Title
A Bayesian Analysis of Serial Reproduction

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 29(29)

ISSN
1069-7977

Authors
Xu, Jing
Griffiths, Thomas L.

Publication Date
2007

Peer reviewed

A Bayesian Analysis of Serial Reproduction

Jing Xu (jing.xu@berkeley.edu)
Thomas L. Griffiths (tom_griffiths@berkeley.edu)
Department of Psychology, 3210 Tolman Hall, MC 1650
University of California at Berkeley, Berkeley, CA 94720 USA

Keywords: serial reproduction; Bayesian modeling; memory; Markov chain; autoregressive time-series.

Bartlett (1932) explored the consequences of "serial reproduction" of information, in which one participant's reconstruction of a stimulus from memory becomes the stimulus seen by the next participant. These experiments were done using relatively uncontrolled stimuli such as pictures and stories, but suggested that serial reproduction could reveal the biases inherent in memory. We formally analyze the process of serial reproduction under a Bayesian account of memory, and test this approach with experiments using simple one-dimensional stimuli.

When people reconstruct a stimulus drawn from a category, they are influenced by the structure of the category. Huttenlocher, Hedges, and Vevea (2000) proposed that this effect can be modeled as a Bayesian inference, in which people combine the inexact fine-grain stimulus information with category information to achieve higher accuracy. We show that if this is the case, serial reproduction can be modeled as an autoregressive time-series, with a predictable trajectory and stationary distribution. Within the same theoretical framework, we also formally analyze how the convergence rate and stationary distribution of this process are influenced by different category distributions, perceptual noise, and different types of response behavior. Our analyses provide a formal justification for the idea that serial reproduction reflects memory biases.

Bayesian Inference

In our experiments, participants face the problem of reconstructing a one-dimensional stimulus from memory. They are trained that the stimuli they see come from a category with a Gaussian distribution over stimulus values, being $N(\mu_0,\sigma_0^2)$. At the *n*th trial of serial reproduction, a participant's noisy observation of a stimulus is $x_n \sim N(\mu,\sigma_x^2)$ where μ is a value drawn from the category distribution, and σ_x^2 determines the noise level. The participant then infers the true value μ by computing the posterior distribution $p(\mu \mid x_n)$ and stores a sample from it as a memory of x_n. The participant then generates x_{n+1} from memory, sampling from $p(x_{n+1} \mid \mu)$. When repeated by many participants, this process of serial reproduction constitutes a Markov chain, since the current trial depends only on the previous trial. The transition probability for going from one stimulus to the next is thus

$$p(x_{n+1} \mid x_n) = \int p(x_{n+1} \mid \mu)\, p(\mu \mid x_n)\, d\mu$$

and the stationary distribution is

$$p(x) = \int p(x \mid \mu)\, p(\mu)\, d\mu .$$

Autoregressive and Density Estimation

The serial reproduction process can be viewed as a first-order autoregressive AR(1) process where the current value depends only on the previous value. Using AR(1) time series analysis, we can find the transition probabilities for the Markov chain. At the (n+1)th iteration

$$x_{n+1} = (1-\lambda)\mu_0 + \lambda x_n + \varepsilon_{n+1}, \tag{1}$$

where $\lambda = 1/(1+\sigma_x^2/\sigma_0^2)$ and $\varepsilon_{n+1} \sim N(0, \sigma_x^2 + \sigma_n^2)$ with $\sigma_n^2 = \lambda\sigma_x^2$. That is, the transition probability is

$$x_{n+1} \mid x_n \sim N(\mu_n, \sigma_x^2 + \sigma_n^2)$$

where $\mu_n = \lambda x + (1-\lambda)\mu_0$. We can also find the stationary distribution of the serial reproduction chain by applying Equation 1 recursively to itself, and as $n \to \infty$

$$x_{n+1} \mid x_0 \sim N(\mu_0, \sigma_x^2 + \sigma_0^2) .$$

The rate at which the Markov chains converge to the stationary distribution depends on the value of λ. Since $\lambda = 1/(1+\sigma_x^2/\sigma_0^2)$, the convergence rate thus depends on the ratio of the participant's perceptual noise and the variance of the prior distribution σ_x^2/σ_0^2. We have tested these predictions through a series of experiments, implementing serial reproduction in the laboratory with a population of university undergraduates.

Acknowledgments

This work was supported by grant number 0704034 from the National Science Foundation.

References

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* Cambridge, England: Cambridge University Press.

Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, 129, 220-241.