

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

An evolutionary model of recombination in social learning

Permalink

<https://escholarship.org/uc/item/7b98h15h>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Whitaker, Eve

Quillien, Tadeo

Zhao, Bonan

Publication Date

2026

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

An Evolutionary Model of Recombination in Social Learning

Eve Whitaker (eve.whitaker@protonmail.com)¹, Tadeq Quillien² & Bonan Zhao¹

¹School of Informatics, University of Edinburgh

²Department of Psychology, University of Edinburgh

Abstract

Social learning is often valued for reducing the costs of individual exploration and avoiding mistakes. Its benefits, however, extend beyond simple error prevention in domains where knowledge is compositional: when ideas are generated by combining existing elements, meetings of minds are a fertile ground for novel innovations. We introduce an evolutionary agent-based model in which agents pursue diverse learning strategies. Some agents create knowledge independently, others exchange and recombine ideas. Agents interact repeatedly in a compositional task environment, building, sharing, and combining knowledge components. We find that social interaction produces a rich pool of partial ideas, and recombination among these ideas can connect disconnected knowledge strings, though only alongside accurate discoveries made by individual learners. Social recombination allows populations to explore combinatorial solution spaces more effectively than individual learning alone, or success-biased social learning. By highlighting the interplay between idea exchange, recombination, and selective copying, our results reveal a novel pathway through which social learning enhances adaptive knowledge accumulation.

Keywords: cumulative culture; social learning; individual learning; compositional knowledge; agent-based model

Introduction

A defining feature of cumulative culture is that knowledge does not merely accumulate. Instead, innovations interact and recombine to create new possibilities. New technologies are often developed from recombining familiar, existing technologies in novel ways (Arts & Veugelers, 2015), and terms for new concepts frequently make use of existing words (Xu et al., 2024). Through these compositional processes, relatively simple building blocks can give rise to rich, open-ended cultural repertoires (Brändle et al., 2023; Zhao et al., 2024). Understanding how such structured knowledge emerges, spreads, and stabilizes is therefore central to explaining the dynamics of human learning and cultural evolution (Bramley et al., 2023; Henrich, 2015; Mesoudi, 2016). Here, we present a model with recombination as the core method with which agents acquire full knowledge of an environment.

Compositionality and Learning Strategies

In complex environments that require compositional knowledge, people do not simply observe or copy. We also innovate, either alone or in collaboration with others, giving rise to a broad ecology of learning strategies. Through independent exploration we can construct deeply nested concepts that are

highly insightful, but may be difficult for others to understand (Chu & Evans, 2021; Foster et al., 2015). By contrast, we can also interact with others, and not only by copying, but also by actively connecting previously disconnected pieces of information, thereby contributing new knowledge to the collective culture (Derex & Boyd, 2016; Wu et al., 2019).

This diverse range of learning strategies, with some people preferring to work independently while others thrive in teams, has been documented across research fields. For example, psychological theories have sought to categorise this diversity along dimensions such as systematizing vs empathizing (Baron-Cohen, 2009; Greenberg et al., 2018) and introversion-extraversion within the Big Five trait taxonomy (Costa Jr & McCrae, 1992; John & Srivastava, 1999). Research on collective intelligence shows that groups benefit from cognitive diversity, particularly when differences are supported by effective social interaction, yielding performance gains that exceed predictions solely based on average individual ability (Duhigg, 2016; Woolley et al., 2010). Together, this evidence highlights the need to understand how diverse learning strategies—especially those that support recombination and transmission of partial knowledge—shape cumulative culture and collective knowledge. Our model directly addresses this question by formalising these strategies as distinct learning types, social and individual, and examining how recombination features in both.

Evolutionary Models

A natural approach for studying how diverse learning strategies interact over time is through evolutionary agent-based models (ABM). ABMs simulate populations of individuals who learn, interact, and reproduce under selective pressures (e.g., Acerbi et al., 2022; Boyd and Richerson, 1982; Henrich, 2004). These models have been instrumental in formalizing trade-offs between social and individual learning, demonstrating how population-level patterns can emerge from simple learning rules. However, many existing models operate with simple learning strategies in simple environments, with binary or low-dimensional knowledge states, and solutions are either acquired or not. Such formulations abstract away the internal structure of knowledge and therefore cannot capture settings in which learning involves constructing, combining, and transmitting partial representations.

