

Inductive Biases Constrain Cumulative Cultural Evolution

Bill Thompson (billthompson@berkeley.edu)
Social Science Matrix, University of California, Berkeley

Thomas L. Griffiths (tgriffiths@princeton.edu)
Departments of Psychology and Computer Science, Princeton University

Abstract

Cumulative cultural evolution is a distinctively human form of information-processing that endows our societies with improbable and efficient technologies. But how objective is this process? A widely held conjecture is that human cognitive biases can constrain cumulative cultural evolution, and therefore shape our discoveries. We present a Bayesian analysis of a simple form of cumulative cultural evolution. This model allows us to formulate and test the theoretical conjecture in an experimental setting. Across a series of behavioural experiments, we show that people's inductive biases constrain a population's ability to discover counter-intuitive virtual technologies in a simple search problem. Our analysis highlights formal relationships between cumulative cultural evolution, Bayesian inference, and stochastic optimization.

Keywords: cumulative cultural evolution; inductive biases; optimization; computation; Bayes; cultural evolution;

Introduction

We are surrounded by bizarre and complex objects that vastly improve our lives. To our recent ancestors, many of the tools and technologies we rely on today were inconceivable, yet the same innovations will soon seem primitive to our descendants. The capacity for cumulative discovery is a uniquely human form of information processing on a breath-taking scale – but how objective is this process? Is technological evolution an unbiased search for optimal solutions to the problems we face? Or is it shaped by the same representational constraints and biases that limit individuals?

This question has been widely discussed in the context of *cultural evolution* (Mesoudi, 2016). Most theories of cultural evolution agree on the conjecture that in some circumstances, human cognitive biases must constrain cumulative cultural evolution (Morin, 2016; Acerbi & Mesoudi, 2015; Claidière, Scott-Phillips, & Sperber, 2014). This hypothesis has been widely debated and examined in formal models (Claidière & Sperber, 2007; Boyd & Richerson, 1985; Henrich & Boyd, 2002; Griffiths, Kalish, & Lewandowsky, 2008), but it has never been tested experimentally. In part, testing this hypothesis has been challenging because it is difficult to quantify an appropriate set of expectations in an experimental setting (Miton & Charbonneau, 2018).

In this paper, we develop a mathematical model of cumulative cultural evolution that allows us to formulate these expectations precisely. Our model is derived from a Bayesian analysis of individual cognition. The model makes quantitative predictions about the circumstances under which induc-

tive biases are likely to stifle discovery. To test these predictions, we adapt a widely studied experimental paradigm in which participants design and transmit a simple artificial technology: virtual arrowheads. Our strategy is to first characterise participants' inductive biases in this context using serial reproduction chain experiments. On the basis of these estimates, we conduct a series of arrowhead-design experiments which differ only in the extent to which task reward structure contradicts participants' biases. In the process of formalizing our predictions, we identify a formal relationship between cumulative cultural evolution and stochastic optimization.

Background

Cumulative Cultural Evolution

Unlike other species, every generation of humans builds on the insights and actions of their ancestors (Henrich, 2015). When an Apple engineer develops iPhone security updates, she makes use of cognitive resources expended by Turing almost a century before. In this sense, people alive today extend the computations initiated by people who faced similar problems in the past. What kind of process allows us to effectively pool computational resources with strangers over seemingly unbounded timescales? Computation over generations depends on a proclivity to learn from the people around us and the artefacts they create. When this kind of learning is repeated over time, a stochastic process is induced. This process is called *cultural transmission* (Boyd & Richerson, 1985). In some species, cultural transmission leads to cumulative innovation. This special case is known as *cumulative culture* (Mesoudi & Thornton, 2018) and is surprisingly rare (Whiten, Caldwell, & Mesoudi, 2016).

