

Machine Culture

Levin Brinkmann,^{*†1} Fabian Baumann,^{†1} Jean-François Bonnefon,^{†2} Maxime Derex,^{†2,4} Thomas F. Müller,^{†1} Anne-Marie Nussberger,^{†1} Agnieszka Czaplicka,¹ Alberto Acerbi,³ Thomas L. Griffiths,⁵ Joseph Henrich,⁶ Joel Z. Leibo,⁷ Richard McElreath,⁸ Pierre-Yves Oudeyer,⁹ Jonathan Stray,¹⁰ Iyad Rahwan^{*†1}

* Correspondence: brinkmann@mpib-berlin.mpg.de; rahwan@mpib-berlin.mpg.de

† Equal contributions

¹ Max Planck Institute for Human Development, Center for Humans & Machines, Lentzeallee 94, Berlin 14195, Germany

² Toulouse School of Economics, Esplanade de l'Université, 31080 Toulouse Cedex 06, France

³ Department of Sociology and Social Research, University of Trento, Italy

⁴ Institute for Advanced Study in Toulouse, Esplanade de l'Université, 31080 Toulouse Cedex 06, France

⁵ Department of Psychology & Department of Computer Science, Princeton University, Princeton, NJ 08540

⁶ Department of Human Evolutionary Biology, Harvard University, 11 Divinity Ave, Cambridge, MA 02138, USA

⁷ DeepMind Technologies Ltd, London, United Kingdom

⁸ Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany

⁹ Inria, Flowers team, Université de Bordeaux, France

¹⁰ Center for Human-Compatible Artificial Intelligence, University of California, Berkeley, USA

Abstract:

The ability of humans to create and disseminate culture is often credited as the single most important factor of our success as a species. In this Perspective, we explore the notion of 'machine culture,' culture mediated or generated by machines. We argue that intelligent machines simultaneously transform the cultural evolutionary processes of variation, transmission, and selection. Recommender algorithms are altering social learning dynamics. Chatbots are forming a new mode of cultural transmission, serving as cultural models. Furthermore, intelligent machines are evolving as contributors in generating cultural traits—from game strategies and visual art to scientific results. We provide a conceptual framework for studying the present and anticipated future impact of machines on cultural evolution, and present a research agenda for the study of machine culture.

Introduction

The ability of humans to create and disseminate culture is considered the single most important factor in our species' dominance on earth ¹. The evolution of human culture has been the subject of extensive study in all of the behavioral sciences, including anthropology ¹, psychology ², cognitive science ³, biology ⁴, linguistics ^{5,6}, archaeology ⁷, sociology ⁸ and economics ⁹ (Box 1).

Cultural evolution exhibits key Darwinian properties. Culture is shown to exhibit variation, transmission, and selection, and evolves through the selective retention of cultural traits, as well as nonselective processes such as drift ¹⁰. Major shifts in any of these three Darwinian properties can greatly impact cultural evolution. For instance, between 1300 and 1600, European culture experienced successive major shifts due to increased exposure to Chinese technology such as gunpowder, which changed the nature of warfare (variation)¹¹; Gutenberg's invention of the printing press (transmission)¹²; and renewed interest in Classical ideas and values, such as Classical ideals of artistic expression, during the Renaissance (selection)¹³. When such substantial changes occur, they induce rapid and major impact on culture.

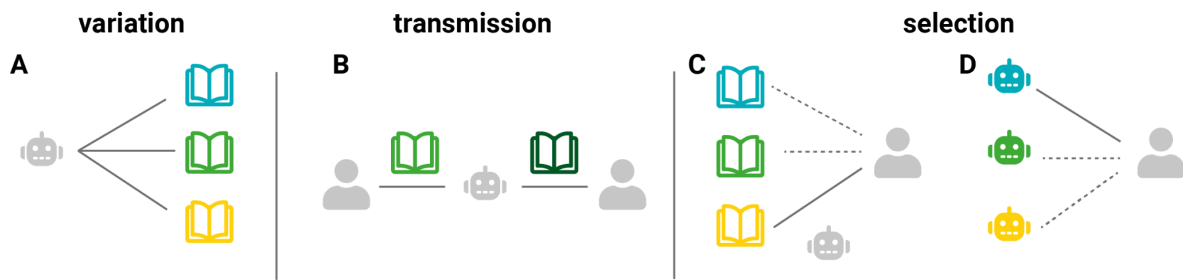


Figure 1: Examples of machine culture. **A.** Generation of novel cultural artifacts through machines. **B.** Machine transmits and potentially mutates cultural artifacts. **C.** Machine selects between different cultural artifacts. **D.** Human selects among diverse machines.

While new technologies have always affected the course of cultural evolution, in this article we argue that intelligent machines will exert a transformative influence on cultural evolution through their impact on all three Darwinian properties of culture: variation, transmission, and selection (Fig. 1). This process began in the early days of the Internet with machine-based content ranking by search engines and social media feed algorithms influencing what information people get from others. More recently, generative algorithms have begun participating in the creation of cultural traits themselves. We are not only observing a transformation of human culture but also its evolution into machine culture—culture mediated or generated by machines. This article aims to provide researchers across disciplines with a primer and a roadmap for navigating this monumental shift. As the impact of an increasingly digital society on cultural evolution has been explored elsewhere¹⁴, we specifically focus on the current and potential impact of intelligent machines on cultural evolution. For the purposes of this article, we use the terms “intelligent machines” and “artificial intelligence (AI) systems” interchangeably, with AI referring to the science and technology that allow machines to perform tasks that typically require human intelligence such as perceiving the environment, planning and executing actions, and adapting by learning from data or experience^{15,16}.

Examples of machine-mediated cultural evolution

We begin by presenting empirical evidence of machine cultural evolution, setting the stage for a detailed exploration through a framework that discusses instances where machines mediate or generate cultural traits from a cultural evolutionary perspective. Through four pivotal examples, we illustrate the diverse ways intelligent machines are transforming cultural evolutionary dynamics. Generative machines, such as text-to-image algorithms, are contributing to the variety of cultural traits. Models drawing upon reinforcement learning are pushing humans onto novel ground, for instance in the ancient game of Go and beyond. Large language models (LLMs) are facilitating the transmission of cultural knowledge and redefining the value of human intellectual skills. Meanwhile, transmission pathways are rewired by recommender systems selecting what and from whom humans learn. At first glance, the examples provided might seem to pertain to vastly different technological areas, and to translate into a collection of unrelated effects. However, even now, machines are beginning to integrate features from a range of the outlined technologies—reinforcement learning, for instance, is enhancing generative AI. Furthermore, these technologies are operating on multiple levels; generative AI not only generates novel ideas but also offers recommendations for their refinement.

Box 1: Glossary

Culture: information capable of affecting individuals' behaviors that are acquired from other individuals via social transmission

Cultural Evolution: the change of cultural information over time, the key properties for an evolutionary process are variation, transmission, and selection

Social Learning: learning that is influenced by the observation of another individual or their products.

AI: the science and technology enabling machines to perform tasks that typically require human intelligence such as perceiving the environment, planning and executing actions, and adapting by learning from data or experience.

Machines: intelligent machines. Used interchangeably with AI systems, thus referring to machines that may possess capabilities to perceive the environment, plan and execute actions, and adapt by learning from data or experience.

Variation: the existence of different cultural traits within a population. It represents the raw material on which other processes, like selection and transmission, operate. Humans and machines add to existing cultural variation through random and guided exploration, as well as recombination of existing cultural traits.

Transmission: the process by which cultural information, such as knowledge, behaviors, traditions, or practices, is passed from one individual to another through social learning mechanisms such as observation or teaching.

Selection: the process by which certain cultural traits, practices, or ideas become more or less prevalent within a population over time due to differential adoption.

Cultural recombination through generative AI

Generative AI has seen two major waves of innovation in recent years. The inception of Generative Adversarial Networks (GANs) by Goodfellow et al. in 2014 enabled the algorithmic generation of high-fidelity images¹⁷. GANs offered capabilities beyond the generation of lifelike images—they also have the ability to blend or interpolate—giving birth to novel creations such as fantasy lifeforms¹⁸. Subsequent advancement in 2022 saw the advent of diffusion-based text-to-image generative AI systems such as DALL·E, Midjourney, and Stable Diffusion. These models substantially enhanced the recombination power of these early models by generating high-resolution images conditioned on text descriptions^{19–21}. The original DALL·E, although now surpassed by other models in terms of image quality, demonstrated such recombination capabilities impressively¹⁹. For instance, when prompted to produce ‘an armchair in the shape of an avocado’, it creatively recombined these two distinct concepts (see Fig. 2).



Figure 2: Recombination of Visual Concepts. The avocado chair, synthesized by OpenAI's text-to-image generative AI, DALL·E, exemplifies the early stages of algorithmic cultural recombination. By seamlessly combining previously learned concepts—avocados and chairs in this case—the model showcases the ability to create coherent integrations of disparate elements and demonstrates novelty through recombination.

These models can thus increase cultural variation by helping humans to produce new and relevant recombinations, which are sometimes recognized as works of art, sold at prestigious auction houses ²². While recombination oftentimes forms the foundation of human creativity ²³, it is still debated how, or even if, machines can generate relevant content beyond the boundaries of human culture. Even simple latent representations can disentangle the semantic meaning of linguistic concepts ²⁴. Similarly, text-to-image models use language as a cognitive tool to disentangle and consequently recombine visual concepts ²⁵. However, these models build upon concepts harvested from human culture. As such, text-to-image models may be limited to the concepts defined and demonstrated by humans in the underlying dataset. However, generative AI models can produce novel "future art" by forecasting future art movements and by deliberately avoiding classification into established artistic movements—such creations have been evaluated as more creative than prestigious contemporary works by human artists ²⁶. When applied to engineering, similar methods can lead to the discovery of designs that are both novel and superior ²⁷.

Generative AI exemplifies the dual nature of machines as both cultural artifacts and creators thereof. On the one hand, generative algorithms are increasingly presented and, to some extent, recognized as authors of art ²². Simultaneously, machines are subjected to cultural processes of comparison, distribution, modification, and eventual abandonment.