This limitation extends to work that has incorporated environmental instability. Both simple (Rogers, 1988) and com-




Figure 1: Individual learning. **1)** The learner observes and saves a digit from the ground truth and combines this digit with an item from their knowledge inventory to create something new. **2)** If the learner is from the ‘checking’ sub-type, they only save this combination if it matches part of the ground truth.

plex, rugged environments (Barkoczi et al., 2016) have been studied, and some flexible forms of social learning have been shown to be adaptive to unstable environments (Morgan et al., 2022), none of these experiments incorporate recombination as a learning mechanism. Instability is particularly relevant for recombination, as environmental change rarely renders all current knowledge obsolete and recombination may offer a path to reconnecting partial knowledge.

Inspired by recent work that models innovation as a process of recombination, not just as the discovery of isolated facts (Zhao et al., 2024), we present an agent-based model set in a compositional task environment. In this environment, agents use diverse learning strategies: some focus on creating knowledge independently, testing it against the environment, while others work together to copy, combine, and propagate knowledge (Figures 1-2). This setting allows us to track how diverse learning strategies evolve and interact, shaping the cumulative culture in this rich, compositional world.

How do Learning Strategies Evolve?

Crucially, we contrast two types of social learning strategies. The first type of social learner copies other agents based on a ‘copy-the-successful’ heuristic (Jiménez & Mesoudi, 2019; Thompson et al., 2022), but otherwise does not attempt to create new ideas. A second type of social learner also copies successful agents, but additionally attempts to combine their own ideas with those of the agent they copy. Concretely, we examine whether recombination of ideas increases the prevalence and usefulness of social learning beyond what success-biased copying already contributes.

Model

We use an evolutionary agent-based model to study how different learning strategies shape the collective discovery of com-




Figure 2: Social learning: **1)** The learner observes a string from their partner’s inventory W_j and add the string to their own inventory W_i . **2)** If the learner is from the ‘combine’ sub-type, they additionally create a new combination based on compatible strings in their own inventory, W_i . Both agents save the new combination.

plex knowledge. Agents face a shared environment with an underlying ground truth that can only be recovered by accumulating and recombining partial knowledge over time. Within the population, agents differ in how they acquire information, relying either on direct interaction with the environment or on learning from others. These strategies vary in how readily new information is generated, recombined, and retained, creating trade-offs between exploration, verification, and transmission.

Task and Agents

Ground truth. The environment is represented by a ground truth string G of length ℓ , composed of unique digits from the set $\{0, 1, \dots, 9\}$, randomly sampled and without replacement. For example, one possible ground truth could be $G = 504678321$. An agent a_i holds an inventory W_i , which is a set of strings composed of digits.

Fitness. For a given ground truth G , an agent’s fitness reflects how much of G it has recovered, defined as the length of the longest common substring between G and any string in the agent’s inventory, normalized by the length of G :

$$F(a_i, G) = \frac{\max_{w \in W_i} \text{LCS}(w, G)}{|G|}. \quad (1)$$

Individual learners. In each learning episode, an individual learner can acquire a digit $d \in G$ uniformly from the ground truth G , and add d to its inventory if not already present:

$$W_i \leftarrow W_i \cup \{d\}, d \notin W_i.$$

Next, the learner can combine this acquired digit d with an existing knowledge element $w \in W_i$ such that $d \notin w$ and $w \neq d$. This combination forms a new string c by concatenating d and

w in random order. For example, ‘1’ and ‘34’ could create the combination ‘134’ or ‘341’.

