

Evolution in Mind:

Evolutionary Dynamics, Cognitive Processes, and Bayesian Inference

Jordan W. Suchow, David D. Bourgin, and Thomas L. Griffiths

Department of Psychology,

University of California, Berkeley

Correspondence: suchow@berkeley.edu (J.W. Suchow).

Abstract

Evolutionary theory describes the dynamics of population change in settings affected by reproduction, selection, mutation, and drift. In the context of human cognition, evolutionary theory is most often invoked to explain the origins of capacities such as language, metacognition, and spatial reasoning, framing them as functional adaptations to an ancestral environment. However, evolutionary theory is useful for understanding the mind in a second way: as a mathematical framework for describing evolving populations of thoughts, ideas, and memories within a single mind. In fact, deep correspondences exist between the mathematics of evolution and of learning, with perhaps the deepest being an equivalence between certain evolutionary dynamics and Bayesian inference. This equivalence permits reinterpretation of evolutionary processes as algorithms for Bayesian inference and has relevance for understanding diverse cognitive capacities, including memory and creativity.

Keywords: evolution, cognitive processes, learning, memory, creativity, Bayesian inference

Linking evolution and cognition

In settings as diverse as cancer biology and color vision, evolutionary dynamics provides a unifying mathematical framework for understanding how populations of replicators evolve when subject to mutation, selection, and random drift [1]. The explanatory power of evolutionary dynamics is not limited to biology and medicine — from the notion of a “meme” [2–4] to influential models of cultural evolution [5–8], evolutionary concepts such as selection and mutation are often appealed to in explaining cultural phenomena such as cooperation [9, 10], social learning, and language change [11–14].

As evolutionary dynamics does for changes in populations, Bayesian inference provides a unifying mathematical framework for describing changes in beliefs [15]. When acquiring a language, learning a concept, or inferring a causal mechanism, the mind is faced with an inductive inference problem and must transform incomplete information into updated thoughts, beliefs, or knowledge. To learn the meaning of a new word, for example, the learner must integrate prior knowledge of plausible word meanings with data in the form of a few observed examples of the word’s usage, a process well described as Bayesian inference [16].

Evolutionary dynamics and Bayesian inference are thus two frameworks for describing change: changes to a population, and changes to beliefs. What happens, then, when the population under consideration is not external to the mind, but within it? In this Opinion, we argue that an equivalence between evolutionary dynamics and Bayesian inference enables reinterpretation of evolutionary processes as algorithms for inductive inference. From this vantage point, evolutionary theory is promoted from a metaphor for the inner workings of the mind to a mathematical framework that can explain the dynamics of diverse cognitive processes. After two case studies showing that evolutionary processes over mental representations have

interpretations as algorithms for Bayesian inference, we close with thoughts on what is gained, and what stands to be gained, by keeping evolution in mind.

Evolutionary dynamics and Bayesian inference

On the surface, evolution and learning have much in common — both can be viewed as optimization processes, finding the best among a set of alternatives. (Strictly speaking, optimization applies only to evolution with haploid organisms, a single locus, and no mutation. Under these circumstances, the replicator dynamics implements gradient descent on the mean fitness of the population.) Mathematical biologists have sometimes pointed out that the theoretical tools they use to study evolution are just as relevant to other adaptive processes. For instance, what came to be known as the Price equation, a mathematical description of how a trait's frequency changes in response to selective pressures, was noted as having possible application to students improving their mathematical abilities through instruction [17]. But there are also deeper correspondences, and one of the deepest is the link between evolutionary dynamics and Bayesian inference.

A standard mathematical model of evolution is the *replicator dynamics*. In the simplest version of this model, generations are assumed to be discrete, and each agent produces a number of offspring proportional to its fitness. Letting $p_i^{(t)}$ denote the proportion of a population corresponding to variant i at time t and f_i denote the fitness of variant i , the replicator dynamics assumes

$$p_i^{(t+1)} = \frac{f_i p_i^{(t)}}{\bar{f}}, \quad (1)$$

where \bar{f} is the mean fitness across all variants, $\bar{f} = \sum_j f_j p_j^{(t)}$. This equation indicates that the proportion of the population of type i will increase if type i has greater than average fitness.

