Hybrid social learning in human-algorithm cultural transmission

Levin Brinkmann^{*1}, Deniz Gezerli¹, Kira von Kleist¹, Thomas F. Müller¹, Iyad Rahwan¹,
and Niccolò Pescetelli¹

¹*Center for Humans and Machines, Max Planck Institute for Human Development*

Abstract

Humans are impressive social learners. Researchers of cultural evolution have studied the many biases shaping cultural transmission by selecting who we copy from and what we copy. One hypothesis is that with the advent of superhuman algorithms a hybrid type of cultural transmission, namely from algorithms to humans, may have long-lasting effects on human culture. We suggest that algorithms might show (either by learning or by design) different behaviors, biases and problem-solving abilities than their human counterparts. In turn, algorithmic-human hybrid problem solving could foster better decisions in environments where diversity in problem-solving strategies is beneficial. This study asks whether algorithms with complementary biases to humans can boost performance in a carefully controlled planning task, and whether humans further transmit algorithmic behaviors to other humans. We conducted a large behavioral study and an agent-based simulation to test the performance of transmission chains with human and algorithmic players. We show that the algorithm boosts the performance of immediately following participants but this gain is lost for participants further down the chain. Our findings suggest that algorithms can improve performance, but human bias may hinder algorithmic solutions from being preserved.

1 Introduction

When the first superhuman computer program in the game of Go —AlphaGo— beat the world champion Lee Sedol in 2016, its gameplay was considered surprising and unconventional, apparently violating longstanding Go traditions. In particular, for move 37, AlphaGo calculated the chance of a human player making the same move as 1 in 10000 [1].

Its unconventional play likely originated from the fact that AlphaGo, and more so its successor AlphaGo Zero [2, 3], learned through self-play with little or no reliance on human historic gameplay. The performance of AlphaGo raises the question, how such novel gameplay would influence human strategies [1]. Replaying historic human matches of the last 300 years showed that an algorithm similar to AlphaGo Zero increasingly often chooses the same move as humans [4], indicating convergence towards a common gameplay. Remarkably, there has been a steep increase in this alignment since 2017 when such an algorithm became available to the public [5, 4]. These observations suggest the fascinating hypothesis that increased alignment is the result of a hybrid form of social learning, where AI solutions are copied and maintained by humans. Similar patterns of increased alignment between human and algorithmic play have been suggested in the game of chess [6].

The use of technology, such as books or software, for human training in games like Go and chess is not a novel phenomenon and represents a common way how knowledge is socially transmitted from a generation to the next. Yet, current development in AI made it possible for algorithms to not only play chess, but to play creatively without the need to rely on human games. This opened up the possibility of social learning— namely learning by observation [7]— between artificial and biological agents. Digital technology already influences the processes of social transmission among people by providing new and faster means of communication and imitation [8, 9]. Going one step further, we argue that rather than a mere means for cultural transmission (such as books or the Internet), algorithmic agents and AI may also play an active role in shaping cultural evolution processes online where humans and algorithms routinely interact.

In this paper, we propose and test the hypothesis that social learning between humans and algorithms may be especially beneficial when biological and artificial problem solvers show diversity in the heuristics

^{*}corresponding author: brinkmann@mpib-berlin.mpg.de

and strategies they adopt to problem-solving. Diversity in information, biases and problem-solving strategies has been suggested to reduce herding and error cascades [10–13]. By self-learning or by design, algorithms showing complementary biases to humans could foster the discovery of new solutions in domain-specific problems and improve outcomes compared to human-only problem solvers. We would expect this effect to be greater in domains where human bias is suboptimal for the problem to be solved. Algorithms learning from interaction with their environment, rather than from human data, may be able to innovate over human solutions, as in our opening anecdote. Similarly, when human biases are known, algorithms can intentionally be designed to exhibit complementary biases to their human counterpart to enhance collective performance [14]. Although several heuristics that humans employ are adaptive under assumptions of cognitive constraints and bounded rationality [15–17], they can be suboptimal under restricted circumstances—*e.g.*, in digital environments for which they are not adapted to [18–20]. In this study we focus on a specific human bias, namely the tendency for myopic behavior when facing a sequential decision [21–23].

Many decision-making tasks (including Go and chess) are composed of sequential decisions that require an agent to explore large decision trees. As the tree grows exponentially large with increasing number of decisions, humans and algorithms rely on various heuristics to avoid exploring the full tree [23, 24, 3]. Huys et al. [23] introduced a goal-directed decision-making task where participants plan and make sequential moves on a directed network (Figure 1). Each move is associated with gains or losses of different magnitude. Participants have full knowledge of the network and the rewards associated with each transition between two nodes. The authors found that people tend to selectively discount decision branches following a large costs [23, 25]. We call this selective discounting aversive pruning bias. Such a heuristic can lead to sub-optimal solutions when an initial cost has to be born before a larger reward can be gained [26]. Lieder et al. showed that an algorithm can augment human decisions by providing pseudo-rewards (*e.g.*, gamified badges and direct rewards) that reshapes the value of each option based on an optimal value function so to render people’s myopic strategy optimal [26]. They showed that such pseudo-rewards can significantly increase participants performance. Similar asocial algorithmic aids to human cognition have been investigated [27], but it remains unclear whether the associated benefits disappear when the aid is removed. We address the question of whether algorithms can durably improve human performance via social learning and whether humans further transmit such new behavior to new human imitators. Social learning can be especially useful in complex problems and under uncertainty [28, 29, 7, 30], and it does not require explicit causal understanding to be effective [31].

This paper explores hybrid social learning in the lab, adapting Huys’s decision-making task to a transmission chain paradigm with human and algorithmic players [32–34]. In a transmission chain, players solve a task in a sequence and can observe (and copy) the solution of the previous player before they enter theirs [35]. We compare a control condition of human-only players with a hybrid treatment condition where an algorithm replaced a human player in the second generation. Transmission chains have been used to investigate how biases in social learning shape cultural evolution [35–37]. Previous research has shown that humans have different biases of what (content bias) and who (context bias) is copied [38, 7]. Both content and context biases are likely to be important in hybrid social learning. For instance, people differ in how they develop and sustain trust in human and algorithmic partners [39, 40]. As we were interested in what people learn from artificial players, we controlled context bias by not revealing whether the previous player in the transmission chain was a human or an algorithm.

Participants repeatedly playing on the same network tend to reuse similar actions [25]. To exclude such asocial learning, we developed a novel randomized version of the task, in which each participant plays the same network only once. We classified environments where the human aversive pruning bias is adaptive (human-rewarding environments) or misadaptive (human-regretful environments). An agent with aversive pruning bias would perform well in the former and poorly in the latter. We designed the algorithmic player to show a bias opposite to humans, namely a tendency to explore decision branches associated with initial costs. We predicted performance improvement over generations and better performance in hybrid chains than in human-only chains (control) due to the increased strategic diversity of the former. In line with our preregistered hypotheses, we found increased performance over generations and a short term performance improvement in the generation after the algorithmic solutions was introduced. However, in contrast to our initial hypothesis, the improvement introduced by the algorithm was not sustained over following generations of players. Solutions that conflict with the human aversive pruning bias had lower copying fidelity and therefore quickly disappeared. We develop an agent-based model that replicates some of our findings and makes novel predictions about untested experimental conditions. We discuss our results in terms of content bias and frequency of encounters with algorithmic solutions.

