorter and increase our ability to accurately estimate category representations.

Keywords: category representation; Markov chain Monte Carlo; directional judgements

Introduction

Categories are an essential component of how people reason about the world, allowing us to act intelligently when we encounter new objects, new people, or new situations. Natural stimuli, such as images or text, tend to be very complicated. In contrast, much of our understanding of categorisation behaviour has been built on experiments using well-controlled stimuli that vary along only one or two dimensions (e.g., Nosofsky, 1986). A major driver of the disconnect between experimental investigations and real-world behaviour is that standard experimental methods do not allow the experimenter to adaptively focus on informative stimuli. Attempting to map out a complicated category with standard methods would require participants to complete an unendurable number of trials.

Identifying connections between cognitive science and statistics can help us develop new methods to extend our experimental reach. Many of the computational models of categorisation developed from carefully-controlled laboratory studies can be interpreted as representing categories as probability distributions over their constituent stimuli (Ashby & Alfonso-Reese, 1995). Statisticians have developed sophisticated methods to sample from high-dimensional distributions and Sanborn and Griffiths (2008) identified how to use one of these methods, Markov Chain Monte Carlo (MCMC), to draw samples from categories that were represented in the mind. This method, termed Markov Chain Monte Carlo with

People (MCMCP), focused trials on informative regions of stimuli – greatly reducing the number of trials a participant needs to perform.

While MCMCP exploited the local nature of categories, it could not exploit the way that many categories lie on manifolds. Many seemingly complex stimulus representations possess a simpler embedded representation. For example, images, when represented as a grid of pixel values, inhabit a high-dimensional space, typically hundreds to hundreds of thousands of dimensions. Yet modifications to these images, such as changing a facial expression or moving a limb, require simultaneously modifying only some of the dimensions. Thus often stimuli possess many irrelevant dimensions and many “linked” dimensions, where changing one implies proportionally changing many other connected dimensions. Since such relationships are often smooth among dimensions, suggesting a manifold structure.

In this paper, we present an extension to MCMCP that is able to take advantage of the manifold structure of continuous-valued stimuli and apply it to a low-dimensional stimulus. We take inspiration from the large literature in statistical computing that has explored how MCMC algorithms can be modified to increase their efficiency in solving specific problems. For example, Hybrid Monte Carlo (Neal, 1996) is a successful method for sampling in high-dimensional spaces that incorporates directional information into MCMC methods. We adapt MCMCP to include directional information, allowing participants to “look ahead” at the consequences of their decisions so they can guide stimulus generation, increasing the efficiency with which we can investigate categories.

The remainder of this paper is organised as follows. We first introduce the MCMC and explain how it was adapted to people. Next we introduce Look-Ahead Monte Carlo with People (LAMCP), first motivating it intuitively and then presenting the technical justification of our method. We compare LAMCP to MCMCP empirically in an experiment exploring the golden ratio in rectangles—a simple example of a low-dimensional manifold in a higher-dimensional space. We then conclude with a brief discussion how our method relates to hybrid Monte Carlo and other sampling techniques.

Background: MCMC with People

Using the connection between computational models of categorisation and probability distributions (Ashby & Alfonso-

Reese, 1995), let $p(x|c)$ denote the probability of stimulus x under the distribution associated with category c . We could attempt to elicit people’s category representation by asking them to rate the typicality of item x in category c , but that would require us to present participants with every stimulus of interest. Instead we would like to draw samples from $p(x|c)$.

The Metropolis sampling scheme (Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953) is a commonly used algorithm for generating a sequence of samples from a designer-supplied probability distribution. It constructs a Markov chain whose stationary distribution is the provided probability distribution and then uses this Markov chain to, eventually, generate a sequence of samples from the stationary distribution. The algorithm has two parts that are provided by the designer: a proposal distribution $q(x^*|x)$, and an acceptance function $A(x^*, x)$. The proposal distribution is typically a simple distribution (such as a Gaussian or uniform distribution) and is used to propose possible samples, x^* , given the current sample, x . The acceptance function tests how similar samples from this proposal distribution are to the desired stationary distribution, guiding the algorithm towards this distribution. Both the proposal distribution and acceptance function must be carefully selected to ensure the Markov chain converges correctly (see Neal, 1993). At each step t of the algorithm, a new state of the chain is proposed, x_t^* , by sampling it from the proposal distribution, $q(x_t^*|x_{t-1})$. With probability $A(x_t^*, x_{t-1})$, the state of the Markov chain at step t is the proposal x_t^* , otherwise it is the previous state, i.e., $x_t = x_{t-1}$. The initial state x_0 of the Markov chain is picked at random. It can be shown that if this procedure is repeated for long enough, the Markov chain it defines will eventually converge on the desired stationary distribution.