Equation 1 should look familiar to those who craft models of human cognition. In fact, it takes exactly the same form as Bayes' rule, which stipulates how a rational agent should update his or her beliefs in light of evidence. If each hypothesis f_i is assigned a *prior* probability $p(h_i)$ before observing data d , then the *posterior* probability $p(h_i|d)$ is given by

$$p(h_i|d) = \frac{p(d|h_i)p(h_i)}{\sum_j p(d|h_j)p(h_j)}, \quad (2)$$

which has the same interpretation as Equation 1: hypothesis i will increase in probability if the probability of the data d under that hypothesis is greater than the average over all hypotheses. If we take proportions of a population of hypotheses to correspond to degrees of belief in those hypotheses ($p_i^{(t)} = p(h_i)$) and take the probability of observed data under each hypothesis to correspond to its fitness ($f_i = p(d|h_i)$), Equations 1 and 2 are isomorphic [18, 19].

Given the link between the replicator dynamics and Bayesian inference, it is sensible to ask whether proposals concerning evolutionary processes as cognitive processes might be reinterpreted in terms of algorithms for Bayesian inference. We explore the consequences of this idea in two case studies of human cognitive processes: memory and creativity. In the first, a population of memories evolve as they are maintained in working memory. In the second, a population of ideas evolve in a creative process.

Memory

Working memory is a core component of our cognitive architecture: it holds and manipulates information in mind, providing a workspace for thought [20]. The capacity of working memory is limited to only a handful of objects or events. Across individuals, working-memory capacity is positively correlated with general intelligence, suggesting a key role in higher-order cognition

[21]. Recent research has provided detailed insight into the dynamics of working memory, and suggested that they have parallels with evolutionary dynamics.

Memory maintenance as an evolutionary process

Since the 1950s, theories of working memory have posited a population of discrete, flexible representations that support ongoing maintenance of information held in mind. These representations go by many names — stimulus samples [22], chunks [23], slots [24], and resources [25], among others [26] — and typically take form as a limited commodity, at least partially shared across memories, whose availability affects our ability to maintain information in working memory. The commodity may be instantiated as (for example) cycles of a time-based refreshing process [27] or populations of neurons in prefrontal cortex representing “token” encodings of visual events [28].

The dynamics of this commodity — this population of representations in the mind — can be considered in the context of evolutionary theory. In the evolutionary interpretation, selection is a cognitive control process that redirects maintenance, giving preference to some information over others, perhaps in accordance with the demands of the task or the salience of the held information [29, 30]. Drift is a process of forgetting by which memories degrade, perhaps due to time or interference, losing fidelity or becoming misattributed. And mutation alters the contents of memory, causing what we remember to become unhinged from our experiences.

Considering working memory in this evolutionary framework can account for diverse effects related to the capacity of working memory and the dynamics of forgetting [30, 31]. For example, neutral drift or a stochastic pure-death process may explain why the fidelity of working memory varies stochastically across objects [32]. And a generalization of the Moran process (see Box 1) outfitted with a stability threshold can capture the full time-course of the forgetting

functions of working memory, including exponential-like decay, the dependence of decay rate on the number of objects, and the curious fact that the asymptote is 1 object (Fig. 2). However, the relationship between evolutionary dynamics and Bayesian inference outlined above offers a different way to understand the success of these models.

The Wright-Fisher model as a particle filter

An interesting relationship can be found between the Wright-Fisher model of evolutionary dynamics in a finite population (Box 1) and the particle filter, a sample-based Monte Carlo inference algorithm that estimates a posterior probability distribution through maintenance of a discrete sample-based approximation that is updated over time [33]. The particle filter starts with a set of particles (hypotheses) sampled from the prior distribution $p(h_i)$. (There is the possibility that the set of particles will contain several copies of a hypothesis if it has a high prior probability, or if there are few hypotheses or many particles.) Upon observing data, each particle (hypothesis) is assigned a weight proportional to the probability of observing the data given that hypothesis, $p(d|h_i)$. These weights are then normalized (i.e., divided by their sum, so that the weights across all hypotheses sum to 1) and used to sample a new set of particles. Each particle in the new set is sampled from the previous set in proportion to their normalized weights. The new hypotheses will reflect both the original proportions of hypotheses and the information contained in the data. This can be seen as being formally equivalent to the Wright-Fisher model with selection, where the “fitness” of each hypothesis corresponds to the probability of the observed data under that hypotheses, as in the general correspondence between the replicator dynamics and Bayesian inference given above.