2 Methods

2.1 Participants

All 177 Participants were recruited through Prolific (www.prolific.co), where they were redirected to an external website to complete the experiment. Before starting, they completed a consent form and read the instructions. The experiment, including two practice rounds, took around 60 minutes in total. Participants were paid £7 for the completion of the experiment. Furthermore, there was a reward of one penny given for every 100 points gained during the experiment. Participants received on average £3.20 bonus payments, depending on their performance. In cases where participants had to drop out because of technical issues (failed network connection, etc.), they were paid a compensation of £3.50. The experiment was run in multiple sessions. Two sessions failed for technical reasons. Data from failed sessions was disregarded entirely and the experiment restarted from the last saved image. The only entry requirement was speaking English. All participants were included in the analyses.

Figure 1: (a) In the first stage of the task, participants saw an animation of the solution entered by the previous player (left hand-side). A snapshot showing the transition from node E to node C is depicted. In the second stage, the participants could enter a path by clicking on the respective nodes in sequence (center). The node with gray background color indicates the current node the participant is in. In the last stage, the total score of the player's sequence is revealed (right hand-side). The network presented here is classified as human-regretful. (b) For each environment class, we constructed two chains of 8 generations of players. In hybrid chains, the second generation player was replaced by an algorithm. The networks depict the solutions of the first 4 generations as well as the last generation for a selected environment (corresponding to (a)). The integer on the arrows denotes the step at which a player was choosing the move. The cumulative reward is shown in the upper right corner of each graphics. In this example, for the human-only chain the cumulative reward increases at first, but quickly reaches a plateau. For the hybrid chain, the algorithm shows a performance greater than observed in the human-only chain, but this improvement gets lost over subsequent human generations.

2.2 Task

The task was an adaptation of Huys et al. [23], in which participants were asked to find an optimal sequence of moves on a carefully designed directed network of 6 nodes. We generalized the task, by randomly sampling networks, instead of using a single network. From each node, there were exactly two possible moves to other nodes, each being associated with one of four possible payoffs (-100, -20, 20, or 140) (see Figure 1a). Possible moves were visualized by directed arrows with colors coding for their respective payoffs (red, orange, blue, and green for increasingly larger rewards). The objective was to find a path of 8 moves beginning at a fixed starting node, which maximizes cumulative payoffs. We called a network together with a specific starting position an environment. The experiment was implemented using a customized version of the Empirica framework [41] and consisted of three consecutive stages (see Figure 1a). In the first stage, participants were asked to watch the solution of a previous player. They saw the score of the previous player and a 15 seconds animation of the 8 moves. Each move between two nodes was animated individually for about two seconds with both nodes being highlighted by a darker color and the corresponding directed arrow and reward thickened (Supplementary Figure 1 and Supplementary Video 1). In the second stage, participants were shown the same environment and were asked to select a path of 8 moves. The path could be entered by clicking on the nodes in sequence. The currently occupied node was displayed in darker color. If a node was selected which could not directly be reached from the current node, the erroneously selected node was colored in bright red. The participant was then able to select a different node instead. Participants were able to see their current accumulated score, the number of steps remaining and a score of the last moves entered. This information was immediately updated whenever a participant clicked a possible target node (Supplementary Figure 2). The answer of the participant was considered to be valid if all 8 moves were played in the allotted time (15s). To strongly incentivise participants to respond even if being uncertain about the solution, participants paid a large cost (-500) for the round if they did not provide a valid answer on time. In the third and final stage, the final score of the current round was displayed for 5 seconds in large fonts (Supplementary Figure 3). Participants were also informed if they had failed to enter a response on time.

2.2.1 Experimental Design

Transmission chains featured 8 different players, who could be human or algorithmic. We call the position in the chain a generation. Within each chain, each player was exposed to the solution of the previous player. Players in the first generation were exposed to a random solution.

We manipulated the chain type (human vs hybrid, Figure 1b) and the environment type (human-rewarding vs human-regretful), in a 2x2 design. In human-only chains (control condition), all 8 generations featured human players. In hybrid chains (treatment condition), an algorithm (described below) replaced a participant in the 2nd generation and provided an algorithmic solution instead. The rest of the chain comprised human players. 800 environments of two different types were investigated. The two types, 'human-rewarding' and 'human-regretful', differed in whether aversive pruning respectively increases or reduces the expected reward (see below for further detail). For each of the 800 environments two chains were constructed, one for each of the two chain conditions. This led to a total of 1600 chains and 12800 games, of which 800 were played by the algorithm.

Participants were assigned to new environments on the fly at random based on availability, with the constraints that (a) the previous generation in the chain had successfully completed all stages and (b) that participants played each environment at most once. If a participant did not enter a path of 8 moves on time, the solution was considered invalid and the corresponding position in the chain was reopened for a new participant. Each participant played a maximum of 80 rounds. Towards the end of the experiment, participants completed less than 80 rounds as no further environments were available. Due to the random assignment procedure, participants were equally likely to play in each of the chain types as well as the environment types throughout the experiment. However, participants entering the experiment at the beginning were more likely to be placed in earlier generations, compared to participants who entered the game at a later stage. Hence, we added to our regression models a random effect for each individual participant and, to control for individual experience with the task, we added a fixed effect for the number of rounds already played in the experiment.

2.2.2 Aversive Pruning Model

Huys et al. described a pruned tree search algorithm for this type of task that best fitted human decisions [23]. The model calculated the state-action value $Q(a, s)$ of each action (move) a in state s . The value of a

particular action is given by the sum of the immediate reward $\mathcal{R}(a, s)$ and the maximum value of the next action a' from the next state $s' = \mathcal{T}(a, s)$ where $\mathcal{T}$ is the deterministic transition function. At each level of depth of the search tree, future rewards are discounted by a factor of $(1 - \gamma_{a,s})$. Together, this leads to the Bellman equation

$$\mathcal{Q}(a, s) = \mathcal{R}(a, s) + (1 - \gamma_{a,s}) \max_{a'} \mathcal{Q}(a', \mathcal{T}(a, s)). \quad (1)$$

The parameter $\gamma_{a,s}$ is interpreted as the rate of pruning of the search tree in a mean field approximation [23]. Correspondingly, rewards k steps ahead are discounted by a factor of $(1 - \gamma_{a,s})^{(1-k)}$. Scaling the state-action value $\mathcal{Q}$ by the inverse temperature β and applying a softmax function leads to the policy

$$\pi(a_t | s_t) = \frac{e^{\beta \mathcal{Q}(a_t, s_t)}}{\sum_{a'} e^{\beta \mathcal{Q}(a', s_t)}}. \quad (2)$$

Central to their work, Huys et al. defined a selective 'Pruning' version of this model to account for stronger pruning when participants encounter a large cost [23]. In our experiment, a large cost is defined as a reward of -100. In this model, which we will call the aversive pruning model, the $\gamma_{a,s}$ parameter takes two different values, a specific pruning rate γ_s in the case of large costs and general pruning rate γ_g in all other cases [3].