Discovering Technologies

There are many forms of cumulative culture, but one simple example has been heavily studied: refinement of technologies towards consistent functional objectives. Outside of the laboratory, examples of this process are easy to find: motorcycles today are faster, more efficient, more reliable, safer, and longer-lasting than the motorcycles people rode during the second World War. In an experimental setting, small-scale analogues of this process have been studied in several domains. For instance, Caldwell and Millen (2008) showed that micro-societies of experimental participants discover cumulatively more effective ways to design a tall-standing tower

of spaghetti. In these experiments, later participants observed the designs of earlier participants, and created taller and taller spaghetti towers as a result. Similar findings have been reported in lineages of participants designing simple knots (Muthukrishna, Shulman, Vasilescu, & Henrich, 2014), paper aeroplanes (Caldwell & Millen, 2008), rice baskets (Zwirner & Thornton, 2015), and fishing nets (Derex, Beugin, Godelle, & Raymond, 2013), for example.

Inductive Biases in Cumulative Culture

Cumulative cultural evolution can be recreated and manipulated in the laboratory. However, it remains unclear whether the products of these processes are shaped by biases in the way people think, or whether inductive biases and representational constraints are effectively washed out over time. This has been difficult to establish empirically, in part because it is often challenging to quantify the influence of people’s inductive biases. Recent reviews have noted that experimental tasks often feature unconstrained or difficult to quantify design spaces (Miton & Charbonneau, 2018), and that there is a need for a better understanding of the information-processing dynamics that link cognition and cultural evolution (Mesoudi & Thornton, 2018; Heyes, 2018). Mathematical analyses have repeatedly identified the potential for human biases to shape cumulative cultural evolution (Claidière & Sperber, 2007; Boyd & Richerson, 1985; Griffiths et al., 2008). However, extending abstract models to an experimental setting remains a challenge. Here, we introduce a formal model that is closely related to these theories of culture, but derived from a Bayesian analysis of cognition, and therefore directly applicable in an experimental context. Our analysis extends prior Bayesian models of cultural evolution (Griffiths & Kalish, 2007; Navarro, Perfors, Kary, Brown, & Donkin, 2018) to the cumulative case. The model we introduce allows us to specify formal predictions about the circumstances in which inductive biases constrain cumulative cultural evolution, and test those predictions experimentally.

Model: Optimization by Cumulative Culture

Our analysis applies to settings in which the design features of a technology can be described in terms of (n) continuous valued parameters $\Theta^t \in \mathbb{R}^n$. This setting offers a natural connection to prior experimental work, in which participants modify design features such as length, height, width, angles, crossing points, mass, or hue.

Induction of a Design

Each new individual estimates these design features from artefacts produced by the previous generation. If this estimation procedure can be given a formulation as Bayesian inference, then an individual’s estimate $\hat{\Theta}$ can be decomposed into a trade-off between two quantities: noisy empirical observation of the true design features Θ ; and inductive biases imposed by cognition. Inductive biases can be expressed as a prior distribution $p(\Theta)$. Using this framework, $\hat{\Theta}^t$ can be treated as

a random variable distributed according to the posterior distribution implied by a Bayesian model of learning.

Innovation

After estimating the existing design, each participant attempts an innovation. Assume a $f: \Theta \rightarrow \mathbb{R}$ is a function that reflects the utility of a technology with respect to its design features. In the literature on cultural evolution, this quantity would sometimes be referred to as a *fitness* landscape. We will make the assumption that individuals are capable of bounded, local innovation. This is appropriate to scenarios in which innovation is largely driven by an ability to identify similar but improved variants of whatever already exists, through limited experimentation with minor design variations for example. Local information about f can be naturally expressed as its gradient with respect to design features, evaluated at Θ^t . We denote this quantity $\nabla_f = \nabla_{\Theta} f(\Theta^t)$.

Diffusion Chains

These assumptions formalize a simple theory of cumulative culture as repeated cycles of observation, induction, and local innovation, leading to the expression:

$$\Theta^{t+1} = \Theta^t - \alpha \nabla_f - (\Theta^t - \hat{\Theta}^t), \quad (1)$$

where $t \in 1, \dots, T$ denotes a specific generation in a transmission chain. This equation describes a single step of a transmission chain in terms of the relationship between an existing technology (Θ^t), its utility (f), its status with respect to human cognitive constraints and the fidelity of transmission ($\Theta^t - \hat{\Theta}^t$), and an innovation rate (α). We examine the properties of this general model under some simplifying assumptions.