The Wright-Fisher model and the Moran model, two ways of simulating finite-population evolutionary dynamics (see Box 1), yield similar conclusions in a wide range of situations [34].

This suggests that we might be able to reinterpret the success of the Moran process as a model of human memory in terms of the solution to a probability estimation problem. Particle filters have already been proposed as models of cognitive processes in other settings [35, 36].

The estimation problem faced by a memory maintenance system is that of representing the environment in a way that is useful for the planning and execution of future action. In the approximate solution to this problem offered by the particle filter, that representation is a set of particles (i.e., sampled hypotheses). Consider, for example, a single trial of a typical memory experiment where a participant is asked to hold in mind a complex shape that is briefly presented. Initially, because direct exposure to the object provides strong evidence in favor of the hypothesis that the object has a particular shape, the filter quickly encodes that hypothesis into memory. However, as time progresses beyond the initial exposure, the circumstances rapidly change because the input ceases to provide new relevant data. In the absence of evidence, the process becomes one of neutral drift.

These features of the particle filter can explain several important features of the encoding and forgetting functions of visual memory within a unified framework. Because particles are initialized with hypotheses sampled from the prior, encoding of information into memory is not instantaneous and will depend on the strength of evidence present in the input, explaining graded benefits of increased encoding time on the fidelity of memory representations [37]. Assuming a fixed information capacity of a particle, increasing the number of objects will decrease the fidelity of any one object's representation across the particles, explaining the fragility and quick decay at higher loads (Fig. 2). Most of all, framing the process of memory maintenance as a particle filter makes it possible to break out of the trial-based nature of typical laboratory

experiments and to consider use of the working memory system in our everyday life, with its stream of experiences embedded in continuous time.

Creativity

The capacity to generate new and useful ideas is one of the defining characteristics of human cognition. It has also been frequently described using ideas borrowed from biological evolution.

The evolution of thoughts

One of the first examples of an analogy between creativity and evolution comes from William James in his 1880 lecture “Great Men and Their Environment”, where he outlines a correspondence between “zoological evolution” in nature and the evolution of ideas during creative synthesis:

"...throughout the whole extent of those mental departments which are highest, which are most characteristically human ... new conceptions, emotions, and active tendencies which evolve are originally produced in the shape of random images, fancies, accidental out-births of spontaneous variation in the functional activity of the excessively instable human brain, which the outer environment simply confirms or refutes, adopts or rejects, preserves or destroys — selects, in short, just as it selects morphological and social variations due to molecular accidents of an analogous sort." [38]

The view of creative ideation as an evolutionary process defined over a population of ideas was expressed more precisely in the Blind Variation and Selective Retention theory (BVSR) [39]. The theory characterizes human creative activity using a two-stage model in which ideational variants were first generated at random (blind variation) and then culled according to their utility within a given domain (selective retention). The utility of this perspective was seen to be primarily metaphysical: just as natural selection did for biological evolution, the BVSR

model offered an explanation for a seemingly purposive process without appealing to dubious teleological arguments. Indeed, these ideas were later developed into a theory of sociocultural evolution and evolutionary epistemology that sought to apply the principles of blind variation and selective retention to the philosophy of science [39, 40].

The evolutionary perspective has enjoyed significant attention within the creativity literature, most notably in work formalizing the mechanisms implicated at each stage of the BVSR model and quantitatively evaluating the model's predictions [41, 42]. This updated theory offers an extended review of the notion of 'blindness' during generation, ultimately leaving room for ideational variants to be generated along a blind-sighted continuum, with both internal and external selective mechanisms occurring sequentially and in parallel according to the nature of the domain and task [41, 42]. These clarifications and additions have permitted the BVSR's predictions to be tested empirically against collections of historical and behavioral data, providing support for many of its claims. As a result, the evolutionary perspective has been described as possessing "a rigor that is unsurpassed by any other major theory of creativity" [43], and has been successfully used to model scientific and artistic career trajectories [44], the acquisition and development of creative expertise [45], and ratios of creative success to total output [46], amongst other phenomena.

Blind variation and selective retention as stochastic search

In the analogy to evolution made in the case of human creativity, the notion of a population of ideas competing against one another is replaced with an account of how individual ideas are generated and retained. Consideration of equivalencies between evolutionary processes and Bayesian inference suggests a different way to think about how this blind variation and

selective retention process can be applied in the context of creativity: as elements of a stochastic search algorithm.