$$\gamma_{a,s} = \begin{cases} \gamma_s, & \text{if } \mathcal{R}(a, s) = -100 \\ \gamma_g, & \text{else} \end{cases} \quad (3)$$

2.3 Environment Generation, Selection and Classification

Before the experiment, we created 800 environments, each one characterised by a directed network of six nodes and a starting node, with each edge of the network defining a possible move between two nodes. First, we created a pool of 60000 strongly connected directed networks and uniformly sampled, for each link between two nodes, one of four possible reward (-100, -20, 20, 140). Considering six possible starting nodes for each network, this yielded 360000 environments. We then calculated for each environment a path maximising the cumulative reward across 8 consecutive moves. To reduce variation in the reward distributions, environments with a maximum reward in the upper and lower quartile were removed from the pool. To avoid trivial solutions (*e.g.*, loops between two nodes), environments were rejected if the maximum path did not cover at least four distinct nodes. Finally, to exclude environments with myopic optimal solutions, we compared for each node on the optimal path, the reward of the optimal move with the reward of the alternative sub-optimal move. We required environments to have at least four moves in which the optimal move has the same or a lower direct reward than the sub-optimal one.

The final selection of environments was based on the sensitivity of the expected total reward to changes in the aversive pruning parameter. The aversive pruning sensitivity for each environments was examined by choosing a reference policy ([2], $\gamma_g = \gamma_s = 0.35$ and $\beta = 0.03$) and calculating the derivative of the expected reward with respect to the aversive pruning parameter γ_s . We then randomly selected 400 environments in the lowest and highest deciles of aversive pruning sensitivity. We defined environments in the lowest decile human-regretful as showing an aversive pruning bias in these environments leads to lower rewards. Environments in the highest decile were called human-rewarding.

2.4 Matching the algorithmic performance

Rather than using an algorithm with super-human performance, we were interested in the effect of hybrid social learning between players that similarly discount future rewards and yet show different biases. We thus tuned the algorithm to have a comparable performance with human player. On a pilot study we estimated with a Bayesian model fit (see Supplementary Material) the model parameter of human player as $\gamma_g = 0.20$ (CI_{90} : (0.15, 0.25)), $\gamma_s = 0.45$ (CI_{90} : (0.32, 0.58)) and $\beta = 0.012$ (CI_{90} : (0.011, 0.014)). Note that γ_g and γ_s are of comparable to the values reported by Huys et al. [23]. However, in our pilot data we observed a lower inverse temperature β . One possible difference is that in Huys and colleagues' work participants were extensively trained on one specific network, while in our work participants played different environments (and each environment only once). Lower inverse temperature β in our study might then indicate more randomness in our participants responses.

We designed a risk seeking algorithm with a bias inverse to humans but with comparable performance. We fixed $\gamma_g = 0.5$ and $\gamma_s = 0.05$ so to have a comparable tree depth as human players, and then fitted $\beta = 0.0264$ to match the performance of human player on the pilot study. To mimic social learning, the

algorithm used an additional heuristic at run time. First, a solution, i.e. a sequence of 8 actions, was sampled using the parameters described above, then the total reward of this solution was compared with the one of the previous player. If the reward of the algorithm’s solution was greater or equal to the previous player’s reward, the algorithmic solution was played. Otherwise, an exact copy of the previous player’s solution was played by the algorithm.

2.5 Statistical Analysis

We ran two separate hypothesis-driven regressions models, one on solution total reward (*i.e.*, the sum of the rewards over the 8 moves of a single round) and the other on whether a solution was optimal or not. Additionally, we ran exploratory regressions on the number of actions copied between solutions, modeled as Poisson distribution with a logarithmic link function. Different models were compared with a likelihood ratio test (*anova* function in R). We used a single model for both types of environments and consequently added interactions between each fixed effect of interest and the environment type. We used the human regretful environments as baseline.

We run a maximum likelihood estimation of the pruning parameters ([3]) and inverse temperature ([2]) for each condition and environment type on 100 bootstrapped resamples of the solutions. We report percentile confidence intervals and p-values based on the rate of samples satisfying the null hypotheses. 95% confidence intervals are reported throughout. The code of the statistical analysis and the corresponding data is published with this work.

2.6 Preregistration

Our preregistered hypotheses were that (H1) in human-only chains, individual solutions will improve across generations, within each environment, via social learning; (H2) We expected that placing the algorithm in the chain at generation two (*GenHy2*) will locally increase performance so that a score boost is observed in generation three (*GenHy3*) and following generations compared to the first generation (*GenHy1*); We expected our algorithmic manipulation (H3) to globally increase performance as measured by normalized score accrued in the game, (H4) to accelerate solution discovery as measured by the slope of score improvement and reduction of error compared to the global optimal solution, and (H5) to increase the likelihood of chains discovering the best solution. Furthermore, we expected (H6) to not have an effect of the algorithm intervention on human rewarding networks as people will be able to judge that their own solution is better than the algorithm’s.

3 Results

3.1 Algorithm impacts following generations, but effect quickly decays

To investigate and compare the evolution of the performance of solutions in the different chains, we run a linear mixed-effects model predicting the reward of a individual solution, by considering (a) the numeric position in the chain (generation), (b) individual generations following the algorithm and (c) the number of rounds participants had previously played (max 80) as additive effects. For the first two effects (a, b) we added an interaction with the environment type. Additionally we added random effects for the (d) individual participants and (e) individual environments. The round of a participant (c) was added to account for non-social learning of participants. We considered the first generation of the human regretful environments as the baseline. Algorithmic solutions were not considered in this analysis because they were part of our treatment.

We encoded the influence of the algorithm on the performance of following generations by adding two independent effects for the two generations directly following the algorithm (*GenHy3* and *GenHy4*). All further generations were assigned a single effect (*GenHy5+*) and we considered solutions not following an algorithm in the chain as the baseline. This includes all solutions in human-only chains as well as the first generation (*GenHy1*) in the hybrid chains, where the algorithm has not yet been introduced. We selected this most parsimonious model (Supplementary Table 1), because others that either included independent effects for all generations following the algorithm ($\chi^2 = 1.99$, $df = 2$, $p = 0.37$) or that included independent effects on the three generations (*GenHy3*, *GenHy4* and *GenHy5*) following the algorithm ($\chi^2 = 5.13$, $df = 6$, $p = 0.52$) did not significantly improve model fit.

As a first validation of our experimental setup, we quantified the effect of social learning by investigating the impact of generation on reward. We found for human regretful environments an improvement of 3.867 (*s.e.* = 1.244, $t = 3.109$, $p = 0.002$, $CI = (1.429, 6.305)$) points from generation to generation, and for

Figure 2: (a) Difference in performance between conditions (hybrid - human-only); (inset) Performance improvement over generations within human-only chains in relation to the first generation; (b) Average number of actions of solutions that match those of generation 2 within the same chain; (c) Maximum likelihood estimates of the pruning parameter for human-only chains; (d) difference of the maximum likelihood estimates of the pruning parameter between conditions (hybrid - human-only). All panels share the same color code. Vertical bars are indicating bootstrapped 95 % confidence intervals. A dashed vertical line shows the algorithm’s position.

the human rewarding environment an additional improvement of 4.859 ($s.e. = 1.390$, $t = 3.495$, $p < 0.001$, $CI = (2.134, 7.583)$) points per generation. The inset in figure 2a depicts the average reward of solutions in human-only chains in relation to the reward of the first player in the chain. The positive slope indicates increase in performance over the 8 generation suggesting the presence of social learning as predicted (H1) and the accumulation of higher performing solutions in later generations. Social learning appears to lead to larger increases in performance for ‘human rewarding’ environments where the human bias is beneficial.