We consider two subtypes of individual learners: *non-checking* and *checking*. Non-checking individual learners add any new combination c to their inventories:

$$W_i \leftarrow W_i \cup \{c\}.$$

And the checking individual learners (Figure 1) only store new combinations that are equal to or within the ground truth string:

$$W_i \leftarrow W_i \cup \{c\}, c \in G.$$

Social copiers. In contrast to learning from the ground truth, social learners learn from other agents, biased by the partner’s fitness score. They use a copy-the-successful heuristic with fitness score indicating whether the teacher-agent has useful knowledge in their inventory but not revealing what knowledge that is. Concretely, they choose another agent a_j from the population $A \setminus \{a_i\}$, with probability proportional to a_j ’s fitness:

$$P(\text{choose } a_j) = \frac{F(a_j, G)}{\sum_{k \in A \setminus \{a_i\}} F(a_k, G)}. \quad (2)$$

If all fitness values are zero, a partner is selected uniformly from $A \setminus \{a_i\}$. After choosing a partner, the social copier randomly retrieves a knowledge element $w_j \in W_j$ from the partner’s inventory and adds it to its own:

$$W_i \leftarrow W_i \cup \{w_j\}.$$

Social recombinators. The compositional nature of the task allows for a richer form of social learning, which we call *copy-and-combine* (see Figure 2). In this strategy, a social learner generates new knowledge by combining elements from their own inventory with elements obtained from a social partner. To capture costs of knowledge transmission, such combination is only possible when the two agents share partial overlap in their representations. Concretely, after selecting a partner and receiving a knowledge element w_j from them (same as social copiers), the focal social recombinators can additionally make a new combination. They do this by selecting one of their own knowledge strings, w_i , such that the first or last digit of w_i matches the last or first digit of w_j . One matching alignment is selected uniformly at random and the two strings are merged by overlapping the matching characters, producing a new string with no repeated digits:

$$c = \text{merge}(w_i, w_j).$$

For example, if $w_j = \text{‘674’}$ and $w_i = \text{‘36’}$, they can form a new element ‘3674’. This new knowledge element c is then added to both the focal agent’s and partner’s inventories:

$$W_i \leftarrow W_i \cup \{c\}, W_j \leftarrow W_j \cup \{c\}.$$

Table 1: Parameters of the model, including both variable and fixed parameters.

Parameter	Values
Individual learning type	non-checking, checking
Social learning type	copy-only, copy-combine
Ground shift generation	[100, 200, 300], None
Population size	200
Ground truth length	9
Mutation rate	0.01
Memory capacity	10
Maximum generations	400

Simulation process

Initialisation. Each simulation starts with initializing the environment and population. A ground truth string of length ℓ is randomly determined, forming the initial environment. Population size is determined by parameter N . Each agent is assigned a learning strategy, either social or individual, the exact subtypes of which are determined by two further parameters (see Table 1). Each agent is created with an empty inventory and a fitness score of zero.

Learning. After initialization, each agent carries out the learning task according to their learning strategy. All agents have a limited inventory. When it is full, adding a new item will cause a randomly-selected element to be deleted to free space. This very simple method of clearing the memory allows all agents to use the same memory function without undermining either learning strategy as it does not rely on access to the ground truth or to fitness.

After a round of learning, each agent’s fitness is recalculated based on their updated inventory (Equation 1).

Reproduction Once all agents have acted and fitness scores have been recalculated, a new generation is created. Agents in the new generation select a parent agent from the previous generation with probability proportional to fitness (Equation 2). After picking a parent, the new agent clones the inventory and learning strategy from the parent. There is a probability μ of mutation, where an agent will switch their learning strategy. The resulting population forms the new generation, and embarks on their learning tasks. This process repeats until the final generation.

Ground truth shift At specified generations, the ground truth G will change by one digit, resulting in a new ground truth which contains no repeated digits. For example, if $G = 123456789$ then it could shift to $G = 123456780$, where the digit that changes is randomly selected.