Markov chain Monte Carlo (MCMC) algorithms are one of the most commonly used methods for stochastic search. Algorithms in this family excel at exploring state spaces that would otherwise be intractable to systematically investigate [47]. The algorithms achieve this feat by offering a method for defining a stochastic process (a Markov chain) over a state space which can be used to generate samples in proportion to a desired probability distribution. This distribution can subsequently be approximated using a finite number of samples from this process (the Monte Carlo principle). This technique has enjoyed success in cognitive modeling as a bridge between optimal but computationally intractable Bayesian accounts of cognition and psychologically plausible approximations which retain many of the same performance guarantees [35].

The evolutionary account of creativity suggests a two-stage search model incorporating both generation and selection components. The Metropolis-Hastings algorithm [48, 49] is an MCMC search algorithm which exhibits such a structure: at each step of the search it proposes a new state based on the value of its current state, and then makes a decision about whether to retain this proposed state based on the utility ratio of the proposal to its current state (Box 2). Critically, the proposal stage proceeds in a way that is independent of the utility of the possibilities, and the retention stage uses the utility function to select proposals in proportion to their overall quality. In the context of creativity, the states that the search operates over might correspond to solutions to a difficult problem or ideas for a creative work. Initial support for such an account of creative search is provided by work inspired by BVS theory in which a Metropolis-Hastings process on a semantic network is used to simulate creative search during

problem solving (Fig. 3) [50]. When the model is evaluated against human responses on the Remote Associates Test (RAT) [51], a verbal test of creativity which required participants to come up with a word which related three seemingly independent cues to one another (e.g., generating the response ‘party’ to relate the cues ‘surprise’, ‘line’, and ‘birthday’), the stochastic search model reproduced many of the patterns present in the intermediate responses generated by humans as they worked on the RAT questions. These include significant correlations between human and model accuracies, local dependencies between responses, and undirected search trajectories through the state space which did not appear to differentially “hone in” on the correct answer [50]. Most importantly, these results serve to highlight the link between creative ideation, evolutionary processes in the mind, and (approximate) Bayesian inference.

Concluding remarks and future perspectives

As illustrated by the account of creativity offered by William James, the suggestion of parallels between evolution and cognition date back almost as far as the idea of natural selection itself. In the years since, both evolutionary biology and cognitive science have made progress in formalizing their respective objects of study. As a consequence, it has become possible to create models of human cognition based on mathematical models from evolutionary biology and consider how those models are related to other frameworks for understanding human cognition. In particular, there is a deep connection between evolutionary dynamics and Bayesian inference that can be used to join these two perspectives.

Rather than suggesting that one perspective should simply be reduced to the other, we see value in being able to go back and forth between these two ways of thinking about cognitive processes. Theoretical evolutionary biology and Bayesian statistics are two quite distinct literatures, offering complementary insights about the mind. For example, the Price equation [17]

is a fundamental observation about how the relationship between the traits of organisms and their fitness influences the progress of selection. Translated into a cognitive context, it indicates how the relationship between properties of hypotheses and the probability of observing data under those hypotheses will govern the degree to which beliefs about those properties will change. Biologists use the Price equation to prove abstract results about the nature of evolution, suggesting that a similar approach could be applied to understanding learning (and Bayesian inference).

Links between evolution and human cognition also have the potential to produce insights that run in the other direction — from our understanding of individual human minds to theories of how species and societies can adapt. One of the deepest questions in the study of cultural evolution is how societies can accumulate knowledge over time. This is fundamentally a question of how a society can exhibit a property that we see displayed by individuals — learning through experience. The link between evolutionary dynamics and Bayesian inference suggests a way to approach this question (for an example of recent work taking this approach, see [52]).

When an evolutionary process is considered as a candidate model for a cognitive process by way of the evolutionary process's connection to Bayesian computation, the evolutionary process is best thought of as an algorithm for approximate inference in the context of a particular computational problem faced by the mind. For example, in the case of memory, the computational goal is to continually infer the state of the world in the absence of new input, and the solution takes the form of a particle filter algorithm where particles are units of a memory-supporting commodity. In the case of creativity, the computational goal is to generate useful solutions to a problem or valuable ideas and the solution takes the form of a stochastic search algorithm over a semantic network. The evolutionary processes are algorithms.