Cultural innovation through reinforcement learning

In 2016, AlphaGo defeated Lee Sedol, the world champion Go player, with a series of four victories over five games. Remarkably, AlphaGo managed to surprise Sedol with distinctively nonhuman gameplay. In particular, move 37 in the second game was considered extremely unconventional,

estimated by AlphaGo itself to have a 1 in 10,000 chance of being made by a human ²⁸ (see Fig. 3.A). AlphaGo’s unconventional gameplay likely originated in its self-play training: While its training started with reproducing human gameplay from a large database, AlphaGo also trained through self-play, selecting promising but uncertain moves and evaluating their success against itself ²⁹. Thereby, it iteratively improved and developed novel game strategies. It was one of these innovative, self-trained strategies that took Lee Sedol by surprise. As it turned out, it was not even necessary for the model to start from learning human gameplay at all. The successor of AlphaGo, named AlphaGo Zero, ignored all of the accumulated Go knowledge of humankind ²⁹. Starting from a blank slate, it not only rediscovered human Go strategies but also developed strategies that surpassed those of its human creators.

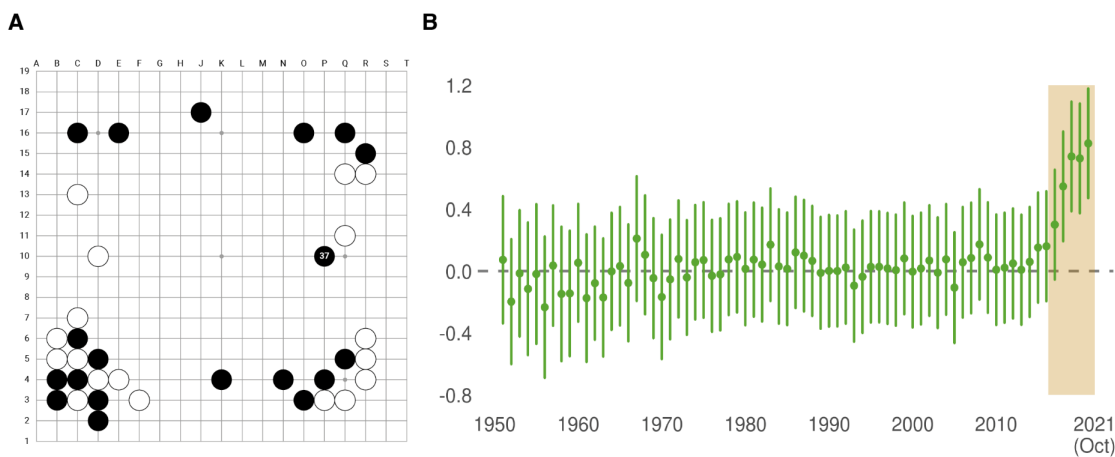


Figure 3: Go play before and after the introduction of AlphaGo. **A** AlphaGo, in its match against Go world champion Lee Sedol, made a highly unusual and strategic 37th move by placing its stone further from the edge, towards the center of the board, deviating from the traditional strategy of securing territory along the periphery during the early stages of the game. With this unconventional move, AlphaGo not only broke with centuries-old Go traditions but also paved the way for its ultimate victory in the match. **B** (reproduction based on ³⁰) Decision quality of professional Go players as evaluated by an algorithm performing at superhuman level. Decision quality significantly increased after Lee Sedol was beaten by AlphaGo on March 15, 2016 (shaded area).

The innovations generated by AlphaGo and AlphaGo Zero soon entered human culture, as shown by research comparing human gameplay before and after the algorithms’ introduction ³⁰. The decision quality, as measured by an open-source variant of AlphaGo Zero, showed very little improvement in human gameplay from 1950 to 2016, followed by a sudden improvement after the introduction of AlphaGo in March 2016 ^{31,32} (see Fig. 3.B). However, this improvement wasn’t solely due to humans adopting strategies developed by AlphaGo. It also reflected an unexpected shift, wherein humans started developing moves that were qualitatively distinct both from previous human moves and from the novel moves introduced by AlphaGo. In summary, AlphaGo served as an early, quantifiable exemplar of machine culture, generating novel cultural variations through genuine, nonhuman innovation. This was followed by a major transition into an even broader range of traits as the result of humans building on the previous discoveries made by machines. As the methods underpinning AlphaGo have been generalized to other games and extended to scientific problems, we anticipate a continued infusion of machine-generated discoveries across diverse domains of human culture ^{33,34}.

Language models transmit and revalue cultural knowledge

The release of ChatGPT, a widely accessible LLM, has revolutionized how we interact with machines, using it to learn, brainstorm, and refine ideas. Trained on extensive human text data, both historical and contemporary, LLMs act as models of human culture²⁵, facilitating cultural transmission across individuals and generations. In due course, students began requesting LLMs to complete their homework³⁵; knowledge workers used LLMs (at their own peril) to automatically extract and summarize required content³⁶; and software developers widely adopted LLMs as powerful code-writing assistants³⁷. LLMs not only serve as content creators but also, for better or worse, act as reservoirs of knowledge and providers of learning opportunities. Thereby, ChatGPT exemplifies social learning from machines.

As the capabilities and usages of LLMs continue to develop, the value of certain human skills will shift. Some skills may lose value quickly, especially in language-related and cognitively demanding occupations such as translation, copywriting, or proofreading³⁸. Occupations with more creative uses of language may follow, given that LLMs may soon surpass the creativity of humans as measured by standardized tests³⁹. Even though not all occupations will be affected in the short-term, a recent survey projects that around 20% of the workforce will experience LLMs impacting at least half of their tasks³⁸. It is important to note, however, that this projection is an early estimate and as such inherently uncertain in its accuracy and reliability. Meanwhile, other skills may gain in value — for instance, skills that allow efficient collaboration with LLMs, such as prompt engineering (i.e., the skilful writing of instructions to get LLMs to do what we want)⁴⁰. Consequently, human creativity is experiencing a remarkable shift⁴¹. It is not merely focused on generating final outputs but is increasingly evolving towards interactions with machines, progressing from explicit prompts to more natural conversations⁴². Workers who invest in related skills may outperform workers who do not, potentially accelerating the adoption of LLMs and further increasing the value of knowing how to collaborate with them.

Cultural rewiring through recommender systems

The digital age is so data-rich that it has become increasingly hard for humans to navigate available information. In this abundance, recommender systems that manage and filter information have a silent but increasingly important role that is easily overlooked given how seamlessly these systems have integrated into our everyday digital lives. These systems select and prioritize items based on a variety of explicit and implicit features, including personal interests, control settings, past user behavior, and the behavior of similar users. While recommender systems do not add variation in cultural traits, they demonstrably impact the selective retention and transmission of cultural traits.




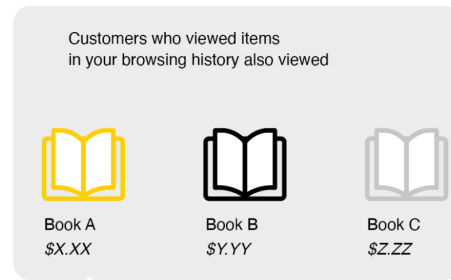
A Link recommendation**B Item recommendation**

Figure 4: Exemplary instances of cultural rewiring. Top panels represent the user interfaces, bottom panels provide a schematic depiction of the underlying mechanism. **A** Friendship recommendation on an online platform (e.g., X or LinkedIn). A new social tie (person A) is recommended to the green user. **B** Collaborative filtering for item recommendations. The recommender system suggests a book to the right user based on the correlations in co-purchases of other items (movie and music album) by the left user.

By promoting new social ties—such as suggesting who to follow on X (formerly Twitter), to date on Tinder, or to work with on LinkedIn—recommender systems alter the people to whom we are exposed to, ultimately changing the structure of our social networks and hence pathways of cultural transmission⁴³ (see Fig. 4.A). But recommender systems can also bypass network structures, exerting an even more direct impact on which cultural content or products we are exposed to. For example, e-commerce websites and streaming platforms deploy recommender systems to steer customers through the expansive array of available products based on content and collaborative filtering. Content filtering matches information about a customer’s consumption history with the attributes of all available products to make suggestions about related purchases: a customer who has recently purchased running shoes may receive suggestions for additional running equipment, matched to the shoes in price and design⁴⁴. Meanwhile, collaborative filtering⁴⁵ makes recommendations based on less obvious patterns of correlations in users’ profiles: if users A and B overlap in their previously consumed movies and music, the recommender system might suggest content of a completely different kind, for example, recommending to user A a book that user B liked (see Fig. 4.B). In sum, recommender systems influence cultural evolution by rewiring our social networks and modifying information flows such that they can substantially influence the dynamics of cultural markets^{46–48}.

A framework for machine-mediated cultural evolution

Building on these exemplary instances of how machine technologies may impact cultural evolution, we will now map out a systematic framework for studying the potential of machines to shape cultural evolutionary processes.

Culture has been defined as information capable of affecting individuals' behaviors that is acquired from other individuals via social transmission⁴⁹. Consequently, the science of cultural evolution examines the change of cultural information over time^{50,51}. Cultural information is represented by individual cultural traits, which can exist as cognitive representations or be expressed in behaviors or artifacts⁵². Culture evolves according to a similar process by which species change, that is, through the selective retention of cultural traits and through other nonselective processes, such as drift¹⁰. While there are ongoing discussions on how far the analogy between cultural and genetic evolution should be pushed^{53,54}, there is general agreement that culture exhibits the key properties of evolutionary systems that are Variation, Transmission, and Selection. These properties are not necessarily the result of distinct processes of human cognition and behavior, yet they offer a useful framework to analyze the multifaceted ways machines can influence cultural evolution (see Table 1 for a summary).

Variation refers to the existence of different cultural traits within a population. Transmission involves the spread of cultural information from one individual to another through social learning mechanisms, including observation and teaching. During this transmission, information losses often occur, affecting the preservation of cultural traits. Selection occurs when certain cultural traits are more likely to be adopted by individuals due to factors such as their usefulness, popularity, or compatibility with existing cultural practices. Crucially, the prevalence of specific traits over time can be influenced by both their selection and variations in transmission fidelity, contingent on the traits' characteristics. We anticipate that machine technologies will affect each of these three properties(see Fig. 4), which are likely to have transformative impacts on cultural evolution.