Assumption 1: Gaussian Prior & Observation Noise Assume individual learning can be modelled as probabilistic inference in a Gaussian model: observations of an existing design are noisy, and this noise can be approximated by Gaussian corruption of the true design features; inductive biases can be approximated by a Gaussian distribution.

Assumption 2: Independent Features Individual design features $\theta_i \in \Theta$ can be treated independently. This is a limiting assumption, but nonetheless appropriate to many relevant contexts. If μ_i is the prior expectation, the posterior expectation is:

$$\mathbb{E}[\hat{\theta}_i^t] = \lambda_i \mu_i + (1 - \lambda_i) \theta_i^t \quad (2)$$

where $\lambda_i = \sigma_i^2 / (\sigma_i^2 + \delta_i^2)$ reflects the relative variance of the prior (δ_i^2) and observation noise (σ_i^2) – in other words, the strength of an inductive bias $p(\theta_i)$ relative to the fidelity of transmission.

Chain Dynamics

The expected change at each generation can be written:

$$\mathbb{E}[\Theta_i^{t+1} - \Theta_i^t] = \lambda_i (\mu_i - \Theta_i^t) - \alpha \nabla_f. \quad (3)$$

which implies no further accumulation in expectation when $\nabla_f = \nabla_f^*$, where $\nabla_f^* \equiv (\mu_i - \theta_i^*)(\lambda_i/\alpha)$. This cultural process will halt if the potential for local innovation drops below a threshold determined by: the distance of the current design from the prior expectation ($\mu_i - \theta$), relative to a willingness to explore (α), weighted by the balance of prior and empirical leniencies in learning (δ_i^2/σ_i^2).

Assumption 3: Quadratic Utility Landscape The fate of the process is closely tied to the utility landscape in which it is operating. A simple but broad class of cases can be captured by the assumption that there is an optimal design Θ^* , and that f can be locally approximated by a quadratic surface with the optimum design at its minimum / maximum. In this regime, utility decreases with squared distance from the optimum at a rate proportional to a parameter a . A utility landscape that can be described in this manner has gradients $\nabla_f = a(\theta - \theta_i^*)$. A transmission process acting on a utility landscape of this form will halt in expectation if it reaches $\theta_i^t = \phi_i^*$:

$$\phi_i^* = \lambda_i^* \mu_i + (1 - \lambda_i^*) \theta_i^* \quad (4)$$

which is a linear combination of the prior mean μ_i and the optimum design θ_i^* with mixing proportions:

$$\lambda_i^* = \sigma_i^2 / (\sigma_i^2 + \alpha a s) \quad (5)$$

where $s = \delta_i^2 + \sigma_i^2$. Equations (4) and (5) represent our main theoretical result. Our analysis predicts that the outcome of a transmission chain is a compromise between the inductive biases of individuals and the optimal design. When these conflict, the balance of the compromise is quantifiable from the relationships between: transmission fidelity (σ_i^2), strength of inductive bias (δ_i^2), an exploration rate (α), and the slope of the utility landscape (a). The weighting factor $0 \leq \lambda_i \leq 1$ interpolates between cultural evolutionary processes that are constrained by inductive biases ($\lambda \rightarrow 1$) and therefore dragged back toward the prior $p(\theta_i)$, and processes that are dominated by information contained in the the utility landscape ($\lambda \rightarrow 0$), and therefore destined to discover an objective optimum.