Having found that social learning does occur, we investigated the impact of the algorithm on following generations in mixed chains. Figure 2a depicts the average within-environment reward difference between hybrid chains with human-only chains. We found for human regretful environments a significant effect ($\beta (s.e.) = 30.786 (7.974)$, $t = 3.861$, $p < 0.001$, $CI = (15.157, 46.415)$) for participants directly following the algorithm and a weak effect ($\beta (s.e.) = 13.225 (7.922)$, $t = 1.669$, $p = 0.095$, $CI = (-2.302, 28.753)$) for the second generation following the algorithm. No effect was found the remaining generations ($\beta (s.e.) = -2.473 (5.094)$, $t = -0.485$, $p = 0.627$, $CI = (-12.457, 7.511)$). We did not find evidence for interactions of these effects with the environment type (see Supplementary Table 2). Participants in the generation following the algorithm (generation 3) gained higher rewards than their counterparts in human-only chains. However, this effect appears to quickly wear off, suggesting a temporary boost in human performance due to hybrid social learning (H2). However, we found no evidence for global performance improvements in hybrid chains (H3) when considering the second half of the transmission chain (generation 5-8).

We then investigated how participants’ behavior changed over the generations. We calculated a maximum likelihood estimate of the aversive pruning model parameters, independently for each generation and each condition. To calculate confidence intervals around each point estimate, we bootstrapped 100 resamples. Figure 2c shows the pruning parameters estimates for human-only chains. For these chains, we ran a linear regression on the estimates and with generation as the only predictor, and found a significant reduction of general pruning rate γ_g over generations for human rewarding ($\beta = -0.017$, $p < 0.01$, $CI = [-0.026, -0.0056]$) and human regretful environments ($\beta = -0.0078$, $p = 0.05$, $CI = [-0.016, 0.0013]$), respectively. For the specific pruning rate γ_s we found a similar reduction in both environments (human rewarding: ($\beta = -0.017$, $p < 0.01$, $CI = [-0.026, -0.0056]$); human regretful ($\beta = -0.0078$, $p = 0.05$, $CI = [-0.016, 0.0013]$)). These findings suggest that social learning led over the generations to solutions, which, on average, required more planning if done by individuals in isolation.

Given that the algorithm had a significant but temporary effect on following human performance, we investigated the effect of the algorithm on participants’ behavior. Figure 2d shows the difference between parameter estimates in human-only and hybrid chains. Not surprisingly, we observed a difference between the algorithm and humans in generation two, as we designed the algorithmic parameter to show a different bias than participants. For solutions following the algorithm we estimated for human regretful environments a lower specific pruning parameter ($\delta = -0.1$, $p = 0.03$, $CI = [-0.2, -0.006]$) and a higher general pruning parameter ($\delta = 0.056$, $p = 0.03$, $CI = [0.0014, 0.13]$). For human rewarding environments, we found in the same generation non significant differences between conditions (human rewarding: ($\delta = -0.046$, $p = 0.24$, $CI = [-0.21, 0.11]$); human regretful: ($\delta = 0.065$, $p = 0.08$, $CI = [-0.019, 0.13]$)). We found no significant difference between the two conditions in the following generations. On the one hand, these findings show that solutions directly following the algorithm are qualitatively different from the ones in human chains and that the algorithmic strategy is partially transmitted to participants following the algorithm. On the other hand, participants further down the transmission chain appear to reverse back to their typical strategies.

Finally, we investigated the rate at which participants in human-only and hybrid chains followed optimal strategies (see Supplementary Figure 4). We ran a logistic regression with the same variables as previously described, on whether a solution was optimal. For hybrid chains, we found an increased rate at which optimal solutions are discovered in generation 3 (β (*s.e.*) = 0.598 (0.147), $Z = 4.059$, $p < 0.001$, $CI = (0.310, 0.887)$) compared to human-only chain. However the difference quickly decayed and we did not find any significant difference in optimal solution discovery in final generations (β (*s.e.*) = 0.598 (0.147), $Z = 4.059$, $p < 0.001$, $CI = (0.310, 0.887)$). Correspondingly these findings do not support the hypotheses of a faster optimal solution discovery (H4) and sustained increase in discovery rate (H5) caused by the algorithm.

3.2 Algorithmic solutions are copied less, after controlling for scoring

Figure 2b depicts the average number of matching moves between second generation solutions (either human or algorithmic) and solutions in following generations (human). Despite their higher performance algorithmic solutions did not appear to be preserved, compared to their human counterparts. This finding may result from two opposite effects being at play. On the one hand, the higher reward of algorithm solutions could lead to a higher rate of copying. On the other hand, the mismatch with the inherent bias of participants might reduce copying.

To examine the mechanisms behind the algorithmic solution decay and a potential human content bias against algorithmic solutions, we conducted a set of exploratory analyses only on the third generation. We modelled the number of actions copied as a Poisson distribution with fixed effects for the previous solution’s (a) creator (algorithm or human) and (b) standardized reward (Supplementary Table 2). We added random effects to account for covariation due to individual participants and environments. We found an increased rate of copying of high scoring solutions (β (*s.e.*) = 0.394 (0.034), $Z = 11.423$, $p < 0.001$, $CI = (0.326, 0.461)$) and a lower rate of copying of algorithmic solutions (β (*s.e.*) = -0.189 (0.041), $Z = -4.625$, $p < 0.001$, $CI = (-0.269, -0.109)$). A model including an interaction between the two effects, did not significantly improved model fit ($\chi^2 = 1.35$, $df = 2$, $p = 0.51$), suggesting that the two effects were additive. These findings suggest that once controlling for reward magnitude, algorithmic solutions were copied at a lower rate than human solutions. We found no of significant interaction between the two main effects and type of environment.

We did not disclose to participants whether the previous solution they see is from an algorithm or another human, which rules out any bias against algorithmic solutions beyond the characteristics of the solution itself. If this were the case, we should expect to find reduced copying also of those human solutions that happened to show a higher number of large costs (and an increased copying rate of high performing solutions). We thus independently tested the hypotheses that both higher rewards and fewer large costs lead to increased

Figure 3: Average reward of the solutions of 100000 modeled agents. Human-only chains are depicted in blue, algorithm-only in orange, hybrid chains with a single algorithm (as in the experiment) in green and randomly mixed hybrid chains in red. On the left panel the environment favors the algorithmic bias, on the right panel it favors the human bias. We compare two type of content bias, one with a bias for higher performing solutions (solid) and a second with an additional bias to match the specific bias of the agent.

copy rates. A model predicting the number of copied actions in generation 2 to 8 in human-only chains (Supplementary Table 3) showed a positive effect of the previous solution’s reward (β (*s.e.*) = 0.393 (0.014), $Z = 28.162$, $p < 0.001$, $CI = (0.366, 0.420)$) and a negative effect of the previous solution’s number of large costs (β (*s.e.*) = -0.040 (0.011), $Z = -3.544$, $p < 0.001$, $CI = (-0.061, -0.018)$). Both factors interacted with environment type, suggesting that they were stronger in human regretful environments.