When will an evolutionary process be a reasonable candidate model of a given cognitive process? Two criteria seem key. First, because the tightest links between evolutionary and cognitive processes have been found through equivalences between evolutionary processes and algorithms for inference, cognitive processes well described as inferential processes are more likely to find an evolutionary counterpart. Second, it helps if existing models or conceptualization of the cognitive process include at least some of the key components of an evolutionary process — a population of replicators, selection, random drift, mutation, etc.

In this Opinion, we suggest that the formal equivalence between certain simple evolutionary processes and algorithms for Bayesian inference is relevant to understanding human cognitive processes such as learning, memory and creativity. In doing so, we hope to encourage new work that keeps evolution in mind: finding equivalencies between particular evolutionary processes and algorithms for inference drawn from statistics and machine learning; creating process models of other cognitive capacities using the framework of evolutionary dynamics; and exploring the rich array of phenomena studied in modern evolutionary biology and what connection, if any, they may have to the similarly rich array of processes studied in the cognitive sciences.

Acknowledgments

This work was supported in part by the National Science Foundation (grants SPRF-IBSS-1408652 to T.L.G. and J.W.S; and BCS-1456709 to T.L.G.).

References

1. Nowak, M. A. (2006) *Evolutionary dynamics*. Harvard University Press
2. Dawkins, R. (1976) *The selfish gene*. Oxford University Press
3. Aunger, R. (2001) *Darwinizing culture: The status of memetics as a science*
4. Shennan, S. (2002) *Genes, memes, and human history: Darwinian archaeology and cultural evolution*
5. Campbell, D. T. (1974) Evolutionary epistemology. In P. A. Schlipp (Ed.), *The philosophy of Karl Popper* (pp. 413–463). La Salle, IL: Open Court
6. Cavalli-Sforza, L. L., & Feldman, M. W. (1981) *Cultural transmission and evolution: a quantitative approach* (No. 16). Princeton University Press
7. Boyd, R., & Richerson, P. J. (1988) *Culture and the evolutionary process*. University of Chicago Press
8. Henrich, J., & McElreath, R. (2003) The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews*, 12(3), 123-135
9. Axelrod, R. M. (1984) *The evolution of cooperation*. New York: Basic Books
10. Doz, Y. L. (1996) The evolution of cooperation in strategic alliances: Initial conditions or learning processes? *Strategic Management Journal*, 17(S1), 55-83
11. Nowak, M. A., & Krakauer, D. C. (1999) The evolution of language. *Proceedings of the National Academy of Sciences*, 96(14), 8028–8033
12. Christiansen, M. H., & Kirby, S. (Eds.). (2003) *Language evolution*. OUP Oxford
13. Ritt, N. (2004) *Selfish sounds and linguistic evolution: A Darwinian approach to language change*. Cambridge University Press

14. Blythe, R. A. (2015) Colloquium: Hierarchy of scales in language dynamics. *European Physics Journal B*, 88(295)
15. Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011) How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–285
16. Xu, F., & Tenenbaum, J. B. (2007) Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272
17. Price, G. R. (1970) Selection and covariance. *Nature*, 227, 520–521
18. Shalizi, C. (2009) Dynamics of Bayesian updating with dependent data and misspecified models. *Electronic Journal of Statistics*, 3, 1039–1074
19. Harper, M. (2009) The replicator dynamic as an inference dynamic. arXiv:0911.1763 [math.DS]
20. Baddeley, A. (1992) Working Memory and Conscious Awareness. *Theories of memory*, 11-20
21. Cowan, N. (2012) Working memory capacity. Psychology Press
22. Estes, W. K. (1950) Toward a statistical theory of learning. *Psychological Review*, 57(2), 94
23. Miller, G. A. (1956) The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63(2), 81
24. Zhang, W., & Luck, S. J. (2008) Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235
25. Bays, P. M., & Husain, M. (2008) Dynamic shifts of limited working memory resources in human vision. *Science*, 321(5890), 851–854