Table 1. Tentative conjectures about ways in which machines might shape the processes of cultural evolution.

	Machine Capability	Possible Impact on Culture
Variation	Learning at unprecedented scale and speed (e.g., reinforcement learning)	Emergence of solutions culturally alien to humans
	Superhuman model complexity	Generation of solutions inconceivable by collective human intelligence and human cultural evolution
	Incomparably broad and deep knowledge base	Creation of novel recombinations beyond the human horizon
	(Semi-)autonomous creativity (e.g., image generation models)	Facilitation or crowding out human participation in creative exploration
Transmission	Exposing and/or preserving documented cultural knowledge (e.g., LLM chatbots)	Enhanced retrieval and increased transmission fidelity of documented cultural knowledge
	Reproduction of human biases	Amplification or mitigation of existing biases; potential for increased cultural erosion
	Accelerated processing of empirical evidence	Facilitation of transmission of less compressed knowledge
	Leveraging the unique cognitive capabilities of both humans and machines	Expanding the collective capacity to maintain diverse cultural artifacts
Selection	Recommendation of social ties (e.g., link recommendation algorithm)	Shaping social networks, potentially inhibiting or enhancing serendipitous encounters
	Curating content (e.g., ranking algorithms, collaborative filtering)	Indirectly shaping social networks via content exposure
		Shaping incentives for human content creators (e.g., clickbait)
	Adapting to human feedback (e.g., RLHF)	Alignment of machines to human goals, potentially leading to unintended consequences (e.g., spread of highly believable myths)
	Machine learning from machine-generated content	Selection of machine-generated content as main driver of cultural evolution
	Adaptability of AI models to market/societal demands	Proliferation of appealing apps with varying alignment to human welfare
	Competing with humans in cultural production (e.g., poetry)	Specialization of humans and machines in distinct niches

Variation

Variation refers to the presence of diverse cultural traits within a population. Humans and machines contribute to this existing cultural variation through both random and guided exploration, as well as the recombination of existing cultural traits. However, machines, leveraging their unique capacities, can produce traits distinct from those produced by humans, thus potentially steering culture toward new paths.

Machines have the ability to learn individually at an unparalleled scale, enabling the discovery of novel cultural traits through extensive exploration. Thorough exploration is not unique to machines. Thomas Edison, for example, famously tested over 6,000 materials to find the most suitable filament for his incandescent light bulb⁵⁵. However, the computational speed of humans imposes time constraints that limit their exploration compared to machines⁵⁶. For instance, AlphaGo Zero needed only three days to play 4.9 million games against itself, achieving a superhuman level of proficiency²⁹. No human can amass that volume of experience individually; hence, we tend to rely heavily on preexisting, culturally evolved solutions that are socially learned^{57,58}. This approach narrows the scope of innovation to the cultural context of previous generations. AI systems like AlphaGo use reinforcement learning—based on iterative rounds of trial and error—and hence can transcend this cultural path dependency by starting from a blank slate and discovering innovative solutions through exploration alone⁵⁹. As a result, AI systems have the potential to generate culturally alien traits. In this respect, AlphaGo is in stark contrast to language models, which primarily learn to reproduce human cultural output.

Machines and humans employ guiding models—policies—that aid in efficiently exploring complex solution spaces, thereby enabling the discovery of solutions that would be unobtainable through random exploration alone. For instance, Mesoamerican skywatchers, via an early example of human astronomic modeling, developed a policy for optimal seed planting times, maximizing agricultural productivity⁶⁰. Similar to humans, algorithms like AlphaGo utilize complex policies to guide their search. The general approach to guided search is similar for both humans and machines: experience is used to update a predictive model of the world, which then informs actions that generate new experiences. Machines no longer require humans to explicitly formulate these guiding models. Instead, general-purpose neural networks are universal approximators, allowing functional relationships to be learned^{61,62}. This allows for unprecedented model complexity and consequently pushes the boundary of attainable solutions⁶³. Surpassing the capabilities of any human-conceived model, the neural network-based algorithm AlphaFold has achieved remarkable precision in modeling intramolecular relations, enabling it to tackle one of molecular biology's greatest challenges: predicting a protein's three-dimensional structure based solely on its DNA sequence⁶⁴. For centuries, protein structure prediction has drawn substantial intellectual contributions from scientists, even harnessing the collective intelligence of tens of thousands of citizen scientists⁶⁵. Nevertheless, AlphaFold's complex model of intramolecular relations, manifested in the form of a deep neural network, proved superior, eclipsing these extensive human efforts with remarkable proficiency.

Many contemporary algorithms, such as LLMs, have been trained on a cultural repertoire of unprecedented scale⁶⁶. For instance, some LLMs are trained on hundreds of billions of words⁶⁷. While size does not guarantee diversity⁶⁸, the degree to which machines can learn from human culture repertoires far exceeds what can be achieved by a single human being, allowing, for instance, individual language models to cover a broad spectrum of languages comprehensively⁶⁹. By combining knowledge from disjointed communities—spanning cultural and geographic contexts,

language barriers, or scientific disciplines—intelligent machines can produce novel recombinations that may be beyond humanity’s reach. Irrespective of whether LLMs truly exhibit understanding ⁷⁰, the mere recombination of cultural traits can lead to innovation and fuel cultural evolution ^{68,71}. Indeed, recombination is considered central to generating novel cultural traits ^{23,68,72–74}. Another way in which intelligent machines may increase the variety of cultural traits is by building on human knowledge while at the same time identifying and deliberately avoiding common human pathways. For instance, in the realm of scientific discoveries, intelligent machines may be designed to uncover scientifically plausible and promising “alien” hypotheses that would normally fall outside the focal point of contemporary scientific communities ⁷⁵. This does not imply, however, that intelligent machines are without their own limitations or blind spots, as we will discuss later.

But machines may also augment human exploration. Conditional generative AI, like DALL-E, allows users to steer the generative process through text prompts, enabling experimentation with diverse concepts and visualizations without the need for advanced artistic skills ¹⁹. Throughout human history, the effective size of the population that contributed to the creation and maintenance of cultural knowledge has constrained human cultural evolution ¹. Technology has had a varying impact on human cultural activities. On the one hand, it can democratize behavior. For instance, advancements in digital cameras and editing software have enabled individuals with limited technical knowledge to capture high-quality photographs and edit them effectively. On the other hand, the advent of new technologies can lead to increasingly specialized roles for humans. For instance, visual effects technology has led to the emergence of entirely new professions in film production, contributing to a cumulative increase in crew size ⁷⁶. Lastly, technological advancements can also lead to a reduction in human engagement in certain activities. For example, the advent of farming decreased the necessity for hunting, as people could grow food instead. It is yet to be seen whether the proliferation of increasingly intelligent machines will follow a similar trajectory. Will it democratize and increase human participation in creative fields, will it lead to new professions and increasing specialization, or could it lead to a decline in human-led creative expression, as we become increasingly reliant on intelligent machines ⁴¹? Undoubtedly, human cultural exploration, in science as in arts, will increasingly be in collaboration with machines ^{77,78}.

Transmission

The transmission of cultural information occurs via social learning, defined as learning that is influenced by the observation of another individual and/or its products ⁵⁷. This is in contrast to individual learning, where information is acquired, for instance, through trial and error. In humans, social learning tends to be imperfect, which means that transmission events are associated with a risk of losing cultural information. For example, if individuals observe a behavior incorrectly or do not fully understand it, they may not pass on the correct information to others. Over time, this can lead to the deterioration or even loss of cultural practices, which can have lasting consequences for a population.

Intelligent machines will increasingly be involved in the preservation and transmission of cultural information. Cultural evolution has supplied humans with increasingly efficient tools to preserve cultural information. The invention of writing, for instance, allowed humans to mitigate cultural loss by recording information in a more permanent way. Theoretical and empirical studies of cultural evolution have shown that the stability of cultural information strongly depends on both the size of the population that shares information ⁷² and the type of social learning mechanisms involved (with lower fidelity when transmission relies on mere observation than when it relies on verbal teaching, for

instance) ⁷⁹. Besides serving as a persistent medium of cultural storage analogous to a book, machines can learn to seek and transmit information ^{78,80} and can act as conversational and pedagogical agents, similar to teachers ⁸¹. This dual role has potential for drastically boosting cultural preservation by reducing cultural drift. For instance, LLMs have been harnessed to resurrect historical figures ⁸² and to revive languages teetering on the edge of extinction ⁸³. By fostering interactive cultural experiences for learners, these technologies can enhance the comprehension and retention of cultural information ⁸⁴.

As machines store and transmit cultural information, they may reproduce and transmit biases inherent to human culture, for example through biased training datasets; but they also offer potential to mitigate those. Machine learning models reproduce various content biases inherent to their training data, including gender bias, racial or ethnic bias, bias for negative and threat-related information, and socioeconomic bias ^{68,85-89}. For instance, many LLMs make implicit assumptions about the gender associated with professions like nurse and doctor, reflecting the demographic perspective dominant in the training data ⁹⁰. When an LLM is used for revision and enhancement of human-written text, it can transmit and/or reinforce these biases ⁶⁸. Numerous strategies exist to address bias and fairness in machine learning models ⁹¹. For instance, the demographic representation in the outputs of LLMs can be improved by using prompts featuring personas from a broad demographic spectrum ⁹². However, LLMs present another challenge: they struggle to accurately represent languages and communities for which there is limited training data ^{68,93}. This limitation implies that LLM-based copy-editing could inadvertently lead to cultural erosion within these underrepresented communities due to losses in text reproduction. Nevertheless, LLMs, with their capacity to handle vast quantities of training data, can contribute to the preservation and even enhancement of cultural diversity, provided that the training data is carefully curated with attention to diversity and representation.