These findings suggest a content bias in human social learning that favors higher rewards and fewer large costs. Consequently, solutions that do not match human bias, such as those of the algorithm, are less well preserved. In the next section, we explore whether reduced copying rates can be overcome by repeated exposure to algorithmic solutions.

3.3 An agent-based model: Sustained performance improvements with repeated algorithmic exposure

We developed a simple agent-based model mimicking our experimental setup to theoretically explore the impact of biases on social learning in hybrid cultural evolution (see Supplementary Methods for details). We modeled task solutions as points in a two-dimensional space with two independent qualities. The dimension s^g represents the general quality of a solution, and the second dimension s^s the specialization of a solution, i.e., how adaptive (or maladaptive) it is to a specific environment. Thus, human-like agents are adapted in ‘human-rewarding’ environments and algorithmic agents in ‘human-regretful’ environments. Notice that we modeled human and algorithmic agents symmetrically. Hence, ‘human-only’ chains on ‘human rewarding’ environments are symmetric in their performance with an ‘algorithmic-only’ on ‘human regretful’ environments and vice versa (Figure 3).

As in the experiment, we constructed chains of 8 agents. Agents first assessed a perceived quality of the previous players’ solution. Depending on this perceived quality, they decide to copy it or to sample an entirely new solution. The perceived quality determines the content bias of the agent. We compare two types of content biases. Agents with an ‘adapted’ content bias consider both the score of the previous solutions and the match with their specialization (dashed in Fig. 3). Agents with a ‘utilitarian’ content-bias consider only the score of the previous solution (solid in Fig. 3). Agents sample new solutions from a distribution screened towards their adaptive bias. We initialized chains with a neutral solution.

Figure 3 shows the average reward over 8 generations. Line color represents chains composed of different agent types: human-only agents (blue), algorithmic-only agents (orange), and two hybrid chain with algorithmic and human agents. The ‘single-algorithm’ condition (green) represents the performance of hybrid chains similar to our experiment, with only one algorithm in generation 2. For ‘human-regretful’ environments (left panel), the algorithm in the second generation shows higher performance than their peers, and this boost is carried over to human-like agents in the next generation. Replicating our experimental results, the performance boost quickly decays back to the level of human-only chains (blue). The rate of decay is

much faster for agents with an adapted content bias (solid green line) than for agents with a utilitarian content bias (dashed green line). For ‘human-rewarding’ environments (right panel), the introduction of the algorithm (green) leads to a performance drop compared to ‘human-only’ chains (blue). Again, performance converges in following generations.

The model allows us to investigate a condition that was not tested in our experiment namely randomly mixing humans and algorithms. Randomly mixed hybrid chains (red) show a performance in-between the performance of solely-adapted, e.g., algorithmic agents in regretful human environments, and solely-misadapted, e.g., human agents in human regretful environments. For these ‘random-hybrid’ chains, agents with a ‘utilitarian’ content-bias (red, dashed) converge to higher average performance than agents with an ‘adapted’ content-bias (red, solid). However, in the first two generations, agents in mixed chains with an ‘adapted’ content-bias have a slight edge over their peers.

The agent-based model shows that different chains converge to a fixed value irrespective of algorithmic participation at the beginning of the chain. We could reproduce the experimental finding that a performance boost of an adapted algorithm is not sustained by following humans—especially for human agents with a content bias against algorithmic solutions. However, our model suggests sustained effects for well-mixed hybrid social learning.

4 Discussion

In this work, we investigated the impact of algorithmic strategies on social learning using a transmission chain experiment. We adapted the decision-making task by Huys et al. [23] to a transmission chain paradigm to test whether introducing an algorithm to increase the diversity in decision strategies can improve collective performance via social learning. In this task, people are known to show an aversive pruning bias in exploring the decision tree. As expected, we found evidence of a performance improvement over generations due to social learning. Adding an algorithm with a different problem-solving bias than humans temporarily improved human performance but improvements were not sustained in following generations. While humans did copy solutions from the algorithm, they appeared to do so at a lower rate than they copied other humans’ solutions with comparable performance.

Our first contribution is expanding previous research in cultural evolution by suggesting a relatively unexplored area of investigation, namely hybrid social learning. Scholars of cultural evolution have long investigated how social learning could lead to the unmatched explosion of human cultural complexity in comparison to non-human animals [42, 29, 43]. Similarly, we might ask if the advent of self-learning algorithms can influence cultural trajectories via hybrid human-algorithm social learning. Going one step beyond prior work that looked into cultural evolution via digital technology [8, 9], we suggest that in a hybrid society, algorithms may not be just a medium for cultural transmission, but they may play an active role in the production of new cultural artifacts. In particular, we suggest that successful hybrid social learning may occur when algorithms, either by design or by self-learning, show different biases than their human counterparts. Although many algorithms quickly adopt human biases [44–46], several others can increasingly learn from direct interactions with the environments rather than from human data, thereby potentially showing new behaviors and biases. Greater variance in problem-solving and copying skills has been associated with greater cultural variance [47, 48] and—as long as there are selection biases with regard to who to copy from—greater innovation. We looked at particular situations where human biases are known to constrain human performance [23], and therefore humans could most benefit from observing an algorithmic strategy.

In our experiment, we tested these hypotheses by introducing algorithmic players that adopted different decision-making strategies than human players. Investigating hybrid groups of human and algorithmic players provides the experimenter with the advantage to closely control the behavior of algorithmic agents while observing the effect on the rest of the group [49, 50, 14], yet, to the best of our knowledge, bots have not prominently featured in transmission chain experiments.

Our second contribution lies in our empirical findings. We showed that participants did not preserve algorithmic solutions if they were incongruent with their bias in avoiding large costs. Although human and algorithmic biases have been thoroughly investigated in their respective fields (psychology/economics and computer science), how the two interact together is still poorly understood. We show that learning from algorithms might be limited by the specific task and cognitive biases characterising human players. In our experiment, higher-performing solutions that were incongruent with human biases showed lower copying rates, and were consequently lost over generations. Such preference for copying congruent solutions may limit the accumulation of algorithmic solutions into human repertoire [48, 42]. These findings are in agreement with Griffiths et al. analytical work, where the authors suggest that when participant bias and solution fitness go against each other, superior solutions will not be maintained, in conditions of moderate

to high transmission noise [51]. In follow-up work, Thompson and Griffiths modeled cultural evolution in transmission chain experiments as being influenced by attraction towards preexisting biases and local innovations [52]. The authors experimentally showed that, if the two are in conflict, participants’ solutions converge to a middle ground. While their work models the effect of inductive biases on artifacts, our work focuses on the effect of content biases on copying. Yet, biases hinder the discovery of optimal solutions in both cases. Our work goes beyond their findings and suggests that even if an algorithm aids humans in achieving optimal solutions, human bias in what to copy can lead to those solutions being quickly lost in successive human-human transmissions, unless repeated exposure takes place.