Parameters

Table 1 lists the parameters used in the simulation. For each simulation, we include one subtype of social learners and one




Figure 3: Graphs showing the proportion of social learners and individual learners across runs. The solid and dashed lines represent the mean, the shaded regions and error bars show the 95% confidence interval of the mean.

subtype of individual learners, leading to $2 \times 2 = 4$ simulation schemes: (a) social copy vs. individual non-checking, (b) social copy vs. individual checking, (c) social copy-combine vs. individual non-checking, (d) social copy-combine vs. individual checking. Finally, we also varied environmental stability: ground truth can either be static or shifts every 100 generations. Each simulation was run 30 times.

Results

We report the evolutionary dynamics of social and individual learning strategies, in both stable and unstable environments, and their effects on population-wide knowledge evolution, by analysing relative proportion of learning strategies, population-wide fitness and the number of innovations made over time.

Evolution of Learning Strategies

Populations exhibit distinct evolutionary dynamics depending on the present social and individual learning subtypes. In simulations with non-checking individual learners, the proportions of social and individual learners rapidly diverge (Figure 3a & 3c), with individual learning dropping to around 10% of the population. Regardless of social learning type, these populations converge on an incomplete understanding of the

environment, reaching a mean fitness of around 0.8 (Figure 4). While social learners adopt useful information quickly, increasing their fitness and therefore proportion in the population, non-checking individual learners continue to save inaccurate combinations, therefore stifling progress.

This divergence disappears when individual learners can verify their combinations against the ground truth. Instead, social and individual learners coexist. Figure 3b shows a strong convergence between the proportion of checking individual learners to copy-only social learners, with each strategy making up about 50% of the population. In simulations with checking individual learners and copy-and-combine social learners, the two strategy types still co-exist, with social learners slightly more prevalent, at around 70% of the population (Figure 3d).

These findings replicate classic results showing that social and individual learning strategies can co-exist in a mixed-strategy equilibrium (Acerbi et al., 2022; Rogers, 1988). They also show that the ability to copy successful individuals is not the only evolutionary advantage of social learning. Social learners that can exchange information with each other, recombining them to generate new ideas, come to dominate the population to a larger extent than social learners with only a ‘copy-the-successful’ bias.

Evolution of Collective Knowledge

Fitness reflects how much of the ground truth an agent has stored in their inventory. Therefore, we can investigate population-wide knowledge evolution via the mean group fitness. As shown in Figure 4, the main difference in group fitness is between populations with non-checking individual learners, which stall at about 0.8, and those with checking individual learners, which reach close to full fitness (see Figure 4). We conjecture that limited memory, together with early spread of partially correct information, prevent populations with non-checking individuals from reaching full fitness.

Overall, group fitness increases faster in populations with social recombinators (copy-combine social learners). As illustrated by Figure 5, these learners rapidly increase the number of innovations in the population, as they can combine disconnected pieces of knowledge earlier than individual learners can (individual learners only ever combine one digit with one inventory item). Since during social recombination both partners save the new combination, recombination accelerates the spread of new knowledge, as creating a much longer combination improves an agent's fitness dramatically. As a result, in populations with copy-combine social learners, peak fitness is reached sooner because the currently available knowledge fragments are joined together earlier.

In populations with checking individual learners, there are a greater number of innovations than in populations with the non-checking type (Figure 5). This is due to the verification mechanism giving the population direct access to the ground truth, therefore facilitating innovation. Populations with both checking individual learners and copy-combine social learners created the most innovations (Figure 5), and reached full recovery of ground truth much faster than other populations: this population discovers all possible substrings of the ground truth by generation 200, whereas other populations continue to find new combinations over a longer period.




Figure 4: The mean fitness of the population for each combination of learning types, across generations.




Figure 5: The mean cumulative innovations for each combination of learning types. Innovations here refers to new combinations that are a subset of, or equal to, the ground truth. This graph does not include the discovery of individual digits. Where the line stops indicates the point at which no more innovations were made.




Figure 6: The mean fitness of the population in an environment where the ground truth changes every 100 generations.