26. Suchow, J. W., Fougny, D., Brady, T. F., & Alvarez, G. A. (2014) Terms of the debate on the format and structure of visual memory. *Attention, Perception, & Psychophysics*, 76(7), 2071–2079
27. Vergauwe, E., Barrouillet, P., & Camos, V. (2009) Visual and spatial working memory are not that dissociated after all: a time-based resource-sharing account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(4), 1012
28. Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory *Psychological Review*, 114(1), 38–70
29. Muther, W. S. (1965) Erasure or partitioning in short-term memory. *Psychonomic Science*, 3, 429– 430
30. Suchow, J. W., & Griffiths, T. L. (2016) Deciding to Remember: Memory Maintenance as a Markov Decision Process. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*
31. Suchow, J. W., Fougny, D., & Alvarez, G. A. (2016) Looking Inward and Back: Real-Time Monitoring of Visual Working Memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(4), 660-668
32. Fougny, D., Suchow, J. W., & Alvarez, G. A. (2012) Variability in the quality of visual working memory. *Nature Communications*, 3, 1229 <http://doi.org/10.1038/ncomms2237>
33. Doucet, A. (1998) On sequential simulation-based methods for Bayesian filtering. Technical Report CUED/F-INFENG/TR.310
34. Ewens, W. J. (2004) *Mathematical Population Genetics I: Theoretical Introduction* (Vol. 27). Springer Science & Business Media

35. Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010) Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review*, 117(4), 1144–1167
36. Abbott, J. T., & Griffiths, T. L. (2011) Exploring the influence of particle filter parameters on order effects in causal learning. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*
37. Bays, P. M., Gorgoraptis, N., Wee, N., Marshall, L., & Husain, M. (2011) Temporal dynamics of encoding, storage, and reallocation of visual working memory. *Journal of Vision*, 11(10), 1–15
38. James, W. (1880). Great men and their environment *Atlantic Monthly*, 46, 441–449
39. Campbell, D. (1960) Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychological Review*, 67(6), 380–400.
40. Campbell, D. T. (1965) Variation and selective retention in socio-cultural evolution. In H. R. Barringer, G. I. Blanksten, & R. W. Mack (Eds.), *Social change in developing areas* (pp. 19–49). Cambridge, MA: Schenkman
41. Simonton, D. (2009) Creativity as a Darwinian phenomenon: The blind-variation and selective-retention model. In M. Krausz, D. Dutton, & K. Bardsley (Eds.), *The Idea of Creativity* (pp. 63–81). Leiden, Netherlands: Brill
42. Simonton, D. (2011) Creativity and discovery as blind variation: Campbell’s (1960) BVSr model after the half-century mark. *Review of General Psychology*, 15(2), 158–174
43. Kozbelt, A., Beghetto, R. A., & Runco, M. A. (2010) Theories of creativity. *The Cambridge Handbook of Creativity*, 20–47

44. Simonton, D. K. (1997) Creative productivity: A predictive and explanatory model of career landmarks and trajectories. *Psychological Review*, 104, 66–89
45. Simonton, D. K. (2000) Creative development as acquired expertise: Theoretical issues and an empirical test. *Developmental Review*, 20(2), 283–318
46. Simonton, D. K. (2004) *Creativity in science: Chance, logic, genius, and zeitgeist*. New York: Cambridge University Press
47. Brooks, S. (1998) Markov chain Monte Carlo method and its application. *Journal of the Royal Statistical Society: Series D (the Statistician)*, 47(1), 69–100
48. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953) Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6), 1087–1092
49. Hastings, W. K. (1970) Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1), 97–109
50. Bourgin, D. D., Abbott, J. T., Griffiths, T. L., Smith, K. A., & Vul, E. (2014) Empirical evidence for Markov Chain Monte Carlo in memory search. *In Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (pp. 224–229)
51. Mednick, S. A. (1962) The associative basis of the creative process. *Psychological Review*, 69(3), 220–232
52. Krafft, P., Zheng, J. Pan, W. Penna, N. D., Altshuler, Y., Shmueli, E., Tenenbaum, J. B. & Pentland, A. (2016) Human collective intelligence as distributed Bayesian inference. arXiv:1608.01987

53. Moran, P. A. P. (1958, January) Random processes in genetics. In *Mathematical Proceedings of the Cambridge Philosophical Society* (Vol. 54, No. 1, pp. 60-71). Cambridge University Press
54. Fudenberg, D., Nowak, M. A., Taylor, C., & Imhof, L. A. (2006) Evolutionary game dynamics in finite populations with strong selection and weak mutation. *Theoretical population biology*, 70(3), 352-363
55. Suchow, J. W. (2014) *Measuring, monitoring, and maintaining memories in a partially observable mind* (Doctoral dissertation)

Box 1. Evolution in discrete, finite populations

The pure death process, the Moran process, and the Wright-Fisher process are three models of evolution in discrete, finite populations:

Pure death. The pure death process is a simple model of population change that provides a starting point for thinking about evolutionary processes. Starting with a set of N individuals, at each time step one individual is selected at random and removed from the population.