Biased Computation by Cumulative Culture

One way to interpret this finding is as a description of the computation that is being implemented by the *process* we have analysed – the computation implemented by a chain of individuals. Two analogies motivate this interpretation. First, equation (5) has the same form as equation (4). At each generation, an individual person performs a computation that we formalized as a sample from the posterior distribution in a Bayesian model of inference. This computation is biased and *local*: the expectation is a linear combination of the individual’s inductive bias and the *currently existing* design θ^t . However, the chain *as a whole* can be understood to implement a biased but *global* computation: equation (5) describes the expectation of a posterior distribution computed by the same kind of learner after observing (a noisy realisation of) the *op-*

timal design θ^* . Second, in the Gaussian case, equation (1) can be rewritten as:

$$\theta_i^{t+1} = \theta_i^t - \alpha \nabla_f - \lambda_i (\theta_i^t - \mu_i) + \varepsilon_i^t \quad (6)$$

which is a form of stochastic gradient descent with regularisation. Stochastic optimization and Bayesian inference are known to be related (Mandt, Hoffman, & Blei, 2017). This highlights a common interpretation of cognitive and cultural processes – they are both forms of information processing. This cultural process solves an optimization problem subject to regularisation by human inductive biases. In the remainder of this paper, we test this prediction.

Experiment: Discovering Virtual Technologies

We adapted an experimental paradigm that has been widely used to study the influence of social learning on cumulative culture (Mesoudi, Chang, Murray, & Lu, 2015). The experimental task involves designing a virtual arrowhead. The arrowhead has a number of attributes (e.g. length, width) that can be modified and achieves a score when deployed on a virtual hunt. The score reflects the number of calories of food earned by the arrowhead. Participants’ goal was simply to test and redesign a single arrowhead they inherited, in an attempt to increase its score. This paradigm allowed us to construct a low-dimensional search problem in which we hypothesised that task-naive participants would display biased expectations. Although there is significant discussion surrounding the definition of cumulative cultural evolution, a central requirement is that over time, it’s products must “enhance some measure of performance...[through]...sequential improvements...” (Mesoudi & Thornton, 2018). In our experiment, unconstrained cumulative cultural evolution would correspond to the chains of participants sequentially designing arrowheads that achieve higher scores until the maximum score is achieved.

Method

Stimuli The experiment was presented as a website. Participants designed a virtual arrowhead using two HTML range sliders which modified its width and length. Figure 1 shows the design space. During experimental trials, the screen was split into left (25% screen width) and right (75% screen width) panels. The left panel displayed the participant’s estimate of the arrowhead they inherited in reduced proportions that nonetheless preserved the design. Underneath was a depiction of two range-sliders positioned in accordance with the arrowhead’s attributes, and text indicating the number of calories that the arrowhead earned. The main panel (Right) displayed an arrowhead in the center, pointing downwards. Two range-sliders were located beneath the arrowhead. Moving a range-slider modified either the length or the width of the arrowhead dynamically. Both dimensions of the arrowhead could take values ranging between 50 and 150 pixels.

Arrowhead scores were determined by a quadratic function of the form $f(\theta_i) = \frac{1}{2}a(\theta_i - \theta_i^*)^2 + c$, where θ_i^* is the




Figure 1: Regularly sampled virtual arrowheads in a design space implied by the ability to modify two features – length and base width.

optimum value for feature θ_i . We chose to use this family of functions because it constructs a smooth utility landscape with a single optimum. Previous work has focused on conditions that allow populations to search effectively through more complex landscapes with both local and global optima (Acerbi, Tennie, & Mesoudi, 2016). In contrast, our analysis focuses on a problem that should be relatively easy to solve optimally if people’s inductive biases do not constrain cultural evolution. Quadratic functions are also the class of landscapes we examined in our theoretical analysis.

Within the family of quadratic landscapes, we required a function which was: smooth over the full range; did not return negative scores; had an optimum (maximum) score that lies within a semantically reasonable range given the framing of the task. To meet these requirements, we set $\alpha = -30$ and $c = 10000$ for all experiments and divided the result by 100. Given our settings of θ_i^* (see below), the maximum available score was 1000 calories. Calories decreased away from θ_i^* at a rate given by $\nabla_f = -\frac{3}{10}(\theta_i^t - \theta_i^*)$. To award a score, we computed f for both arrowhead features (length and base width) and awarded the mean.