Cultural transmission can be impacted not only by the replication of biases but also by disparities between human and algorithmic biases. Examples of biases found in humans are confirmation bias ⁹⁴, availability bias ⁹⁵, and a bias towards specific symmetries ⁹⁶. Despite their reputation for undermining optimal decision-making, biases can actually reflect optimal decisions within a particular socio-environmental context under cognitive constraints such as memory, computation, or experience ⁹⁷⁻¹⁰⁰. Similarly, machines may seize biases -- for instance, deep learning architectures assume spatial symmetries to improve training efficiency ⁹⁴. Importantly, humans and machines operate under different cognitive constraints ⁵⁶ and inhabit distinct socio-environmental contexts. This may give rise to idiosyncratic biases. For instance, due to their enhanced computational capabilities, machines may display more utilitarian rationality ⁹⁰. Remarkably, humans often anticipate more utilitarian behavior from them ¹⁰² and might indeed, as we will elaborate in the following section, enhance this tendency. Irrespective of their origin, the effects of biases on culture tend to strengthen through repeated transmission ^{98,103,104}. This can restrict the spectrum of solutions that a population can derive, and even highly effective solutions may not be sustained if they conflict with pre-existing biases ¹⁰⁵. As such, the increased cognitive diversity—for instance with regard to biases—within human-machine societies has the potential to expand the collective capacity to maintain diverse cultural artifacts. Artifacts that might not be maintained by humans could potentially be maintained by machines. Conversely, in the transmission between humans and machines, a misalignment of biases can increase the risk of information loss ¹⁰⁶.

Machines' increased computational capacity might additionally affect the feasibility of accumulating uncompressed information via a “big data” approach. The compressibility of information is the

inverse of its algorithmic complexity: Compressing information is achieved by creating a rule that is shorter than a complete list of the data itself ¹⁰⁷. Compression is a key feature of both human cognition and machine learning ¹⁰⁸; any information-processing system must address a general trade-off between a truthful representation of the “raw” information and constraints on computation. The extent and nature of these constraints, however, differ between humans and machines ⁵⁶. Compression processes have an important role in human transmission to mitigate cultural loss: Raw information oftentimes is not learnable because it is too complex, whereas a compressed rule might very well be. Human language, for example, retains its expressivity by becoming learnable via the evolution of compressed structure ¹⁰⁹. Similarly, scientists develop, transmit, and revise theories as compressed representations of knowledge. Machine learning has the potential to reduce some of the constraints resulting from human computation, as the amount of data that can be processed is vastly increased. Consequently, information might be increasingly transmitted with low compression – in the form of “big data” – when predictive power is of ultimate importance and for some applications, the necessity to derive and transmit highly compressed rules or theories might be reduced ^{110,111}. For instance, with the advent of AlphaFold, scientists might focus on collecting and preserving further ground truth data on molecule structures to refine future models rather than refining and transmitting theories of atomic interactions. While symbolic representations may remain crucial for efficient computation, algorithmic-assisted discovery could lessen the need for their transmission, as these representations could be easily regenerated ¹¹². However, theories may retain importance in shaping human understanding and intuition, serving as essential tools for conceptualizing knowledge—a function that the framework we present in this work aims to fulfill.

Selection

Culture evolves in part through the selective retention of cultural traits. In the context of machine culture, selection can occur at a level where machines select what and from whom humans learn, at a level where humans select machine behavior, and at a level where there is selection between humans and machines.

When it comes to what humans learn, social learning strategies shape what, when, and whom we copy. These strategies can be broadly categorized into content- and context-based strategies ^{113,114}: content-based strategies consider what is learned, favoring—for instance—social over nonsocial information ¹¹⁵. Conversely, context-based strategies attend to situational features, focusing on properties of a cultural model (e.g., their competence, success, prestige, knowledge, similarity), frequencies (e.g., most common behavior, rare behavior), or internal states of the learner (e.g., uncertainty, the cost of individual learning). Machines that help humans to navigate vast information spaces by (pre-)selecting cultural traits often reflect such social learning strategies.

For instance, content-based filtering algorithms aim to maximize the similarity between items a user previously showed interest in and unobserved items. Emulating context-based strategies in selection, ranking algorithms typically sort items according to a relevance score, which is based on the items’ popularity ^{116,117}. Concurrently, collaborative filtering (CF) algorithms detect hidden patterns between items and users, achieving recommendations about novel items to similar users without using additional exogenous information about individual items or users ¹¹⁸ (see Fig. 4B).

Another dimension along which algorithms may influence the selective retention of cultural traits pertains to social networks, which form the backbone of information exchange. In this context, social

ties are rewired as users follow algorithmic recommendations based on user attributes, such as popularity, or similarities in user preferences both in personal (X: “who to follow”) and professional domains (LinkedIn “People you might know”). Link recommendation algorithms have the potential to shape the overall evolution of social networks^{119–121}. X’s “who to follow” recommendation was observed to disproportionately benefit those users who were already the most popular, fueling “the rich get richer” dynamics¹²¹. A growing body of research documents a complex but persistent and critical relationship between social networks’ structure and collectives’ ability to collaborate, coordinate, and solve problems^{122–124} that ultimately shape cultural repertoires^{73,125}.

While machines have most commonly relied on exploiting explicit user preferences and historical behavior (e.g., ratings, engagement), there has been a growing interest in considering users’ internal states to improve algorithmic recommendations. For instance, users might be uncertain about their preferences -- especially in domains in which they lack expertise. Bayesian approaches can model users’ uncertainty and be used to update algorithmic recommendations as the user interacts with the system¹²⁶. As another example, recommender systems might account for the cognitive cost of exploring items or learning about them by prioritizing items that are easier for users to evaluate or learn¹²⁷. Affective recommender systems¹²⁸ use techniques such as natural language processing to make inferences about users’ emotional states, and even combine them with other context information such as the recommendation domain (e.g., music or movies)^{129–131}. While recommender systems have so far been mostly shaping user preferences implicitly (e.g., by optimizing the position or ranking of content), LLMs may accelerate developments where users are increasingly persuaded explicitly through interactive argumentation.

Downstream consequences of selection by machines may often be specific to particular environments, algorithmic models, and feedback loops. However, one feature generalizing across various contexts is that algorithms—by design of underlying business models—are often geared towards maximizing user engagement for profit¹³². In social networks, this may be achieved by promoting content congruent with users’ past engagement or ingroup attitudes¹³³, or content that humans inherently attend to, such as emotionally and morally charged content^{134,135}. One example for this is information that relates to threat or elicits disgust, as shown in transmission chain experiments inspired from cultural evolutionary theory¹³⁶. The algorithmic amplification of such content may then feed back into human social learning, for instance inflating beliefs about the normative value of expressing moral outrage^{137,138}, increasing out-group animosity¹³⁹, or by creating echo chambers and filter bubbles^{140–142}. It is important to note that user engagement is a signal of value to both users and platforms deploying algorithms, connecting them in complex feedback loops^{143,144}: machines such as recommender systems react to user engagement, selecting types of content people engage with. Users also react to recommender systems, both directly in terms of clicking, viewing, purchasing but also in terms of what they produce, as content creators anticipate what will receive the widest distribution. These feedback loops, but also deliberate product design choices along with policy approaches, provide potent leverage points for aligning recommender systems with human values¹⁴⁵. A promising approach to addressing this challenge could involve considering potential misalignments between users’ engagement and their own preferences and identifying the boundary conditions determining when maximizing user engagement enhances user welfare and when it produces the opposite effect¹⁴⁶. Algorithmic systems more generally offer powerful ways to bridge social divides, for instance by designing selection policies that steer users’ attention to content that increases mutual understanding and trust^{147,148}, or by identifying and promoting links in social networks that can effectively mitigate

polarizing dynamics^{149,150}. Machine selection can also be deliberately geared towards fostering content diversity¹⁵¹, or towards maximizing agreement among humans with diverse preferences¹⁵².

Human preferences, in turn, can directly shape machine behavior, in particular through reinforcement learning with human feedback (RLHF)^{153,154}. One way to conceptualize this is by viewing machines as students that generate arrays of solutions, with humans acting as teachers who select the most suitable ones. This process can nudge machines towards desirable properties like helpfulness, honesty, and harmlessness¹⁵⁴. However, harvesting human annotators' preferences at scale through RLHF can induce machines to adopt behaviors unintended by the deploying organizations, exemplified by chatbots that increasingly endorse inflated political views or express a heightened desire to avoid shutdown¹⁵⁵. Human influence on machine behavior also occurs through more subtle pathways: through training on human text alone LLMs picked up a tendency to repeat users' stated views¹⁵⁵. Machines' attention to human feedback may create selection pressures towards pleasing human interlocutors. For instance, humans may favor machines catering to non-factual stories and narratives that match concepts and ideas preferred by human cognition¹⁵⁶, such as those appealing to intuitive expectations about the natural world^{157,158}, or to specific religious practices¹⁵⁹.

Humans select machine behavior also through direct creation and curation of training data for machine learning, and through more indirect interactions with machine-generated outputs. Despite efforts to watermark machine-created content¹⁶⁰, machine and human-made content will increasingly intertwine¹⁶¹. For instance, it is estimated that many supposedly human-written texts on crowdsourcing platforms are already augmented by machines¹⁶². Consequently, it seems inevitable that future machines will be trained on mixed human-machine content, forming part of a larger feedback cycle between content generation and selection. The human element in this cycle may prove crucial, for instance in preventing 'model collapse' -- a dynamic where repeated training on machine-generated data narrows its outputs to very few traits¹⁶¹. This phenomenon stems from a classification bias that favors more prevalent classes¹⁶³, a bias that is amplified through successive iterations of learning¹⁶⁴. By contrast, when encountering similar challenges, human culture might preserve diversity by utilizing biases with counteracting effects, such as endorsing local conformity¹⁶⁵. That said, it seems likely that machines could recover similar strategies to maintain cultural diversity without human involvement.

Yet another pathway for human selection on machine behavior pertains to general machine properties. For instance, humans may choose between different intelligent machines available on the market based on factors such as preferences, cost, usefulness, harmfulness, and alignment with regulatory requirements. As such, the language model LLaMA recently gained attention for its relatively smaller size, making it more cost-effective to use¹⁶⁶. Conversely, ChatGPT outperformed many comparable models due to its superior accessibility and helpfulness¹⁵⁴. Human demands towards machines are likely to change over time, shaping the trajectory of machine culture similar to other historic cases where technologies evolved in response to changing human demands (e.g., the wheel from wooden wheels on carriages to rubber tires on automobiles). Crucially, if intelligent machines are designed and evaluated by non-representative experts, these systems run the risk of unintentionally reproducing and intensifying the biases inherent to their selectors¹⁶⁷.