MCMCP (Sanborn & Griffiths, 2008; Sanborn, Griffiths, & Shiffrin, 2010) transformed the Metropolis sampling scheme into an experimental method for cognitive science. MCMCP is a sequential paradigm in which participants construct the stimuli themselves, in small, manageable steps. The equilibrium distribution of interest is the distribution over stimuli belonging to a single category $p(x|c)$, which we shall write as $\pi(x)$ for short-hand, the particular category being implicit.

On each trial of the MCMCP procedure, a new stimulus sample x_t^* is generated by a computer from an experimenter-provided proposal distribution, such as a Gaussian or uniform distribution. Participants are presented with a choice of two possible samples; the current state, x_{t-1} , or the proposed new state, x_t^* . Their selection becomes the new state of the Markov chain, x_{t+1} . The next trial has the same form, and the procedure repeats until the Markov chain is deemed to have equilibrated.

If participants choose x_t^* according to a valid acceptance function, then one can show that, just like Metropolis sampling, this scheme forms a Markov chain whose stationary distribution is the distribution over stimuli in a particular category $\pi(x)$. Fortunately there is an exact correspondence be-

tween the ratio rule of human decision making (Luce, 1963) and a valid acceptance function for use with Metropolis sampling, namely the Barker acceptance function (Barker, 1965):

$$A(x_t^*, x_{t-1}) = \frac{\pi(x_t^*)}{\pi(x_t^*) + \pi(x_{t-1})} \quad (1)$$

Assuming participants’ choices are Markov, the paradigm forms a Metropolis sampling scheme whose states are samples from people’s distribution of objects in a particular category at the equilibration of the Markov chain.

The proposal distribution used by MCMCP is fixed and provided by the experimenter; typically it is an isotropic Gaussian distribution. This induces a random-walk Markov kernel, which, as Neal (1993) notes for general MCMC, and as Martin, Griffiths, and Sanborn (2012) note for MCMCP, can be inefficient for exploring large, correlated spaces. Gains in efficacy are to be had by removing this random walk, whilst maintaining the properties that allow it to converge to the category distribution.

An Intuitive View of MCMCP and LAMCP

Before presenting our novel Look-Ahead Monte Carlo with People (LAMCP) method, we will provide some intuitions as to how it differs from MCMCP. Intuitively, we can think about exploring a probability distribution by following a Markov chain in terms of a hiker attempting to travel along a ridge path. The ridge represents an interesting low-dimensional manifold embedded in a largely irrelevant higher-dimensional space, and we wish participants to explore this manifold efficiently.

Suppose our hiker is standing upon a bumpy ridge in an otherwise large flat landscape. He wishes to explore the ridge, whilst only descending into the lower parts of the landscape fleetingly. Suppose also that the he cannot see anything, and must be told about the terrain by MCMCP or LAMCP.

MCMCP allows our hiker to know about the terrain where he currently stands and also at another location, picked at random by MCMCP. He must then choose whether he wishes to step to this new location or stay where he is. MCMCP does not know of our hiker’s intention to follow the ridge, and so when MCMCP proposes locations far away from him, the new location is likely to be in the flat and so he will often elect not to move. Proposed locations very close to our hiker are likely to be on the ridge so he will be willing to make a step. Walking along the ridge in this random fashion will take a great many small steps.

LAMCP gives our hiker a guide. This guide walks away from the hiker along a randomly oriented straight line. The guide then returns to the hiker and tells him on average, how much of the ridge she saw on her travels. If the hiker feels that the guide saw a lot of the ridge, then LAMCP randomly picks a location along this straight line. The hiker can then either stay where he is, or walk to this new location offered by LAMCP. With LAMCP, the hiker can take advantage of a longer view. In addition, our guide will propose the same

direction in which our hiker just travelled, which will be advantageous in following a straight ridge path. Both aspects allow our hiker to travel more quickly.