Procedure Participants were informed that: they would go on a virtual hunt; their task was to design a virtual arrowhead that will earn as many calories of food as possible on the hunt; a bonus payment would be made in proportion to the number of calories their arrowhead earned. After consenting to the experiment, participants completed an Information Trial (IT), during which they observed the arrowhead they had inherited (first generation participants inherited an arrowhead with a randomly sampled design). This arrowhead was displayed in the center of the right panel for 3000 milliseconds. Participants then recreated the arrowhead as accurately as possible. Participants then proceeded to the first Modification Trial (MT). The participant’s estimate of the arrowhead design was displayed as the arrowhead they inherited in the left

panel. Participants completed four MTs. During each MT, $\hat{\theta}_i^t$ was displayed in the left panel. At the beginning of an MT, no arrowhead was displayed in the right panel, and the positions of the range-sliders were randomised. Participants could not proceed to the next trial until at least one range-slider had been modified. Upon modifying any of the range-sliders positions, the arrowhead was redrawn. Modifications were limited to a range of ± 30 pixels around $\hat{\theta}_i^t$. This enforced a weak restriction on the innovations participants could make in accordance with our theoretical model. There was no limit to the number of times participants could modify the range-sliders in a given MT. Once satisfied, participants could click Submit to obtain feedback – the number of calories earned by the current arrowhead design. Arrowheads evaluated during previous MTs were displayed (in reduced proportions) in the left panel, in trial order. After completing four MTs, participants were informed that their opportunity to test arrowheads was complete and proceeded to the test trial (TT). Participants were reminded that the arrowhead they designed during this trial would determine a bonus payment. TT was identical to MT in all other respects.

Participants Participants ($n = 1000$) were recruited online using Amazon’s Mechanical Turk. The experimental protocol was approved by the University of California, Berkeley’s Committee for the Protection of human Subjects. Participants were paid \$0.50 to complete the experiment, and awarded a performance-based bonus of up to \$0.50. Most participants completed the experiment in less than three minutes. Data from any participants who completed the experiment in less than 20 seconds were rejected.




Figure 2: Participant view of the experiment. Screenshot shows the fourth Modification Trial. At this point in the experiment, the participant has completed the Information trial (and recreated their inherited arrowhead, shown first in the left panel) and three Modification Trials (the arrowheads tested by the participant so far and their scores are shown in the left panel).




Figure 3: All arrowheads produced by participants at generation 3 or later in 20 serial reproduction chains of 10 generations each. A sample based approximation to participants’ inductive biases, $p(\theta)$.

Results

Reproduction Chains We first ran a simpler experiment using the same stimuli. This experiment used *serial reproduction chains* to characterise participants’ inductive biases in our stimulus set. In these chains, each participant completed an IT, but did not proceed to MT and TT. Each participant observed the arrowhead designed by the previous participant, and was asked to reproduce it as accurately as possible. Previous mathematical (Griffiths & Kalish, 2007) and experimental (Griffiths et al., 2008; Xu & Griffiths, 2010) research has established that serial reproduction chains characterise participants’ inductive biases. Figure 3 shows the distribution of arrowheads produced by all participants at generation 3 or later (in all our analyses, the first two generations of a chain are excluded as *burn-in* generations to minimise the effects of random initial conditions), across 20 serial reproduction chains of 10 generations each. This collection of arrowheads can be understood as a sample-based approximation to the prior distribution $p(\Theta)$. The empirical mean of this distribution is $\hat{\mu}_{width} = 111$, $\hat{\mu}_{length} = 118$. People favoured arrowheads that are relatively long and relatively wide.