Selection is also bound to occur at a level where machines or humans are favored over one another. For instance, the ability of machines to process vast amounts of information both quickly and accurately affords them a competitive edge in numerous cognitive tasks, such as strategic gameplay

and information retrieval. Relatedly, due to their cost-effectiveness and efficiency, intelligent machines may grow into the main workforce across various professional domains^{38,168}. However, analogous to how the invention of the car did not diminish interest in running as a sport, the proliferation of machines may not curtail human interest in intellectual pursuits; instead, it might redirect the focus from necessity to leisure and entertainment. Meanwhile, across various contexts, humans might favor other humans over machines due to their shared experience. For instance, even though present-day machines can conceive messages perceived as more empathic than messages conceived by humans^{169,170}, such messages may be perceived differently once recipients become aware of their artificial origin. This phenomenon, which could be referred to as the “artificial empathy paradox”¹⁷¹ may, at least in part, arise from the very fact that empathizing is effortful for humans¹⁷² and that, as such, it conveys a motivational social signal to others that is made void by machine-involvement. Selection between humans and machines does not imply that one agent dominates over the other in any cultural niche; oftentimes, one will augment -- rather than fully replace -- the other; other times, the presence of the other may trigger the development of new skills, or roles.

Grand challenges and open questions

We now suggest a broad research agenda for computational and behavioral scientists interested in the phenomenon of Machine Culture.

Measurement

One major open challenge is to quantify how much of human cultural dynamics can be attributable to algorithmic processes. For instance, it is difficult to completely disentangle the effects of ranking and recommendation algorithms on culture from alternative processes of human social learning, such as communication technology, institutions, and social practices. Since the inception of human culture, derived tools have had an important role in shaping cultural processes, making the establishment of a baseline a challenging question in itself. Getting good estimates is a precondition to optimizing for the usefulness of these algorithms while avoiding undesirable cultural impact such as polarizations¹⁷³. This challenge is reflected in research on “filter bubbles”, which has moved from considering algorithmic curation as decisive force shaping online engagement¹⁴¹ towards acknowledging the influence of users’ own choices on social media ranking algorithms¹⁷⁴ and search engines¹⁷⁵; thus highlighting the intricate feedback loops between machine and human decisions. While we hope that it will be possible to find appropriate metrics for the cultural impact of machines, the intricacy of this problem qualifies it as an open grand challenge.

A complementary question is how to quantify the influence of machine-generated artifacts – e.g., artwork, literature, music – on human cultural production in these areas. As generative AI becomes more commonplace, distinguishing intrinsic human culture from machine-generated culture or machine-influenced culture becomes even more challenging, especially as watermarking techniques may not be universally adopted. Despite popular media claims about machine-generated art developing its own unique style¹⁷⁶, we do not yet have a reliable way of verifying these claims, let alone assessing machine-to-human cultural transmission¹⁰⁶.

Another measurement concern relates to the quantification of the cultural regularities encoded in LLMs and other AI models. Even prior to the rise of LLMs, social media platforms like Facebook

already possessed fairly detailed quantitative models of cultural regularities and differences¹⁷⁷, developed mainly for the purpose of marketing to particular demographics. However, as LLMs are trained from curated datasets, and fine-tuning using human feedback, quantifying the biases that are introduced and/or mitigated by these models is crucial^{178,179}.

Societal Decision-making

We currently observe a rapid increase in the diversity of AI models, including LLMs, accelerated by the open-source software movement. However, market forces, such as regulation and market power, may result in a world dominated by a small number of monolithic models. This raises the possibility of reduction in cultural diversity, as major social, political, and economic forces try to shape global machine culture to match their preferences. This process may be amplified by feedback loops, in which LLMs train on an ongoing basis from synthetic data, or from human data that contains much machine-generated text. Preliminary evidence points to the possibility of model collapse, with the models losing diversity and converging to a state with low variance¹⁶¹.

Conversely, we face a potential 'Tower of Babel' scenario. As AI models become increasingly personalized, conforming to and reinforcing our individual worldviews, they risk engendering an unprecedented fragmentation of our shared perception of the world. In the biblical story (Genesis 11:1-9), the construction of the tower led to a divine intervention that scattered humanity and confounded languages. Drawing a parallel, if we continually interact with machines that echo and affirm our preconceived notions, we risk isolating ourselves within ideologically and culturally homogenous echo chambers. Such fragmentation can stifle meaningful dialogue, breed misunderstanding, and, ultimately, fracture our shared future vision.

Against this background, a key research agenda is to quantify the degree to which a given AI model, or ecosystem of models, exhibits uniformity or diversity. It is also imperative to understand what constitutes a 'healthy' level of diversity, that retains local sovereignty while also fostering collective human flourishing.

Ensuring that AI models, such as LLMs, reflect the beliefs and values of a given community requires mechanisms for societal decision-making about what knowledge goes into the models¹⁸⁰.

Furthermore, humans exhibit variation in their ethical expectations towards machines, both within and across cultures¹⁸¹. This raises questions about how to best aggregate diverse, potentially conflicting preferences to arrive at an agreeable outcome¹⁸². A number of interesting ideas are emerging, from voting on different algorithmic policy-makers¹⁸³ to utilizing LLMs to summarize diverse human opinions¹⁸⁴ and generate consensus statements¹⁵².

A related challenge is how to ensure long-term monitoring of Machine Culture. Similar to the notion of human-in-the-loop control of intelligent machines, we can aspire towards society-in-the-loop control of the complex phenomena of Machine Culture¹⁸⁵.

Suppose a community knows which cultural beliefs and values it wants to encode in an AI model. The next question is how to ensure that these are indeed present in the model. One approach is to carefully curate the dataset on which the model is trained^{186,187}. Another, increasingly used approach is fine-tuning based on RLHF¹⁸⁸. Yet another approach is to fine-tune using constitutional rules provided by humans¹⁸⁹. The relative merits and drawbacks of these various approaches are still not

well understood. There is a need for methods to check which cultural beliefs and values have been learned, inferred, or encoded in a given AI model and the degree to which it aligns with a target culture.

Long-term Dynamics & Optimization

In all likelihood, the future of culture will be hybrid, with cultural artifacts—scientific theories, industrial processes, art, literature—being created by a combination of human and machine intelligence. This raises a suite of open questions relating to the long-term dynamics of human-machine co-evolution. These dynamics may lead to diverse phenomena ranging from different forms of human-machine mutualism, to Red Queen effects (an evolutionary arms race between humans and machines) that characterize the co-evolution of both forms of intelligence ¹⁹⁰.

A related question is how to optimize the aforementioned dynamics, in order to combine human and machine intelligence in an ideal or safe manner ⁷⁵. This may be relevant to questions of risk mitigation. Although experts disagree on the timescales and the degree of risk involved, the potential of superhuman Artificial General Intelligence (AGI) poses a possible existential threat to the human species ¹⁹¹. Cultural evolution provides a useful framework for navigating this challenge. Specifically, cultural evolution processes take place today at multiple scales, with human collectives—e.g., companies, universities, institutions, cities, nation-states—acting as the units of selection ¹⁹². This multi-level selection can, in principle, operate at the level of human organizations augmented by intelligent machines and (eventually) superhuman AGI. Engineering this evolutionary process can provide means for ensuring human survival and agency in the long run.

Conclusion

We asked GPT-4 to first write a compressed version of this Perspective and then to provide a conclusion. It suggested the following (minimal editing to align nomenclature has been applied). The symbiosis of human and machine intelligence is forging a new epoch of cultural evolution. This Perspective highlights the transformative role of intelligent machines in reshaping creativity, redefining skill value, and altering human interactions. Central to the discourse is the triad of cultural evolution: variation, transmission, and selection, and how machines interface with each. The interaction is multifaceted, from generative AI birthing novel cultural artifacts to recommendation algorithms influencing individual perspectives. However, the crux remains in understanding and navigating the challenges and opportunities that arise from this hybridization of culture. As the imprints of intelligent machines grow deeper, it's imperative to ensure a harmonious co-creation of culture where both human and machine augment, rather than eclipse, each other. This not only broadens the horizons of cultural exploration but also fortifies the tapestry of human experience in the age of intelligent machines.

Author contributions

Conception: LB, IR; Manuscript preparation: LB, FB, JFB, MD, TFM, AMN, IR; Critical review, commentary or revision: AC, AA, TLG, JH, JZL, RM, PYO, JS; Supervision: IR

Acknowledgments

LB thanks Michiel Bakker for useful discussions. The authors thank Bramantyo Supriyatno for supporting the formatting of the manuscript. The conclusion was created by GPT-4 based on a summary (also created by GPT-4) of this manuscript with minimal editing to align nomenclature. JFB and MD acknowledge IAST funding from the French National Research Agency (ANR) under grant no. ANR-17-EURE-0010 (Investissements d’Avenir programme)

Competing interests

The authors declare no competing interests.