Look-Ahead Monte Carlo with People

Look-Ahead Monte Carlo with People (LAMCP) is a sequential paradigm that, like MCMCP, samples from a participant's distribution over stimuli for a particular category. Unlike MCMCP, LAMCP has two kinds of trials. The first kind are just like the trials of MCMCP; participants are asked to choose between a generated stimulus and the previous stimulus as the next state of a Markov chain. In the second kind of trial, however, participants are asked to choose a direction in the stimulus space to explore. This directional information is then used to generate a stimulus to be presented to the participant in the first kind of trial. Thus LAMCP produces two kinds of samples; stimuli, which we shall denote x , and also directions, which we shall denote with d .

LAMCP alternates between proposing stimuli using the current direction and previous stimulus and proposing directions using the current stimulus and previous direction. In this way, LAMCP is able to capture and to some extent remember local manifold structure when generating stimuli.

Recalling the analogy of the hiker trying to follow a ridge; new stimuli are generated from directions by starting at the current stimulus and advancing some distance according to the current direction. The distance advanced is sampled at random. More precisely, the two kinds of trials of LAMCP operate as follows:

Stimulus trial: Suppose that a direction d_t has been sampled already. The direction has an additive effect upon the current stimulus and so the proposal for the new stimulus, x_t^* , is:

$$x_t^* = x_{t-1} + \varepsilon_t d_t$$

where ε_t is a random value, sampled once for each stimulus trial, with distribution which we shall denote $q_\varepsilon(\varepsilon_t)$, determining for how far the direction d_t should be followed. The variable ε_t is the distance travelled, in direction d_t to obtain the new stimulus.

Participants are then presented with the previous stimulus x_{t-1} and the proposal x_t^* and asked which of these stimuli looked like they belonged more to the category of interest $\pi(x)$. As in MCMCP, we assume this choice is made by participants according to the ratio rule and thus corresponds to the Barker acceptance function (Equation 1).

Direction trial: In this step, participants pick a suitable direction for advancing the current stimulus x . A direction proposal is sampled from the direction proposal distribution $q_d(d_t^* | d_{t-1}, x_{t-1})$. Direction values are presented to participants as animations, showing how the proposed stimuli will be derived from the direction. This is the look-ahead part of our paradigm: this animation provides participants with insight into how selecting a direction would

affect the proposal of stimuli during future stimulus trials. Participants can decide to continue along a single direction for multiple stimulus trials by continuing to select the current direction during a direction trial.

Suppose N frames are to be generated for a proposed direction d_t^* . First N values of ε_t are selected, typically uniformly spaced within some task-specific interval: $\varepsilon_t^1, \varepsilon_t^2, \dots, \varepsilon_t^N$. Then frame $n \in \{1, \dots, N\}$ is the stimulus produced by $x_t + \varepsilon_t^n d_t^*$. The animation loops, first increasing the frame number n from 1 to N , then decreasing the frame number from N to 1.

Participants are presented with two animations: one corresponding to the proposed direction d_t^* and another for the previous direction d_{t-1} . Participants are asked to select the animation in which the stimuli look most like they belongs to the category of interest. By doing so, participants are picking the direction that is most likely to offer a stimulus belonging to the category during the next stimulus trial.

Again, as in MCMCP, we assume this choice is made by participants according to the ratio rule and thus corresponds to the Barker acceptance function (Equation 1).

In summary, the LAMCP paradigm is as follows:

1. Generate new direction proposal from direction proposal distribution q_d :

$$d_t^* \sim q_d(d_t^* | d_{t-1}) \quad (2)$$

The experimenter provides q_d : the distribution is required by the Barker acceptance function to be symmetric, $q_d(d^* | d) = q_d(d | d^*)$.

2. Generate an N -frame animation of the resulting samples: $x_t^* = x_{t-1} + \varepsilon_n d_t^*$ for each $n \in \{1, 2, \dots, N\}$.
3. Participants asked to choose between d_t^* and d_{t-1} for the new direction d_t , based upon the animated stimuli. In particular they select the new direction with probability:

$$\frac{\pi(d_t^* | x_{t-1})}{\pi(d_t^* | x_{t-1}) + \pi(d_{t-1} | x_{t-1})}$$

4. Generate new stimulus proposal from direction d_t and previous stimulus x_{t-1} :

$$\varepsilon_t \sim q_\varepsilon(\varepsilon_t) \quad (3)$$

$$x_t^* = x_{t-1} + \varepsilon_t d_t \quad (4)$$

The experimenter provides q_ε . Typically it will be a Gaussian distribution or a scale mixture of Gaussian distributions.