We suggest that hybrid social interaction among human and algorithmic players may play an increasingly critical role in today’s digital society. A tradition in cultural evolution tries to generalise lab findings to the real-world [32, 53], although the limitations of such extrapolations are known [54]. Yet, investigating human-algorithmic social learning in the lab is the first step to study how these phenomena might unfold in the real-world, and how interactions in hybrid social systems may foster or hinder innovations and collective performance.

Designing algorithms to nudge collective behaviour may add to an already long list of ethical concerns in AI [55, 44, 46, 56]. Our results further suggest that even algorithms that could objectively improve human performance might be limited. Content and context biases (i.e. what people are more likely to copy and who they are willing to copy from) might limit hybrid social learning, especially in uncertainty, high cognitive demand, or high time pressure (*i.e.*, high transmission noise [52]). Under these conditions, humans are more likely to follow well-known and adaptive biases [57, 16].

Importantly, we acknowledge the limitations of our study, both in terms of generalisability and sample size. Future studies will need to address whether AI-human collaboration may be more successful in other domains or simpler tasks. In our experiment, we were interested in isolating cultural transmission by exposing participants to one previous solution only. This may limit the generalisability of our study. Outside the lab, people can copy from multiple models, which may give them the option to compare alternative solutions. Also, while in our experiment we tested the effect of a single algorithmic player, the frequency of encountering algorithmic generated solutions in the real-world may be higher. Our agent-based model (Figure 3) predicts that sustained improvement in performance might be observed with greater chances to copy from algorithms, although more work is needed. Finally, in our experiment people visited each environment only once. This likely reduced the effect of individual learning as well as gave participants inadequate feedback on their performance. Repeated unsuccessful feedback with the same environment before being exposed to an algorithmic solution might give participants additional opportunities to copy the algorithm, when algorithmic solutions are valuable.

In this work we focused on the transmission of behaviour, rather than the transmission of strategy itself. Social learning seems to be more effective when copying exact behaviors rather than reasoning and decision strategies [58, 29, 59]. Understanding why a solution works is not a prerequisite for successful cultural evolution [31]. Yet, more explicit communication between model and observer—*e.g.*, in the case of teaching [7]—could allow for better transmission of strategy. For example, communication of intention can improve human-algorithm cooperation [60]. Correspondingly, we speculate that an algorithm that communicates the reasoning behind a solution might be copied at a higher rate and allow following humans to better critically appraise their preexisting beliefs.

To conclude, in this work we found limited influence of bots on long-term human cultural evolution. Our findings do not exclude the possibility of algorithmic influences on human culture, but draw some limiting conditions. The relationship between biased human strategies and algorithmic strategies derived by self-play might look different outside the lab where more complex AI algorithms are at play. However, studying these phenomena in controlled environment is an important first step to understand hybrid social learning. In this study, we suggested that differences between human and algorithmic behavior might be relevant for the emerging properties of cultural evolution.

ethics

The study was approved by the ethics committee of the Max Planck Institute for Human Development.

Data Access

Raw data, processed data, preregistration, and code for the analysis is available on OSF: <https://osf.io/5j6es/>

Contributions

LB, IR and NP originated the idea; LB and NP developed the experimental and analytical methodology; LB and DG implemented the experiment; KvK ran the experiment; DG processed the data and performed the MLE; LB and DG visualized the data; LB, DG and KvK performed the statistical analysis; LB performed the simulation; LB, DG, KvK, and NP analysed the data; LB, DG, KvK, and NP wrote the original draft and preregistration; LB, TFM and NP reviewed and edited the manuscript; LB and NP supervised and administered the work. All authors provided critical feedback that helped to shape the research, analysis and manuscript, as well as, gave final approval for publication.

Acknowledgements

We are thanking Alex Hustler and Moataz Elmasry for their support on implementing and deploying the experiment, respectively. We are thanking Nicolas Paton for his support regarding the empirica framework. We are thanking Sarah Otterstetter for the great work on the illustrations.

References

- [1] Cade Metz. In Two Moves, AlphaGo and Lee Sedol Redefined the Future. *Wired*, 2016. URL <https://www.wired.com/2016/03/two-moves-alpha-go-lee-sedol-redefined-future/>.
- [2] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George Van Den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 10 2017. ISSN 14764687. doi: 10.1038/nature24270. URL <https://www.nature.com/articles/nature24270>.
- [3] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 1 2016. ISSN 14764687. doi: 10.1038/nature16961. URL <https://www.nature.com/articles/nature16961>.
- [4] Yuandong Tian, Jerry Ma, Qucheng Gong, Shubho Sengupta, Zhuoyuan Chen, James Pinkerton, and Larry Zitnick. Open-sourcing a new ELF OpenGo bot and related Go research. *AI Facebook*, 2019. URL <https://ai.facebook.com/blog/open-sourcing-new-elf-opengo-bot-and-go-research/>.
- [5] Minkyu Shin, Jin Kim, and Minkyung Kim. Measuring Human Adaptation to AI in Decision Making: Application to Evaluate Changes after AlphaGo. *arXiv*, 2020. URL <http://arxiv.org/abs/2012.15035>.
- [6] Anthony Strittmatter, Uwe Sunde, and Dainis Zegners. Life cycle patterns of cognitive performance over the long run. *Proceedings of the National Academy of Sciences of the United States of America*, 117(44):27255–27261, 11 2020. ISSN 10916490. doi: 10.1073/pnas.2006653117. URL <https://www.pnas.org/content/117/44/27255>.
- [7] Rachel L. Kendal, Neeltje J. Boogert, Luke Rendell, Kevin N. Laland, Mike Webster, and Patricia L. Jones. Social Learning Strategies: Bridge-Building between Fields. *Trends in Cognitive Sciences*, 22(7):651–665, 7 2018. ISSN 1879307X. doi: 10.1016/j.tics.2018.04.003. URL <http://www.cell.com/article/S1364661318300949>.
- [8] Alberto Acerbi. *Cultural Evolution in the Digital Age*. Oxford University Press, 2019. doi: 10.1093/oso/9780198835943.001.0001.
- [9] Simon Carrignon, R. Alexander Bentley, and Damian Ruck. Modelling rapid online cultural transmission: evaluating neutral models on Twitter data with approximate Bayesian computation. *Palgrave Communications*, 5(1):1–9, 12 2019. ISSN 20551045. doi: 10.1057/s41599-019-0295-9. URL <https://doi.org/10.1057/s41599-019-0295-9>.