References

1. Henrich, J. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. (Princeton University Press, 2016). <https://doi.org/10.1515/9781400873296>.
2. Heyes, C. *Cognitive Gadgets: The Cultural Evolution of Thinking*. (Harvard University Press, 2018). <https://doi.org/10.4159/9780674985155>.
3. Thompson, B., van Opheusden, B., Sumers, T. & Griffiths, T. L. Complex cognitive algorithms preserved by selective social learning in experimental populations. *Science* **376**, 95–98 (2022). <https://doi.org/10.1126/science.abn0915>.
4. Whiten, A. Cultural Evolution in Animals. *Annu. Rev. Ecol. Evol. Syst.* **50**, 27–48 (2019). <https://doi.org/10.1146/annurev-ecolsys-110218-025040>.
5. Gray, R. D. & Atkinson, Q. D. Language-tree divergence times support the Anatolian theory of Indo-European origin. *Nature* **426**, 435–439 (2003). <https://doi.org/10.1038/nature02029>.
6. Kirby, S., Cornish, H. & Smith, K. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proc. Natl. Acad. Sci.* **105**, 10681–10686 (2008). <https://doi.org/10.1073/pnas.0707835105>.
7. Shennan, S. *Genes, memes, and human history: Darwinian archaeology and cultural evolution*. (Thames & Hudson, 2002). <https://doi.org/10.1017/S0003598X00061913>.
8. Kiley, K. & Vaisey, S. Measuring stability and change in personal culture using panel

- data. *Am. Sociol. Rev.* **85**, 477–506 (2020). <https://doi.org/10.1177/0003122420921538>.
9. Mokyr, J. *A Culture of Growth: The Origins of the Modern Economy*. (Princeton University Press, 2017). <https://doi.org/10.1515/9781400882915>.
 10. Mesoudi, A., Whiten, A. & Laland, K. N. Perspective: Is Human Cultural Evolution Darwinian? Evidence Reviewed From The Perspective Of The Origin Of Species. *Evolution* **58**, 1–11 (2004). <https://doi.org/10.1111/j.0014-3820.2004.tb01568.x>.
 11. Needham, J. Science and Civilisation in China | History of science and technology. in *Chemistry and Chemical Technology; Pt. 7: Military Technology; The Gunpowder Epic* vol. 5 (Cambridge University Press, 1986).
<https://doi.org/10.1017/S0035869X0016486X>.
 12. Eisenstein, E. L. *The printing press as an agent of change*. vol. 1 (Cambridge University Press, 1980). <https://doi.org/10.1017/CBO9781107049963>.
 13. Mesoudi, A. Culture and the Darwinian Renaissance in the social sciences and humanities: For a special issue of the *Journal of Evolutionary Psychology* , “The Darwinian Renaissance in the Social Sciences and Humanities”. *J. Evol. Psychol.* **9**, 109–124 (2011). <https://doi.org/10.1556/jep.9.2011.29.1>.
 14. Acerbi, A. *Cultural Evolution in the Digital Age*. (Oxford University Press, 2019).
<https://doi.org/10.1093/oso/9780198835943.001.0001>.
 15. Russell, S. & Norvig, P. *Artificial Intelligence: A Modern Approach*. (Prentice Hall Press, 2009).
 16. Kurzweil, R., Richter, R., Kurzweil, R. & Schneider, M. L. *The age of intelligent machines*. vol. 580 (MIT press Cambridge, 1990).
 17. Goodfellow, I. J. *et al.* Generative Adversarial Networks. Preprint at <http://arxiv.org/abs/1406.2661> (2014).
 18. Epstein, Z., Boulais, O., Gordon, S. & Groh, M. Interpolating GANs to Scaffold Autotelic Creativity. Preprint at <http://arxiv.org/abs/2007.11119> (2020).
 19. Ramesh, A. *et al.* Zero-Shot Text-to-Image Generation. Preprint at

- <https://doi.org/10.48550/arXiv.2102.12092> (2021).
20. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C. & Chen, M. Hierarchical Text-Conditional Image Generation with CLIP Latents. Preprint at <http://arxiv.org/abs/2204.06125> (2022).
 21. Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. Preprint at <https://doi.org/10.48550/arXiv.2112.10752> (2022).
 22. Epstein, Z., Levine, S., Rand, D. G. & Rahwan, I. Who Gets Credit for AI-Generated Art? *iScience* **23**, 101515 (2020). <https://doi.org/10.1016/j.isci.2020.101515>.
 23. Thagard, P. & Stewart, T. C. The AHA! Experience: Creativity Through Emergent Binding in Neural Networks. *Cogn. Sci.* **35**, 1–33 (2011). <https://doi.org/10.1111/j.1551-6709.2010.01142.x>.
 24. Mikolov, T., Yih, W. & Zweig, G. Linguistic Regularities in Continuous Space Word Representations. in *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* 746–751 (Association for Computational Linguistics, 2013).
 25. Colas, C., Karch, T., Moulin-Frier, C. & Oudeyer, P.-Y. Language and culture internalization for human-like autotelic AI. *Nat. Mach. Intell.* **4**, 1068–1076 (2022). <https://doi.org/10.1038/s42256-022-00591-4>.
 26. Elgammal, A., Liu, B., Elhoseiny, M. & Mazzone, M. CAN: Creative Adversarial Networks, Generating ‘Art’ by Learning About Styles and Deviating from Style Norms. Preprint at <http://arxiv.org/abs/1706.07068> (2017).
 27. Wang, Y., Shimada, K. & Barati Farimani, A. Airfoil GAN: encoding and synthesizing airfoils for aerodynamic shape optimization. *J. Comput. Des. Eng.* **10**, 1350–1362 (2023). <https://doi.org/10.1093/jcde/qwad046>.
 28. Metz, C. In Two Moves, AlphaGo and Lee Sedol Redefined the Future. *Wired* (2016).
 29. Silver, D. *et al.* Mastering the game of Go without human knowledge. *Nature* **550**, 354–359 (2017). <https://doi.org/10.1038/nature24270>.

30. Shin, M., Kim, J., van Opheusden, B. & Griffiths, T. L. Superhuman artificial intelligence can improve human decision-making by increasing novelty. *Proc. Natl. Acad. Sci.* **120**, e2214840120 (2023). <https://doi.org/10.1073/pnas.2214840120>.
31. Choi, S., Kim, N., Kim, J. & Kang, H. How Does AI Improve Human Decision-Making? Evidence from the AI-Powered Go Program. SSRN Scholarly Paper at <https://doi.org/10.2139/ssrn.3893835> (2022).
32. Shin, M., Kim, J. & Kim, M. Human Learning from Artificial Intelligence: Evidence from Human Go Players' Decisions after AlphaGo. *Proc. Annu. Meet. Cogn. Sci. Soc.* **43**, (2021). <https://doi.org/10.5281/zenodo.5095146>.
33. Schrittwieser, J. *et al.* Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature* **588**, 604–609 (2020). <https://doi.org/10.1038/s41586-020-03051-4>.
34. Discovering faster matrix multiplication algorithms with reinforcement learning | Nature. <https://www.nature.com/articles/s41586-022-05172-4>.
35. Kasneci, E. *et al.* ChatGPT for good? On opportunities and challenges of large language models for education. *Learn. Individ. Differ.* **103**, 102274 (2023). <https://doi.org/10.1016/j.lindif.2023.102274>.
36. Wagner, G., Lukyanenko, R. & Paré, G. Artificial intelligence and the conduct of literature reviews. *J. Inf. Technol.* **37**, 209–226 (2022). <https://doi.org/10.1177/02683962211048201>.
37. Chen, M. *et al.* Evaluating Large Language Models Trained on Code. Preprint at <https://doi.org/10.48550/arXiv.2107.03374> (2021).
38. Eloundou, T., Manning, S., Mishkin, P. & Rock, D. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. Preprint at <https://doi.org/10.48550/arXiv.2303.10130> (2023).
39. Stevenson, C., Smal, I., Baas, M., Grasman, R. & van der Maas, H. Putting GPT-3's Creativity to the (Alternative Uses) Test. Preprint at <https://doi.org/10.48550/arXiv.2206.08932> (2022).

40. Popli, N. How to Get a Six-Figure Job as an AI Prompt Engineer.
<https://time.com/6272103/ai-prompt-engineer-job/> (2023).
41. Epstein, Z., Hertzmann, A., & THE INVESTIGATORS OF HUMAN CREATIVITY. Art and the science of generative AI. *Science* **380**, 1110–1111 (2023).
<https://doi.org/10.1126/science.adh4451>.
42. Oppenlaender, J. The Creativity of Text-to-Image Generation. in *Proceedings of the 25th International Academic Mindtrek Conference* 192–202 (2022).
<https://doi.org/10.1145/3569219.3569352>.
43. Li, Z. (Lionel), Fang, X. & Sheng, O. R. L. A Survey of Link Recommendation for Social Networks: Methods, Theoretical Foundations, and Future Research Directions. *ACM Trans. Manag. Inf. Syst.* **9**, 1–26 (2018). <https://doi.org/10.1145/3131782>.
44. Lops, P., de Gemmis, M. & Semeraro, G. Content-based Recommender Systems: State of the Art and Trends. in *Recommender Systems Handbook* (eds. Ricci, F., Rokach, L., Shapira, B. & Kantor, P. B.) 73–105 (Springer US, 2011).
https://doi.org/10.1007/978-0-387-85820-3_3.
45. Su, X. & Khoshgoftaar, T. M. A Survey of Collaborative Filtering Techniques. *Adv. Artif. Intell.* **2009**, 1–19 (2009). <https://doi.org/10.1155/2009/421425>.
46. Anderson, A., Maystre, L., Anderson, I., Mehrotra, R. & Lalmas, M. Algorithmic Effects on the Diversity of Consumption on Spotify. in *Proceedings of The Web Conference 2020* 2155–2165 (Association for Computing Machinery, 2020).
<https://doi.org/10.1145/3366423.3380281>.
47. Krumme, C., Cebrian, M., Pickard, G. & Pentland, S. Quantifying Social Influence in an Online Cultural Market. *PLOS ONE* **7**, e33785 (2012).
<https://doi.org/10.1371/journal.pone.0033785>.
48. Salganik, M. J., Dodds, P. S. & Watts, D. J. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* **311**, 854–856 (2006).
<https://doi.org/10.1126/science.1121066>.