5. Participants are asked to choose between x_t^* and x_{t-1} as the new stimulus x_t , choosing x_t^* with probability:

$$\frac{\pi(x_t^*)}{\pi(x_t^*) + \pi(x_{t-1})}$$

In terms of Markov chains, LAMCP is a Metropolis sampler, where each kind of trial can be understood to be updates to either the stimulus or the direction. The overall stationary distribution is $\pi(x, d) = \pi(x)\pi(d|x)$. Participants are asked to pick directions whose animation show a high probability stimulus for the longest. Stated more formally, the stationary distribution for directions is:

$$\pi(d|x) \propto \int \int \pi(x^*)\delta(x^* = x + \epsilon d)dx^*d\epsilon \quad (5)$$

where x^* are the look-ahead points and $\pi(x^*)$ is the stimulus distribution evaluated at these look-ahead points, whilst ϵ is their distance along the direction d from the original stimulus x .

Extending LAMCP to Many Choices

As originally developed, MCMCP and LAMCP consisted of two alternative forced choice (2AFC) trials. This follows naturally from the Barker acceptance function (Equation 1). However, a larger number of alternatives could provide even more informative judgements.

At the start of each trial, instead of generating one proposal x_t^* , a set of n proposals is generated $\{x_t^p : p \in \{1, \dots, n\}\}$. Let c denote the index of the selected choice; instead of the Barker acceptance function in Equation 1, we use:

$$A(x_t^c, x_{t-1}) = \frac{\pi(x_t^c)}{\pi(x_{t-1}) + \sum_{p=1}^n \pi(x_t^p)} \quad (6)$$

Just as the Barker acceptance function (Equation 1) naturally arises in people’s decision making from the Luce’s choice axiom (Luce, 1963), so too does Equation 6.

It is not obvious, however, whether such an acceptance function leads to a valid Markov chain Monte Carlo sampler. By considering the multiple choice Markov transition kernel, one can show that the condition of *detailed balance* is satisfied by this acceptance function, and hence the above acceptance function maintains the equilibrium distribution $\pi(x)$. In particular, the Markov transition kernel is:

$$T(x \rightarrow x') = \sum_{p=1}^n \delta(x', x^p)A(x^p, x)S(x, x^p) + \delta(x', x) \left[1 - \sum_{p=1}^n \int A(x^p, x)S(x, x^p)dx^p \right] \quad (7)$$

where S is the symmetric proposal distribution for generating new proposals. Detailed balance requires one to show that $\pi(x)T(x \rightarrow x') = \pi(x')T(x' \rightarrow x)$ holds for all x and x' , which is easily achieved by substitution of definitions and algebra.

Though we have shown that the connection between the ratio rule and a valid acceptance function holds for more than two choices, there is evidence that people do not follow the ratio rule in this case (Wills, Reimers, Stewart, Suret, & McLaren, 2000). Rouder (2004) suggests that behaviour is better explained by raising the choice probabilities to a power

that varies from trial to trial. In our experiment, we shall only use this new regime for direction trials, where we do not need to sample from a stable distribution. In future, however, it would be interesting to explore how robust MCMCP is to the deviations from the ratio rule that people display.

Experiment: Testing LAMCP

We evaluated whether LAMCP is able to produce proposals that are more commonly accepted than MCMCP, generated higher quality samples than MCMCP, and use fewer trials to achieve the same quality of estimates as MCMCP. Each trial shall correspond to one decision made by a participant. For LAMCP, this means it will take two trials to produce one stimulus.

We applied LAMCP and MCMCP to a simple task where stimuli are parameterized by two dimensions—rectangles with width and height. Our aim is to elicit from participants samples from the category of golden rectangles, where the height is equal to the golden ratio (1.618) times the width. This scenario is a manifestation the ridge example that we motivated our approach with earlier. The golden ratio lies along a ridge in the two dimensional space form by all widths and heights.

Participants were asked to select the rectangle that looked most like a golden ratio rectangle on stimulus trials, or in the case of direction trials, to pick the animation that looked most like a golden ratio rectangle for the longest amount of time. We evaluated how their output deviated from ideal golden ratio rectangles.