- [10] Ilan Yaniv. Group diversity and decision quality: Amplification and attenuation of the framing effect. *International Journal of Forecasting*, 27(1):41–49, 1 2011. ISSN 01692070. doi: 10.1016/j.ijforecast.2010.05.009. URL <https://linkinghub.elsevier.com/retrieve/pii/S0169207010000944>.
- [11] Ramsey M. Raafat, Nick Chater, and Chris Frith. Herding in humans. *Trends in Cognitive Sciences*, 13(10):420–428, 10 2009. ISSN 13646613. doi: 10.1016/j.tics.2009.08.002.
- [12] Lu Hong and Scott E. Page. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16385–16389, 11 2004. ISSN 00278424. doi: 10.1073/pnas.0403723101.
- [13] Philip E. Tetlock and Philip E. Tetlock. *Expert Political Judgment*. Princeton University Press, stu - stud edition, 3 2017. ISBN 9780691128719. doi: 10.1515/9781400888818. URL <http://www.jstor.org/stable/j.ctt7spbt>.
- [14] Margaret L. Traeger, Sarah Strohkorb Sebo, Malte Jung, Brian Scassellati, and Nicholas A. Christakis. Vulnerable robots positively shape human conversational dynamics in a human-robot team. *Proceedings of the National Academy of Sciences of the United States of America*, 117(12):6370–6375, 3 2020. ISSN 10916490. doi: 10.1073/pnas.1910402117. URL <https://www.pnas.org/content/117/12/6370>.
- [15] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974. ISSN 00368075. doi: 10.1126/science.185.4157.1124. URL <http://science.sciencemag.org/>.
- [16] Peter M. Todd and Gerd Gigerenzer. Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, 16(3):167–171, 6 2007. ISSN 09637214. doi: 10.1111/j.1467-8721.2007.00497.x. URL <http://journals.sagepub.com/doi/10.1111/j.1467-8721.2007.00497.x>.
- [17] Herbert A. Simon. Bounded rationality. In *Utility and probability*, pages 15–18. Palgrave Macmillan Ltd., London, 1990.
- [18] Noel G. Machado and Jerrin Samuel. *Predictably Irrational Hiring*. Harper Audio New York, NY, 2021. doi: 10.1177/2631454120988415.
- [19] Thomas Gilovich, Dale Griffin, and Daniel Kahneman. *Heuristics and biases: The psychology of intuitive judgment*. Cambridge university press, 2002.
- [20] Niccolò Pescetelli and Nicholas Yeung. The Role of Decision Confidence in Advice-Taking and Trust Formation. *Journal of Experimental Psychology: General*, 150(3), 2020. ISSN 00963445. doi: 10.1037/xge0000960. URL <https://pubmed.ncbi.nlm.nih.gov/33001684/>.
- [21] George Loewenstein and Jon Elster. *Choice over time*. Russell Sage Foundation, 1992.
- [22] Joel Myerson and Leonard Green. Discounting of Delayed Rewards: Models of Individual Choice. *Journal of the Experimental Analysis of Behavior*, 64(3):263–276, 11 1995. ISSN 0022-5002. doi: 10.1901/jeab.1995.64-263. URL <http://doi.wiley.com/10.1901/jeab.1995.64-263>.
- [23] Quentin J.M. Huys, Neir Eshel, Elizabeth O’Nions, Luke Sheridan, Peter Dayan, and Jonathan P. Roiser. Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3):1002410, 3 2012. ISSN 1553734X. doi: 10.1371/journal.pcbi.1002410.
- [24] Arthur Guez, David Silver, and Peter Dayan. Scalable and efficient bayes-adaptive reinforcement learning based on Monte-Carlo tree search. *Journal of Artificial Intelligence Research*, 48:841–883, 11 2013. ISSN 10769757. doi: 10.1613/jair.4117. URL <https://www.jair.org/index.php/jair/article/view/10853>.
- [25] Quentin J.M. Huys, Níall Lally, Paul Faulkner, Neir Eshel, Erich Seifritz, Samuel J. Gershman, Peter Dayan, and Jonathan P. Roiser. Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences of the United States of America*, 112(10):3098–3103, 2015. ISSN 10916490. doi: 10.1073/pnas.1414219112. URL <https://www.pnas.org/content/112/10/3098>.

- [26] Falk Lieder, Owen X. Chen, Paul M. Krueger, and Thomas L. Griffiths. Cognitive prostheses for goal achievement. *Nature Human Behaviour*, 3(10):1096–1106, 10 2019. ISSN 23973374. doi: 10.1038/s41562-019-0672-9. URL <https://doi.org/10.1038/s41562-019-0672-9>.
- [27] Sam J. Gilbert, Arabella Bird, Jason M. Carpenter, Stephen M. Fleming, Chhavi Sachdeva, and Pei Chun Tsai. Optimal Use of Reminders: Metacognition, Effort, and Cognitive Offloading. *Journal of Experimental Psychology: General*, 149(3), 2019. ISSN 00963445. doi: 10.1037/xge0000652. URL <https://pubmed.ncbi.nlm.nih.gov/31448938/>.
- [28] L. Rendell, R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland. Why copy others? insights from the social learning strategies tournament. *Science*, 328(5975):208–213, 4 2010. ISSN 00368075. doi: 10.1126/science.1184719. URL <https://science.sciencemag.org/content/328/5975/208>.
- [29] L. G. Dean, R. L. Kendal, S. J. Schapiro, B. Thierry, and K. N. Laland. Identification of the social and cognitive processes underlying human cumulative culture. *Science*, 335(6072):1114–1118, 3 2012. ISSN 10959203. doi: 10.1126/science.1213969. URL <http://science.sciencemag.org/>.
- [30] Daniel Barkoczi and Mirta Galesic. Social learning strategies modify the effect of network structure on group performance. *Nature Communications*, 7(1):1–8, 10 2016. ISSN 20411723. doi: 10.1038/ncomms13109. URL www.nature.com/naturecommunications.
- [31] Maxime Derex, Jean François Bonnefon, Robert Boyd, and Alex Mesoudi. Causal understanding is not necessary for the improvement of culturally evolving technology. *Nature Human Behaviour*, 3(5):446–452, 2019. ISSN 23973374. doi: 10.1038/s41562-019-0567-9. URL <http://dx.doi.org/10.1038/s41562-019-0567-9>.
- [32] Alex Mesoudi and Andrew Whiten. The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509):3489–3501, 2008. ISSN 09628436. doi: 10.1098/rstb.2008.0129. URL <http://dx.doi.org/10.1098/rstb.2008.0129>.
- [33] Christine A. Caldwell and Ailsa E. Millen. Studying cumulative cultural evolution in the laboratory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509):3529–3539, 11 2008. ISSN 09628436. doi: 10.1098/rstb.2008.0133. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2607341/>.
- [34] Christine A. Caldwell and Ailsa E. Millen. Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, 29(3):165–171, 5 2008. ISSN 10905138. doi: 10.1016/j.evolhumbehav.2007.12.001.
- [35] F. C. Bartlett and Cyril Burt. Remembering: a Study in Experimental and Social Psychology. *British Journal of Educational Psychology*, 3(2):187–192, 6 1933. ISSN 00070998. doi: 10.1111/j.2044-8279.1933.tb02913.x. URL <http://doi.wiley.com/10.1111/j.2044-8279.1933.tb02913.x>.
- [36] Yoshihisa Kashima. Maintaining cultural stereotypes in the serial reproduction of narratives. *Personality and Social Psychology Bulletin*, 26(5):594–604, 5 2000. ISSN 01461672. doi: 10.1177/0146167200267007. URL <http://journals.sagepub.com/doi/10.1177/0146167200267007>.
- [37] Michael L. Kalish, Thomas L. Griffiths, and Stephan Lewandowsky. Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin and Review*, 14(2):288–294, 2007. ISSN 10699384. doi: 10.3758/BF03194066. URL <https://link.springer.com/article/10.3758/BF03194066>.
- [38] Joseph Henrich and Richard McElreath. The Evolution of Cultural Evolution. *Evolutionary Anthropology*, 12(3):123–135, 1 2003. ISSN 10601538. doi: 10.1002/evan.10110. URL <https://onlinelibrary.wiley.com/doi/full/10.1002/evan.10110>.
- [39] Berkeley J. Dietvorst, Joseph P Simmons, and Cade Massey. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1):114–126, 2015. ISSN 00963445. doi: 10.1037/xge0000033.
- [40] César A Hidalgo, Diana Orghian, Jordi Albo Canals, Filipa De Almeida, and Natalia Martin. *How Humans Judge Machines*. MIT Press, 2021.