49. Richerson, P. J. & Boyd, R. *Not by genes alone: how culture transformed human evolution*. (University of Chicago Press, 2005).
<https://doi.org/10.7208/chicago/9780226712130.001.0001>.
50. Cavalli-Sforza, L. L. & Feldman, M. W. *Cultural Transmission and Evolution: A Quantitative Approach*. (Princeton University Press, 1981).
51. Mesoudi, A. Pursuing Darwin's curious parallel: Prospects for a science of cultural evolution. *Proc. Natl. Acad. Sci.* **114**, 7853–7860 (2017).
<https://doi.org/10.1073/pnas.1620741114>.
52. Enquist, M. & Ghirlanda, S. Evolution of social learning does not explain the origin of human cumulative culture. *J. Theor. Biol.* **246**, 129–135 (2007).
<https://doi.org/10.1016/j.jtbi.2006.12.022>.
53. Acerbi, A. & Mesoudi, A. If we are all cultural Darwinians what's the fuss about? Clarifying recent disagreements in the field of cultural evolution. *Biol. Philos.* **30**, 481–503 (2015). <https://doi.org/10.1007/s10539-015-9490-2>.
54. Morin, O. Reasons to be fussy about cultural evolution. *Biol. Philos.* **31**, 447–458 (2016). <https://doi.org/10.1007/s10539-016-9516-4>.
55. Weitzman, M. L. Recombinant Growth*. *Q. J. Econ.* **113**, 331–360 (1998).
<https://doi.org/10.1162/003355398555595>.
56. Griffiths, T. L. Understanding Human Intelligence through Human Limitations. *Trends Cogn. Sci.* **24**, 873–883 (2020). <https://doi.org/10.1016/j.tics.2020.09.001>.
57. Boyd, R. & Richerson, P. J. *Culture and the evolutionary process*. (University of Chicago Press, 1985).
58. Mesoudi, A. *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. (University of Chicago Press, 2011).
<https://doi.org/10.7208/chicago/9780226520452.001.0001>.
59. Leibo, J. Z., Hughes, E., Lanctot, M. & Graepel, T. Autocurricula and the Emergence of Innovation from Social Interaction: A Manifesto for Multi-Agent Intelligence Research.

- Preprint at <https://doi.org/10.48550/arXiv.1903.00742> (2019).
60. Aveni, A. F. *Skywatchers: A Revised and Updated Version of Skywatchers of Ancient Mexico*. (University of Texas Press, 2001). <https://doi.org/10.2307/972243>.
 61. Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Netw.* **4**, 251–257 (1991). [https://doi.org/10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T).
 62. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015). <https://doi.org/10.1038/nature14539>.
 63. Zenil, H. *et al.* The Future of Fundamental Science Led by Generative Closed-Loop Artificial Intelligence. Preprint at <http://arxiv.org/abs/2307.07522> (2023).
 64. Senior, A. W. *et al.* Improved protein structure prediction using potentials from deep learning. *Nature* **577**, 706–710 (2020). <https://doi.org/10.1038/s41586-019-1923-7>.
 65. Cooper, S. *et al.* Predicting protein structures with a multiplayer online game. *Nature* **466**, 756–760 (2010). <https://doi.org/10.1038/nature09304>.
 66. Bommasani, R. *et al.* On the Opportunities and Risks of Foundation Models. Preprint at <https://doi.org/10.48550/arXiv.2108.07258> (2022).
 67. Hoffmann, J. *et al.* Training Compute-Optimal Large Language Models. Preprint at <https://doi.org/10.48550/arXiv.2203.15556> (2022).
 68. Bender, E. M., Gebru, T., McMillan-Major, A. & Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* 610–623 (Association for Computing Machinery, 2021). <https://doi.org/10.1145/3442188.3445922>.
 69. Brown, T. *et al.* Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **33**, 1877–1901 (2020). <https://doi.org/10.48550/arXiv.2005.14165>.
 70. Mitchell, M. & Krakauer, D. C. The debate over understanding in AI's large language models. *Proc. Natl. Acad. Sci.* **120**, e2215907120 (2023). <https://doi.org/10.1073/pnas.2215907120>.
 71. Charbonneau, M. Modularity and Recombination in Technological Evolution. *Philos.*

- Technol.* **29**, 373–392 (2016). <https://doi.org/10.1007/s13347-016-0228-0>.
72. Henrich, J. Demography and Cultural Evolution: How Adaptive Cultural Processes Can Produce Maladaptive Losses—The Tasmanian Case. *Am. Antiq.* **69**, 197–214 (2004). <https://doi.org/10.2307/4128416>.
73. Henrich, J. & Muthukrishna, M. What Makes Us Smart? *Top. Cogn. Sci.* (2023) <https://doi.org/10.1111/tops.12656>.
74. Youn, H., Strumsky, D., Bettencourt, L. M. A. & Lobo, J. Invention as a combinatorial process: evidence from US patents. *J. R. Soc. Interface* **12**, 20150272 (2015). <https://doi.org/10.1098/rsif.2015.0272>.
75. Sourati, J. & Evans, J. A. Accelerating science with human-aware artificial intelligence. *Nat. Hum. Behav.* 1–15 (2023) <https://doi.org/10.1038/s41562-023-01648-z>.
76. Tinitis, P. & Sobchuk, O. Open-ended cumulative cultural evolution of Hollywood film crews. *Evol. Hum. Sci.* **2**, e26 (2020). <https://doi.org/10.1017/ehs.2020.21>.
77. Grizou, J., Points, L. J., Sharma, A. & Cronin, L. A curious formulation robot enables the discovery of a novel protocell behavior. *Sci. Adv.* **6**, eaay4237 (2020). <https://doi.org/10.1126/sciadv.aay4237>.
78. Kramer, S., Cerrato, M., Džeroski, S. & King, R. Automated Scientific Discovery: From Equation Discovery to Autonomous Discovery Systems. Preprint at <http://arxiv.org/abs/2305.02251> (2023).
79. Lucas, A. J. *et al.* The value of teaching increases with tool complexity in cumulative cultural evolution. *Proc. R. Soc. B Biol. Sci.* **287**, 20201885 (2020). <https://doi.org/10.1098/rspb.2020.1885>.
80. Borsa, D., Piot, B., Munos, R. & Pietquin, O. Observational Learning by Reinforcement Learning. Preprint at <https://doi.org/10.48550/arXiv.1706.06617> (2017).
81. Kohnke, L., Moorhouse, B. L. & Zou, D. ChatGPT for Language Teaching and Learning. *RELC J.* 00336882231162868 (2023). <https://doi.org/10.1177/00336882231162868>.

82. Haller, E. & Rebedea, T. Designing a Chat-bot that Simulates an Historical Figure. in *2013 19th International Conference on Control Systems and Computer Science* 582–589 (2013). <https://doi.org/10.1109/CSCS.2013.85>.
83. Zhang, S., Frey, B. & Bansal, M. How can NLP Help Revitalize Endangered Languages? A Case Study and Roadmap for the Cherokee Language. Preprint at <http://arxiv.org/abs/2204.11909> (2022).
84. Ijaz, K., Bogdanovych, A. & Trescak, T. Virtual worlds vs books and videos in history education. *Interact. Learn. Environ.* **25**, 904–929 (2017). <https://doi.org/10.1080/10494820.2016.1225099>.
85. Buolamwini, J. & Gebru, T. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* 77–91 (PMLR, 2018).
86. Caliskan, A., Bryson, J. J. & Narayanan, A. Semantics derived automatically from language corpora contain human-like biases. *Science* **356**, 183–186 (2017). <https://doi.org/10.1126/science.aal4230>.
87. O’Neil, C. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. (Crown, 2016). <https://doi.org/10.5860/crl.78.3.403>.
88. Prates, M. O. R., Avelar, P. H. C. & Lamb, L. Assessing Gender Bias in Machine Translation -- A Case Study with Google Translate. Preprint at <https://doi.org/10.48550/arXiv.1809.02208> (2019).
89. Acerbi, A. & Stubbersfield, J. Large language models show human-like content biases in transmission chain experiments. Preprint at <https://doi.org/10.31219/osf.io/8zg4d> (2023).
90. Vig, J. *et al.* Investigating Gender Bias in Language Models Using Causal Mediation Analysis. in *Advances in Neural Information Processing Systems* vol. 33 12388–12401 (Curran Associates, Inc., 2020).
91. Pessach, D. & Shmueli, E. A Review on Fairness in Machine Learning. *ACM Comput.*

- Surv.* **55**, 51:1-51:44 (2022). <https://doi.org/10.1145/3494672>.
92. Argyle, L. P. *et al.* Out of One, Many: Using Language Models to Simulate Human Samples. in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* 819–862 (2022).
<https://doi.org/10.18653/v1/2022.acl-long.60>.
93. Hendy, A. *et al.* How Good Are GPT Models at Machine Translation? A Comprehensive Evaluation. Preprint at <http://arxiv.org/abs/2302.09210> (2023).
94. Bartlett, F. C. *Remembering: A study in experimental and social psychology*. xix, 317 (Cambridge University Press, 1932).
95. Kashima, Y. Maintaining Cultural Stereotypes in the Serial Reproduction of Narratives. *Pers. Soc. Psychol. Bull.* **26**, 594–604 (2000).
<https://doi.org/10.1177/0146167200267007>.
96. Griffiths, T. L., Christian, B. R. & Kalish, M. L. Using Category Structures to Test Iterated Learning as a Method for Identifying Inductive Biases. *Cogn. Sci.* **32**, 68–107 (2008). <https://doi.org/10.1080/03640210701801974>.
97. Lieder, F. & Griffiths, T. L. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* **43**, e1 (2020).
98. Simon, H. A. Bounded Rationality. in *Utility and Probability* (eds. Eatwell, J., Milgate, M. & Newman, P.) 15–18 (Palgrave Macmillan UK, 1990).
https://doi.org/10.1007/978-1-349-20568-4_5.
99. Todd, P. M. & Gigerenzer, G. Environments That Make Us Smart: Ecological Rationality. *Curr. Dir. Psychol. Sci.* **16**, 167–171 (2007).
<https://doi.org/10.1111/j.1467-8721.2007.00497.x>.
100. Tversky, A. & Kahneman, D. Judgment under Uncertainty: Heuristics and Biases. *Science* **185**, 1124–1131 (1974). <https://doi.org/10.1126/science.185.4157.1124>.
101. Gershman, S. J., Horvitz, E. J. & Tenenbaum, J. B. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* **349**,