Participants. A total of 43 participants were recruited from Amazon Mechanical Turk: each having at least a 95% task approval rate, and had at least 100 approved tasks. Each participant was required to contribute at least 100 decisions to the task to be paid \$0.05. Participants’ decisions were incorporated in real-time, and so some participants discontinued the experiment before 100 decisions but their results were still included (five participants for LAMCP, six for MCMCP24, and two for MCMCP).

Stimuli. Stimuli were black rectangles rendered in the participant’s web browser. Each stimulus was drawn in a 232 pixel by 232 pixel light grey box, with an internal border of 5 pixels. The light grey box had a border of 5 pixels surrounding it, and was on a white background. For animated stimuli, 25 frames were generated, and the frame was advanced, in a loop, changed approximately every 100 ms. For 2AFC trials, stimuli were shown side by side. For 4AFC trials, stimuli were shown in a grid of two rows and two columns. The stimuli parameters (width and height) were generated by either MCMCP or LAMCP, using truncated Gaussians for the initial stimulus and proposal distributions to ensure the parameters remained within the range 0 to 1. For MCMCP, the variance of the truncated Gaussians was randomly chosen at each trial to either be 0.01 or 0.25. For LAMCP, the variance

of the truncated Gaussians was randomly chosen at each trial to either be 0.1 or 0.5. Variance parameters for LAMCP were higher than those for MCMCP.

Procedure. Participants were asked to study 24 examples of rectangles, six of which were golden ratio rectangles and 18 non-golden ratio rectangles. Six of the counter-examples were the six examples rotated by 90 degrees, with explicit instructions highlighting that golden ratio rectangles are tall and thin, not short and wide.

We ran three regimes—MCMCP with just 2AFC trials, MCMCP with alternating 2AFC and 4AFC trials (to account for the effects of including alternating 2AFC/4AFC trials, which we shall call MCMCP24), and LAMCP with 2AFC for stimulus trials and 4AFC for direction trials. For each regime, we ran 10 chains, each chain being 100 trials long. Each trial consisted of choosing between either the stimulus selected by the previous trial and one or three new stimuli.

Results. The median acceptance rates of MCMCP and MCMCP24 were 38% ($\pm 2\%$; semi-interquartile range) and 37% ($\pm 4\%$), respectively, whilst the median acceptance rate of stimulus for LAMCP was 46% ($\pm 2\%$), suggesting that the proposals generated by LAMCP were typically more representative than those generated by MCMCP.

The median absolute difference between the estimated golden ratio and the true golden ratio was 0.72 (± 0.44 semi-interquartile range) and 0.92 (± 2.73) for MCMCP and MCMCP24, respectively. The median absolute difference of LAMCP, 0.52 (± 0.38) is closer to the golden ratio than both MCMCP methods. The absolute differences of MCMCP and MCMCP24 are significantly ($p < 0.001$) different to the absolute differences of LAMCP under the Wilcoxon rank-sum test.

Effective sample size is a heuristic for determining the number of independent samples yielded by an MCMC procedure. We compared effective sample size estimates of LAMCP, MCMCP and MCMCP24 using R-CODA (Plummer, Best, Cowles, & Vines, 2006). We found that the median effective sample size for LAMCP was 5 (± 1) whilst it was 11 (± 4) and 19 (± 15) for MCMCP and MCMCP24, respectively. This suggests that whilst LAMCP produces more favourable samples, the samples are more correlated with one another than those produced by MCMCP. This could be a consequence a linear correlation introduced by the generative process of the stimuli used by LAMCP. Interestingly the effective sample size for directional samples was 15 ± 3 with an acceptance rate of $75\% \pm 2\%$.

Figure 1 shows the estimated golden ratio and distance of samples to the golden ratio produced by MCMCP and LAMCP. The top row of Figure 1 shows that participants using LAMCP quickly find the golden ratio and are easily able to explore and follow this correlation in the stimulus parameters, compared to MCMCP participants. The bottom row of Figure 1 shows that throughout the evolution of both Markov

chains, LAMCP participants generate samples that are closer to the golden ratio than MCMCP participants.

We recorded the time between trials for each participant, and for LAMCP, compared the difference between these times for stimulus and direction trials. We could find no significant difference in these times under a variety of two-sample tests (Wilcoxon rank-sum, t-, and Kolmogorov-Smirnov tests; $p > 0.05$, $n = 500$). A likely explanation for this is that network transmission time dominates participants decision time: the median time between trials was 2.8 seconds (± 1 second semi-interquartile range). Thus for Mechanical Turk experiments, using directional trials do not appear to take more of a participants' time than stimulus trials.