- [41] Abdullah Almaatouq, Joshua Becker, James P. Houghton, Nicolas Paton, Duncan J. Watts, and Mark E. Whiting. Empirica: a virtual lab for high-throughput macro-level experiments. *Behavior Research Methods*, pages 1–14, 3 2021. ISSN 15543528. doi: 10.3758/s13428-020-01535-9. URL <https://doi.org/10.3758/s13428-020-01535-9>.
- [42] Claudio Tennie, Josep Call, and Michael Tomasello. Ratcheting up the ratchet: On the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1528): 2405–2415, 2009. ISSN 14712970. doi: 10.1098/rstb.2009.0052.
- [43] Alex Mesoudi and Alex Thornton. What is cumulative cultural evolution? *Proceedings of the Royal Society B: Biological Sciences*, 285(1880), 2018. ISSN 14712954. doi: 10.1098/rspb.2018.0712.
- [44] Francis Schrag. *Review of Weapons of math destruction: How big data increases inequality and threatens democracy*, volume 24. Crown Publishing Group, USA, 2017. ISBN 0553418815. doi: 10.14507/er.v24.2197.
- [45] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, 4 2017. ISSN 10959203. doi: 10.1126/science.aal4230. URL <http://science.sciencemag.org/>.
- [46] Joy Buolamwini and Timnit Gebru. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, 81:1–15, 1 2018. ISSN 2640-3498. URL <http://proceedings.mlr.press/v81/buolamwini18a.html>.
- [47] Michael Muthukrishna and Joseph Henrich. Innovation in the collective brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1690):20150192, 3 2016. ISSN 14712970. doi: 10.1098/rstb.2015.0192. URL <https://royalsocietypublishing.org/doi/10.1098/rstb.2015.0192>.
- [48] Joseph Henrich. Demography and Cultural Evolution: How Adaptive Cultural Processes Can Produce Maladaptive Losses—The Tasmanian Case. *American Antiquity*, 69(2):197–214, 4 2004. ISSN 0002-7316. doi: 10.2307/4128416.
- [49] Hirokazu Shirado and Nicholas A. Christakis. Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, 545(7654):370–374, 5 2017. ISSN 14764687. doi: 10.1038/nature22332. URL <https://www.nature.com/articles/nature22332>.
- [50] Malte Jung and Pamela Hinds. Robots in the Wild. *ACM Transactions on Human-Robot Interaction*, 7(1), 2018. ISSN 25739522. doi: 10.1145/3208975. URL <https://doi.org/10.1145/3208975>.
- [51] Thomas L. Griffiths, Michael L. Kalish, and Stephan Lewandowsky. Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509):3503–3514, 11 2008. ISSN 0962-8436. doi: 10.1098/rstb.2008.0146. URL <https://royalsocietypublishing.org/doi/10.1098/rstb.2008.0146>.
- [52] Bill Thompson and Thomas L. Griffiths. Human biases limit cumulative innovation. *Proceedings of the Royal Society B: Biological Sciences*, 288(1946): 20202752, 3 2021. ISSN 0962-8452. doi: 10.1098/rspb.2020.2752. URL <https://royalsocietypublishing.org/doi/10.1098/rspb.2020.2752>.
- [53] Andrew Whiten and Alex Mesoudi. Establishing an experimental science of culture: Animal social diffusion experiments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509):3477–3488, 11 2008. ISSN 09628436. doi: 10.1098/rstb.2008.0134. URL <https://royalsocietypublishing.org/doi/10.1098/rstb.2008.0134>.
- [54] Helena Miton and Mathieu Charbonneau. Cumulative culture in the laboratory: Methodological and theoretical challenges. *Proceedings of the Royal Society B: Biological Sciences*, 285(1879):20180677, 5 2018. ISSN 14712954. doi: 10.1098/rspb.2018.0677. URL <https://royalsocietypublishing.org/doi/10.1098/rspb.2018.0677>.
- [55] Shoshana Zuboff. The age of surveillance capitalism: The fight for a human future at the new frontier of power. *Yale Law Journal*, 129(5):1460–1515, 2020. ISSN 00440094. doi: 10.26522/brooklyn.v29i2.849.

- [56] Miles Brundage, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, Hyrum Anderson, Heather Roff, Gregory C Allen, Jacob Steinhardt, Carrick Flynn, Seán HÉigeartaigh, Simon Beard, Haydn Belfield, Sebastian Farquhar, Clare Lyle, Rebecca Crootof, Owain Evans, Michael Page, Joanna Bryson, Roman Yampolskiy, and Dario Amodei. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation, 2018. ISSN 23318422.
- [57] Mike Oaksford and Nick Chater. A Rational Analysis of the Selection Task as Optimal Data Selection. *Psychological Review*, 101(4):608–631, 1994. ISSN 0033295X. doi: 10.1037/0033-295X.101.4.608.
- [58] Iyad Rahwan, Dmytro Krasnoshtan, Azim Shariff, and Jean Francois Bonnefon. Analytical reasoning task reveals limits of social learning in networks. *Journal of the Royal Society Interface*, 11(93), 4 2014. ISSN 17425662. doi: 10.1098/rsif.2013.1211. URL <http://dx.doi.org/10.1098/rsif.2013.1211>.
- [59] Edwin J.C. Van Leeuwen, Josep Call, and Daniel B.M. Haun. Human children rely more on social information than chimpanzees do. *Biology Letters*, 10(11), 11 2014. ISSN 1744957X. doi: 10.1098/rsbl.2014.0487. URL <http://dx.doi.org/10.1098/rsbl.2014.0487> or via <http://rsbl.royalsocietypublishing.org>.
- [60] Jacob W. Crandall, Mayada Oudah, Tennom, Fatimah Ishowo-Oloko, Sherief Abdallah, Jean François Bonnefon, Manuel Cebrian, Azim Shariff, Michael A. Goodrich, and Iyad Rahwan. Cooperating with machines. *Nature Communications*, 9(1):1–12, 12 2018. ISSN 20411723. doi: 10.1038/s41467-017-02597-8. URL <https://www.nature.com/articles/s41467-017-02597-8>.