Eventually, none remain.

Moran. The Moran process is a model of evolution that was originally used to describe the dynamics of allele frequencies [53] and which has more recently been leveraged to describe evolutionary processes in diverse settings, including frequency-dependent selection, emergence of cooperative behavior, and cultural evolution of language [54; 55]. The Moran process begins with a population of N individuals. At each time step, one individual is replaced by a copy of another individual, selected at random. The long-term dynamics of the Moran process dictate that, in the limit, the entire population reaches fixation, converging to one type.

Wright-Fisher. The Wright-Fisher model is similar to the Moran process, except that it has non-overlapping generations. Like the Moran process, it begins with a population of N individuals. However, in the Wright-Fisher model, at each time step the entire population is replaced by a new generation of N individuals. Critically, each member of the new generation inherits from a randomly selected member of the previous generation.

Box 2. The Metropolis-Hastings algorithm

The Metropolis Hastings (M-H) algorithm [48, 49] is a general-purpose algorithm for defining a Markov chain on a state space whose stationary distribution corresponds to an arbitrary distribution of interest. Given a (potentially infinite) state space $\{1, \dots, M\}$ and a distribution over proposed transitions conditioned on the currently occupied state, $Q(x', x)$, the M-H algorithm proposes a tentative state to transition to, x' from the proposal density $Q(x', x)$ defined by its currently occupied state, x . To determine whether to transition its current state to the proposed value, the M-H algorithm uses an acceptance procedure:

$$a = \frac{P^*(x')Q(x, x')}{P^*(x)Q(x', x)}$$

where $P^*(\cdot)$ is a distribution of interest, evaluated up to a multiplicative constant. If $a \geq 1$, the new state is accepted automatically. Otherwise, the new state is accepted with probability a . It can be shown that for any proposal distribution Q such that $Q(x'; x) > 0 \forall x, x'$, as $t \rightarrow \infty$, the probability distribution over states, $P(x)$, approaches $\frac{P^*(x)}{Z}$, where Z is the normalizing constant.

Glossary

Evolutionary dynamics: A mathematical framework for describing a reproducing population subject to selection, mutation, and drift.

Bayesian inference: A normative framework for updating beliefs in light of evidence through application of Bayes' rule.

Forgetting function. A function that tracks the amount of information remembered as it falls over time after initial exposure to an event or experience.

Algorithm. A procedure for carrying out a computation.

Wright-Fisher process: A neutral model from population genetics with a fixed population size. In each successive generation, individuals copy a randomly selected member of the previous generation.

Figure 1. Simple evolutionary models from population genetics. Instantiations of the pure death process (left), the Moran process (center), and the Wright–Fisher process. In all three, the population begins with a population of 9 individuals, evenly distributed across three types (red, purple, and orange). Figure adapted from [55].

Figure 2. Particle filters and the forgetting functions of visual memory. Each subpanel shows a forgetting function, which tracks the number of remembered objects as it falls over time, for a particular load of k objects. The data are fit with a particle filter model adapted to have a threshold number of particles, below which the memory is inaccessible (48 particles, 200 time steps, maintenance threshold of 5). Data originally from [55].

Figure 3. Creativity as stochastic search in a semantic network. A sample trajectory from the Metropolis-Hastings model of creative search on the RAT problem ‘Falling Actor Dust’. The model begins by selecting a single problem cue (Actor) and initializing its search at the corresponding node in a semantic network. Search proceeds iteratively by proposing a node to transition to from the set of neighbors for the current state, and either accepting or rejecting the transition in proportion to their relative utility (operationalized as the inverse of the average distance between a proposal and the three problem cues in the semantic network). If a proposal is rejected, the model generates a new proposal and retries until a transition is accepted. In the diagram above, rejected proposals are shown with dashed lines, while accepted transitions are shown as solid lines. The model moves through the network in this fashion until it arrives at the node corresponding to the correct answer to the question — in this case, ‘Star’.