Evolution in Changing Environments

Figure 6 shows the mean fitness of the populations in simulations where the ground truth shifts every 100 generations. Populations with copy-combine social learners are able to recover from a ground truth shift much more effectively, but only if the individual learners in that population are ones that verify their combinations. In contrast, populations with copy-only social learners are not able to recover as fast or to the fitness levels they had before the shift. In both populations where social learners only copy, fitness increases steadily but not fast enough to regain the pre-shift fitness levels. In both cases with checking individual learners, fitness can increase much faster. Together, these results suggest that recombination during social learning facilitates faster innovation, which protects against unstable environments, though only alongside individual learners who check.

Discussion

Social learning is a powerful way of acquiring knowledge about the world. One well-known advantage of social learning is that social learners avoid the costs and potential dangers associated with extracting knowledge directly from the world (Henrich, 2015). Here, we highlight a different evolutionary advantage of learning from others: when knowledge is compositional and new ideas are generated by combining existing ideas (Muthukrishna & Henrich, 2016; Zhao et al., 2024), the large space of possible innovations creates a unique advantage for social learners, who can exchange ideas with each other to fuel the creative process.

We investigated this hypothesis by conducting evolutionary agent-based simulations in a complex environment where knowledge has compositional structure. We find that social learners who exchange, combine, and recombine ideas with others enjoy more success (in terms of relative population frequency) than social learners who merely copy successful others. Additionally, populations with these ‘social recombinators’ see their fitness increase quicker, and are most robust to environmental shocks.

Social recombinators have no immediate way of telling which novel combination of ideas is the right one, but sometimes will hit upon a very good new combination by chance. Because the fitness gains of such events are high, the social innovation strategy is a successful one. The social recombination strategy also allows populations to increase their overall knowledge faster.

At the same time, we find that individual learners also have a role to play in helping knowledge grow. Specifically, populations achieve more complete knowledge in the long run if individual learners can directly test their new hypotheses against reality. We can understand this phenomenon as follows. Social recombinators can create novel ideas rapidly, but the usefulness of these ideas is only tested *indirectly*. Good ideas spread via social learning because agents with good ideas enjoy higher fitness and are copied more often. This means that ideas that are only *approximately* correct (e.g. digit strings that differ from the ground truth by one or two digits) can easily spread, and it is difficult for social learning alone to guide the population to fully accurate knowledge. In contrast, individual learners who can check their hypotheses inject small, but perfectly correct ideas into the population knowledge pool, helping the population converge to the ground truth in the long run. In sum, our work suggests that the growth of knowledge in a population depends both on empirical access to reality and the creative exchange of ideas.

Long-run population fitness in our simulations is determined by the particular type of individual learning strategy and does not depend on the social learning strategy. This finding is reminiscent of earlier models in which social learning does not increase the equilibrium fitness of a population (Rogers, 1988). However, social recombinators can impact a population’s efficiency in reaching the fitness ceiling, which can help a population adapt in an unstable environment.

Our finding that individual and social-learning strategies can coexist within a population suggest that evolutionary dynamics favour diversity in learning strategies. At the same time, we have assumed that individuals come in discrete types. We could also have included agents who can flexibly switch from one strategy to the next. Future research should investigate whether evolution favours flexible agents in our setting, or whether it can actually foster cognitive variability, where agents specialize for one particular learning strategy (cf. Wolf & McNamara, 2012).

Our populations also lack network structure, which has been shown to impact the complexity of innovations developed over generations (Derex & Boyd, 2016; Derex et al., 2018). Introducing variations in connectivity might change the dynamics seen in our results, particularly in how well social recombinators can connect disconnected knowledge. Finally, although the ground truth in our simulations has some complex structure, it is still a relatively short and simple string of digits. Simulating a more complex ground truth, coupled with variations in network structure, could be used to investigate situations where the population does not converge to a unique picture of reality, but where stable “schools of thought” co-exist.