Discussion

Look-Ahead Monte Carlo with People is an extension to MCMCP that exploits local manifold structure found in continuous valued stimuli by soliciting direction judgements from participants. This method will allow us to more efficiently explore and understand complicated categories in higher dimensions than previously attempted, by extracting more useful information per trial from participants. Whilst our simple experiment demonstrated the efficacy of our techniques in even the simplest case, for tasks such as images of faces, similar local structure likely exists in stimuli and so we can hope for similar gains.

Compared to other procedures for eliciting distributions from people, LAMCP's two kind of trial paradigm is also similar to iterated learning (Kirby, 1998; Griffiths & Kalish, 2007) where people either sample an internal representation or a physical manifestation of that representation. Directions are akin to an internal representation of what lies ahead, whilst the stimuli are the manifestations of the directions. The analogy is particularly apt if different people participate in each pair of LAMCP trials.

LAMCP not only produces samples from people's stimulus distribution, but also samples from their distribution over directions in stimulus space. Whilst our motivation for sampling from this distribution is purely incidental—we wish to obtain directions in some principled fashion so as to inform the stimulus generation process—the directions give a direct hint as to what people estimate to be the shape of the manifold in which samples lie. This extra piece of statistical information may be useful in gaining a better estimate of structure of stimuli, as well as aiding the estimation process.

Acknowledgements: We wish to thank Peter Dayan and Charles Sutton for several useful discussions.

References

- Ashby, F. G., & Alfonso-Reese, L. A. (1995). Categorization as probability density estimation. *Journal of Mathematical Psychology*, 39, 216-233.
- Barker, A. A. (1965). Monte Carlo calculations of the radial




Figure 1: Estimates of the golden ratio (top) and distance of samples to the golden ratio manifold (bottom) under three regimes: MCMCP (left), MCMCP24 (middle), and LAMCP (right). Heavy solid lines correspond to the median over 10 chains, whilst lighter solid lines are the interquartile range. Dashed straight lines in the top plots correspond to the golden ratio (top) and its reciprocal (bottom).

distribution functions for a proton-electron plasma. *Australian Journal of Physics*, 18, 119133.

Belisle, C. E., Romeijn, H. E., & Smith, R. L. (1993). Hit-and-run algorithms for generating multivariate distributions. *Mathematics of Operations Research*, 18(2).

Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive Science*, 31, 441–480.

Kirby, S. (1998). *Language evolution without natural selection: From vocabulary to syntax in a population of learners* (Tech. Rep.). Language Evolution and Computation Research Unit, University of Edinburgh.

Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology*, i (p. 103-189). Wiley.

Martin, J. B., Griffiths, T. L., & Sanborn, A. N. (2012). Testing the efficiency of Markov Chain Monte Carlo with People using facial affect categories. *Cognitive Science*, 36, 150-162.

Metropolis, A. W., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21, 10871092.

Neal, R. M. (1993). *Probabilistic inference using Markov chain Monte Carlo methods* (Tech. Rep. No. CRG-TR-93-1). Department of Computer Science, University of Toronto.

Neal, R. M. (1996). Bayesian learning for neural networks. *Lecture Notes in Statistics*, 118.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology*, 115(1), 39–57.

Plummer, M., Best, N., Cowles, K., & Vines, K. (2006). CODA: Convergence diagnostics and output analysis for MCMC. *R News*, 6(1), 7-11.

Rouder, J. N. (2004). Modeling the effects of choice-set size on the processing of letters and words. *Psychological Review*, 111, 80-93.

Sanborn, A. N., & Griffiths, T. L. (2008). Markov chain Monte Carlo with people. In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems 20* (pp. 1265–1272). MIT Press.

Sanborn, A. N., Griffiths, T. L., & Shiffrin, R. M. (2010). Uncovering mental representations with Markov Chain Monte Carlo. *Cognitive Psychology*, 60, 63-106.

Wills, A. J., Reimers, S., Stewart, N., Suret, M., & McLaren, I. P. L. (2000). Tests of the ratio rule in categorization. *The Quarterly Journal of Experimental Psychology*, 53(4), 983-1011.