Optimization Chains In light of participants’ inductive biases, we conducted four experiments. We treat these as separate experiments rather than experimental conditions because they were carried out sequentially. Each experiment (20 chains of 10 generations) featured a utility landscape with a differently located optimum but the same calorie gradient surface. Our prediction was that differently located optimums would lead to differential discovery of those designs, and differential task success (number of calories). Our mathematical

analysis identified the distance between the optimum arrowhead and the mean of the prior distribution $p(\theta)$ as the crucial predictive quantity: optimal arrowheads that are farther from the prior distribution should be harder to find because they are less intuitive. Figure 4 shows our results. Experiment 1 ($\theta_{width}^* = 115$, $\theta_{length}^* = 115$) was designed to be most consistent with people’s inductive biases. In this experiment, the arrowheads people designed scored well (mean calories $M = 948$, $SD = 46$). Experiment 4 was least consistent with people’s biases, contradicting people’s expectations in both dimensions. Success in the task suffered as a result ($M = 794$, $SD = 183$). Experiments 2 ($\theta_{width}^* = 75$, $\theta_{length}^* = 115$, $M = 912$, $SD = 98$) and 3 ($\theta_{width}^* = 115$, $\theta_{length}^* = 75$, $M = 845$, $SD = 132$) were designed to contrast with people’s biases in one of the two dimensions – width and length respectively.

We combined data from all four experiments and computed the difference between the optimum arrowhead and the mean of the prior distribution. The main prediction of our formal model (equation 4) can be rearranged into a linear model of the form $\phi_i^* = \hat{\mu}_i + \beta(\theta_i^* - \hat{\mu}_i)$. This allowed us to perform an ordinary least squares regression analysis of this model in our experimental data. The prediction was upheld. Accounting for the mean of the prior distribution ($\hat{\beta} = 1.0$, $p < .001$) and the difference between the mean of the prior and the optimum design ($\hat{\beta} = 0.42$, $p < .001$) accounted 96% of the variance in the features of the arrowheads people produced ($R^2 = 0.962$). We also analysed task success, and found significant differences in the distribution of arrowhead scores in all pairwise comparisons of our four experiments (at $\alpha = .05$). Only the comparison between experiments 3 and 4 ($t(197) = 3.2$, $p = 0.0017$) was not significant at $\alpha = .001$. Figure 4 (b) shows how task success reduced over the four experiments. Finally, we computed the predictions of our mathematical model under the inferred mixing proportions λ explicitly. Figure 4 (c) shows these predictions.

Conclusion

We introduced a simple formal theory of cumulative cultural evolution. We used this theory to predict how the inductive biases of individuals would constrain a cultural process. We tested this prediction in a series behavioural experiments. We found that discovery of an optimal virtual technology in a simple search problem was impeded by people’s inductive biases. These results reinforce a theoretical conjecture that had previously not been studied empirically. Our analysis highlighted formal connections between cumulative cultural evolution, Bayesian inference, and stochastic optimization. Our results suggest a more general insight: identifying the algorithm that is implemented by a cultural process can allow us to characterise the computation it performs, yielding a cognitive interpretation of the process in information-processing terms. Our results showed that computation by cumulative culture can be biased. This naturally raises the question: under what circumstances is computation by culture *unbiased*?




Figure 4: Results of all optimization experiments. Above (a): kernel density estimates of the distribution of arrowheads produced by all participants at generation 3 or later. Dotted lines show the location of the optimum arrowhead in design space. Middle (b): The distribution of scores obtained by the arrowheads produced by participants at generation 3 or later in all optimization chains. The Euclidean distance between the optimum arrowhead and the mean of the prior distribution predicts task success (No. calories). Below (c): Mean arrowhead design produced by all participants at generation 3 or later in all optimization chains (yellow cross), alongside the mean design predicted by our mathematical model (after fitting λ to experimental data, blue circles), the empirical mean of the prior distribution ($\hat{\mu}$, green plus), and the experiment-specific optimum design (black arrowhead).

Acknowledgements

This work was funded in part by NSF grant 1456709 and DARPA Cooperative Agreement D17AC00004.