References

- Acerbi, A., Mesoudi, A., & Smolla, M. (2022). *Individual-based models of cultural evolution: A step-by-step guide using R*. Routledge.
- Arts, S., & Veugelers, R. (2015). Technology familiarity, recombinant novelty, and breakthrough invention. *Industrial and Corporate Change*, 24(6), 1215–1246.
- Barkoczi, D., Analytis, P. P., & Wu, C. M. (2016). Collective search on rugged landscapes: A cross-environmental analysis. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 38.
- Baron-Cohen, S. (2009). Autism: The Empathizing–Systemizing (E–S) Theory. *Annals of the New York Academy of Sciences*, 1156(1), 68–80.
- Boyd, R., & Richerson, P. J. (1982). Cultural transmission and the evolution of cooperative behavior. *Human Ecology*, 10(3), 325–351.
- Bramley, N. R., Zhao, B., Quillien, T., & Lucas, C. G. (2023). Local Search and the Evolution of World Models. *Topics in Cognitive Science*, tops.12703.
- Brändle, F., Stocks, L. J., Tenenbaum, J. B., Gershman, S. J., & Schulz, E. (2023). Empowerment contributes to exploration behaviour in a creative video game. *Nature Human Behaviour*, 7(9), 1481–1489.
- Chu, J. S., & Evans, J. A. (2021). Slowed canonical progress in large fields of science. *Proceedings of the National Academy of Sciences*, 118(41), e2021636118.
- Costa Jr, P. T., & McCrae, R. R. (1992). Four ways five factors are basic. *Personality and individual differences*, 13(6), 653–665.

- Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, 113(11), 2982–2987.
- Derex, M., Perreault, C., & Boyd, R. (2018). Divide and conquer: Intermediate levels of population fragmentation maximize cultural accumulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1743), 20170062.
- Duhigg, C. (2016). What google learned from its quest to build the perfect team. *The New York Times Magazine*.
- Foster, J. G., Rzhetsky, A., & Evans, J. A. (2015). Tradition and Innovation in Scientists' Research Strategies. *American Sociological Review*, 80(5), 875–908.
- Greenberg, D. M., Warrier, V., Allison, C., & Baron-Cohen, S. (2018). Testing the Empathizing–Systemizing theory of sex differences and the Extreme Male Brain theory of autism in half a million people. *Proceedings of the National Academy of Sciences*, 115(48), 12152–12157.
- Henrich, J. (2004). Demography and Cultural Evolution: How Adaptive Cultural Processes Can Produce Maladaptive Losses—The Tasmanian Case. *American Antiquity*, 69(2), 197–214.
- Henrich, J. (2015). The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter. In *The secret of our success*. Princeton University press.
- Jiménez, Á. V., & Mesoudi, A. (2019). Prestige-biased social learning: Current evidence and outstanding questions. *Palgrave Communications*, 5(1).
- John, O. P., & Srivastava, S. (1999). The big five trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (2nd, pp. 102–138). Guilford Press.
- Mesoudi, A. (2016). Cultural Evolution: A Review of Theory, Findings and Controversies. *Evolutionary Biology*, 43(4), 481–497.
- Morgan, T. J. H., Suchow, J. W., & Griffiths, T. L. (2022). The experimental evolution of human culture: Flexibility, fidelity and environmental instability. *Proceedings of the Royal Society B: Biological Sciences*, 289(1986), 20221614.
- Muthukrishna, M., & Henrich, J. (2016). Innovation in the collective brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1690), 20150192.
- Rogers, A. R. (1988). Does Biology Constrain Culture? *American Anthropologist*, 90(4), 819–831.
- Thompson, B., Van Opheusden, B., Sumers, T., & Griffiths, T. (2022). Complex cognitive algorithms preserved by selective social learning in experimental populations. *Science*, 376(6588), 95–98.
- Wolf, M., & McNamara, J. M. (2012). On the evolution of personalities via frequency-dependent selection. *The American Naturalist*, 179(6), 679–692.
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688.
- Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, 566(7744), 378–382.
- Xu, A., Kemp, C., Frermann, L., & Xu, Y. (2024). Word reuse and combination support efficient communication of emerging concepts. *Proceedings of the National Academy of Sciences*, 121(46), e2406971121.
- Zhao, B., Vélez, N., & Griffiths, T. (2024). A rational model of innovation by recombination. *Proceedings of the annual meeting of the cognitive science society*, 46, 290–296.