References

- Acerbi, A., & Mesoudi, A. (2015). If we are all cultural Darwinians what's the fuss about? Clarifying recent disagreements in the field of cultural evolution. *Biology & Philosophy*, 30(4), 481–503. doi: 10.1007/s10539-015-9490-2
- Acerbi, A., Tennie, C., & Mesoudi, A. (2016). Social learning solves the problem of narrow-peaked search landscapes: experimental evidence in humans. *Royal Society Open Science*, 3(9), 160215. doi: 10.1098/rsos.160215
- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago, IL: University of Chicago Press.
- Caldwell, C. A., & Millen, A. E. (2008). Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, 29(3), 165–171. doi: 10.1016/J.EVOLHUMBEHAV.2007.12.001
- Claidière, N., Scott-Phillips, T. C., & Sperber, D. (2014). How Darwinian is cultural evolution? *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 369(1642), 20130368. doi: 10.1098/rstb.2013.0368
- Claidière, N., & Sperber, D. (2007). The role of attraction in cultural evolution. *Journal of Cognition and Culture*, 7(1), 89–111.
- Dere, M., Beugin, M.-P., Godelle, B., & Raymond, M. (2013). Experimental evidence for the influence of group size on cultural complexity. *Nature*, 503(7476), 389–391. doi: 10.1038/nature12774
- Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with bayesian agents. *Cognitive science*, 31(3), 441–80. doi: 10.1080/15326900701326576
- Griffiths, T. L., Kalish, M. L., & Lewandowsky, S. (2008). Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 363(1509), 3503–14. doi: 10.1098/rstb.2008.0146
- Henrich, J. (2015). *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
- Henrich, J., & Boyd, R. (2002). On Modeling Cognition and Culture: Why cultural evolution does not require replication of representations. *Journal of Cognition and Culture*, 2(2), 87–112. doi: 10.1163/156853702320281836
- Heyes, C. (2018). Enquire within: cultural evolution and cognitive science. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 373(1743), 20170051. doi: 10.1098/rstb.2017.0051
- Mandt, S., Hoffman, M. D., & Blei, D. M. (2017). Stochastic Gradient Descent as Approximate Bayesian Inference.

- Mesoudi, A. (2016). Cultural evolution: integrating psychology, evolution and culture. *Current Opinion in Psychology*, 7, 17–22. doi: 10.1016/J.COPSYC.2015.07.001
- Mesoudi, A., Chang, L., Murray, K., & Lu, H. J. (2015). Higher frequency of social learning in China than in the West shows cultural variation in the dynamics of cultural evolution. *Proceedings. Biological sciences*, 282(1798), 20142209. doi: 10.1098/rspb.2014.2209
- Mesoudi, A., & Thornton, A. (2018). What is cumulative cultural evolution? *Proceedings. Biological sciences*, 285(1880), 20180712. doi: 10.1098/rspb.2018.0712
- Miton, H., & Charbonneau, M. (2018). Cumulative culture in the laboratory: Methodological and theoretical challenges. *Proceedings of the Royal Society B: Biological Sciences*, 285(1879), 20180677. doi: 10.1098/rspb.2018.0677
- Morin, O. (2016). Reasons to be fussy about cultural evolution. *Biology and Philosophy*, 31(3). doi: 10.1007/s10539-016-9516-4
- Muthukrishna, M., Shulman, B. W., Vasilescu, V., & Henrich, J. (2014). Sociality influences cultural complexity. *Proceedings. Biological sciences*, 281(1774), 20132511. doi: 10.1098/rspb.2013.2511
- Navarro, D. J., Perfors, A., Kary, A., Brown, S. D., & Donkin, C. (2018). When Extremists Win: Cultural Transmission Via Iterated Learning When Populations Are Heterogeneous. *Cognitive Science*. doi: 10.1111/cogs.12667
- Whiten, A., Caldwell, C. A., & Mesoudi, A. (2016). Cultural diffusion in humans and other animals. *Current Opinion in Psychology*, 8, 15–21. doi: 10.1016/J.COPSYC.2015.09.002
- Xu, J., & Griffiths, T. L. (2010). A rational analysis of the effects of memory biases on serial reproduction. *Cognitive Psychology*, 60(2), 107–126. doi: 10.1016/J.COGPSYCH.2009.09.002
- Zwirner, E., & Thornton, A. (2015). Cognitive requirements of cumulative culture: teaching is useful but not essential. *Scientific reports*, 5, 16781. doi: 10.1038/srep16781