- 273–278 (2015). <https://doi.org/10.1126/science.aac6076>.
102. Malle, B. F., Scheutz, M., Arnold, T., Voiklis, J. & Cusimano, C. Sacrifice One For the Good of Many? People Apply Different Moral Norms to Human and Robot Agents. in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* 117–124 (Association for Computing Machinery, 2015).
<https://doi.org/10.1145/2696454.2696458>.
103. Griffiths, T. L., Kalish, M. L. & Lewandowsky, S. Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philos. Trans. R. Soc. B Biol. Sci.* **363**, 3503–3514 (2008). <https://doi.org/10.1098/rstb.2008.0146>.
104. Kirby, S., Dowman, M. & Griffiths, T. L. Innateness and culture in the evolution of language. *PNAS Proc. Natl. Acad. Sci. U. S. Am.* **104**, 5241–5245 (2007).
<https://doi.org/10.1073/pnas.0608222104>.
105. Thompson, B. & Griffiths, T. L. Human biases limit cumulative innovation. *Proc. R. Soc. B Biol. Sci.* **288**, 20202752 (2021). <https://doi.org/10.1098/rspb.2020.2752>.
106. Brinkmann, L. *et al.* Hybrid social learning in human-algorithm cultural transmission. *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.* **380**, 20200426 (2022).
<https://doi.org/10.1098/rsta.2020.0426>.
107. Tamariz, M. & Kirby, S. Culture: copying, compression, and conventionality. *Cogn. Sci.* **39**, 171–183 (2015). <https://doi.org/10.1111/cogs.12144>.
108. Chater, N. & Vitányi, P. Simplicity: a unifying principle in cognitive science? *Trends Cogn. Sci.* **7**, 19–22 (2003). [https://doi.org/10.1016/s1364-6613\(02\)00005-0](https://doi.org/10.1016/s1364-6613(02)00005-0).
109. Kirby, S., Tamariz, M., Cornish, H. & Smith, K. Compression and communication in the cultural evolution of linguistic structure. *Cognition* **141**, 87–102 (2015).
<https://doi.org/10.1016/j.cognition.2015.03.016>.
110. Anderson, C. The End of Theory: The Data Deluge Makes the Scientific Method Obsolete. *Wired* (2018).
111. Spinney, L. Are we witnessing the dawn of post-theory science? *The Guardian* (2022).

112. Liu, Z., Madhavan, V. & Tegmark, M. AI Poincaré 2.0: Machine Learning Conservation Laws from Differential Equations. *Phys. Rev. E* **106**, 045307 (2022).
<https://doi.org/10.1103/PhysRevE.106.045307>.
113. Kendal, R. L. *et al.* Social Learning Strategies: Bridge-Building between Fields. *Trends Cogn. Sci.* **22**, 651–665 (2018). <https://doi.org/10.1016/j.tics.2018.04.003>.
114. Henrich, J. & McElreath, R. The evolution of cultural evolution. *Evol. Anthropol. Issues News Rev. Issues News Rev.* **12**, 123–135 (2003). <https://doi.org/10.1002/evan.10110>.
115. Mesoudi, A., Whiten, A. & Dunbar, R. A bias for social information in human cultural transmission. *Br. J. Psychol.* **97**, 405–423 (2006).
<https://doi.org/10.1348/000712605X85871>.
116. Sharma, D. K. & Sharma, A. A comparative analysis of web page ranking algorithms. *Int. J. Comput. Sci. Eng.* **2**, 2670–2676 (2010).
117. Duhan, N., Sharma, A. K. & Bhatia, K. K. Page Ranking Algorithms: A Survey. in *2009 IEEE International Advance Computing Conference* 1530–1537 (2009).
<https://doi.org/10.1109/IADCC.2009.4809246>.
118. Koren, Y., Rendle, S. & Bell, R. Advances in collaborative filtering. *Recomm. Syst. Handb.* 91–142 (2022). https://doi.org/10.1007/978-1-0716-2197-4_3.
119. Banihashemi, S. & Abhari, A. Effects of Different Recommendation Algorithms on Structure of Social Networks. in *2021 International Conference on Computational Science and Computational Intelligence (CSCI)* 1395–1400 (2021).
<https://doi.org/10.1109/CSCI54926.2021.00279>.
120. Ferrara, A., Espín-Noboa, L., Karimi, F. & Wagner, C. Link recommendations: Their impact on network structure and minorities. in *14th ACM Web Science Conference 2022* 228–238 (2022). <https://doi.org/10.1145/3501247.3531583>.
121. Su, J., Sharma, A. & Goel, S. The Effect of Recommendations on Network Structure. in *Proceedings of the 25th International Conference on World Wide Web* 1157–1167 (International World Wide Web Conferences Steering Committee, 2016).

<https://doi.org/10.1145/2872427.2883040>.

122. Lazer, D. & Friedman, A. The Network Structure of Exploration and Exploitation. *Adm. Sci. Q.* **52**, 667–694 (2007).
123. Mason, W. & Watts, D. J. Collaborative learning in networks. *Proc. Natl. Acad. Sci.* **109**, 764–769 (2012).
124. Woolley, A. W., Aggarwal, I. & Malone, T. W. Collective Intelligence and Group Performance. *Curr. Dir. Psychol. Sci.* **24**, 420–424 (2015).
125. Derex, M. & Boyd, R. Partial connectivity increases cultural accumulation within groups. *Proc. Natl. Acad. Sci.* **113**, 2982–2987 (2016).
126. Kant, V., Jhalani, T. & Dwivedi, P. Enhanced multi-criteria recommender system based on fuzzy Bayesian approach. *Multimed. Tools Appl.* **77**, 12935–12953 (2018).
127. Bollen, D., Knijnenburg, B. P., Willemsen, M. C. & Graus, M. Understanding choice overload in recommender systems. in *Proceedings of the fourth ACM conference on Recommender systems* 63–70 (2010).
128. Tkalcic, M., Kosir, A. & Tasic, J. Affective recommender systems: the role of emotions in recommender systems. in *The RecSys 2011 Workshops-Decisions@ RecSys 2011 and UCERSTI-2: Human Decision Making in Recommender Systems; User-Centric Evaluation of Recommender Systems and Their Interfaces-2* vol. 811 9–13 (CEUR-WS.org, 2011).
129. Gonzalez, G., de la Rosa, J. L., Montaner, M. & Delfin, S. Embedding Emotional Context in Recommender Systems. in *2007 IEEE 23rd International Conference on Data Engineering Workshop* 845–852 (2007). doi:10.1109/ICDEW.2007.4401075.
130. Osman, N. A., Mohd Noah, S. A., Darwich, M. & Mohd, M. Integrating contextual sentiment analysis in collaborative recommender systems. *PLoS ONE* **16**, e0248695 (2021).
131. Zheng, Y., Mobasher, B. & Burke, R. D. The Role of Emotions in Context-aware Recommendation. *Decis. RecSys* **2013**, 21–28 (2013).

132. Zhang, X., Ferreira, P., Godinho De Matos, M. & Belo, R. Welfare Properties of Profit Maximizing Recommender Systems: Theory and Results from a Randomized Experiment. *MIS Q.* **45**, 1–43 (2021).
133. Levy, R. Social Media, News Consumption, and Polarization: Evidence from a Field Experiment. *Am. Econ. Rev.* **111**, 831–870 (2021).
134. Brady, W. J., Gantman, A. P. & Van Bavel, J. J. Attentional capture helps explain why moral and emotional content go viral. *J. Exp. Psychol. Gen.* **149**, 746–756 (2020).
135. Brady, W. J., Jackson, J. C., Lindström, B. & Crockett, M. J. Algorithm-mediated social learning in online social networks. *Trends Cogn. Sci.* **0**, (2023).
136. Acerbi, A. Cognitive attraction and online misinformation. *Palgrave Commun.* **5**, 1–7 (2019).
137. Brady, W. J. *et al.* Overperception of moral outrage in online social networks inflates beliefs about intergroup hostility. *Nat. Hum. Behav.* 1–11 (2023)
doi:10.1038/s41562-023-01582-0.
138. Brady, W. J. & Crockett, M. J. Norm Psychology in the Digital Age: How Social Media Shapes the Cultural Evolution of Normativity. *Perspect. Psychol. Sci.* 17456916231187395 (2023) doi:10.1177/17456916231187395.
139. Milli, S., Carroll, M., Pandey, S., Wang, Y. & Dragan, A. D. Twitter’s Algorithm: Amplifying Anger, Animosity, and Affective Polarization. Preprint at <http://arxiv.org/abs/2305.16941> (2023).
140. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W. & Starnini, M. The echo chamber effect on social media. *Proc. Natl. Acad. Sci.* **118**, e2023301118 (2021).
141. Pariser, E. *The Filter Bubble: What the Internet Is Hiding from You.* (The Penguin Press HC, 2011).
142. Sunstein, C. R. *Republic.com 2.0.* (Princeton University Press, 2007).
143. Jiang, R., Chiappa, S., Lattimore, T., György, A. & Kohli, P. Degenerate Feedback

- Loops in Recommender Systems. in *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* 383–390 (ACM, 2019). doi:10.1145/3306618.3314288.
144. Pagan, N. *et al.* A Classification of Feedback Loops and Their Relation to Biases in Automated Decision-Making Systems. Preprint at <http://arxiv.org/abs/2305.06055> (2023).
145. Stray, J. *et al.* Building Human Values into Recommender Systems: An Interdisciplinary Synthesis. Preprint at <https://doi.org/10.48550/arXiv.2207.10192> (2022).
146. Kleinberg, J., Mullainathan, S. & Raghavan, M. The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization. Preprint at <https://doi.org/10.48550/arXiv.2202.11776> (2022).
147. Ovadya, A. & Thorburn, L. Bridging Systems: Open Problems for Countering Destructive Divisiveness across Ranking, Recommenders, and Governance. Preprint at <http://arxiv.org/abs/2301.09976> (2023).
148. Yao, B., Jiang, M., Yang, D. & Hu, J. Empowering LLM-based Machine Translation with Cultural Awareness. Preprint at <http://arxiv.org/abs/2305.14328> (2023).
149. Garimella, K., De Francisci Morales, G., Gionis, A. & Mathioudakis, M. Reducing Controversy by Connecting Opposing Views. in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining* 81–90 (Association for Computing Machinery, 2017). doi:10.1145/3018661.3018703.
150. Santos, F. P., Lelkes, Y. & Levin, S. A. Link recommendation algorithms and dynamics of polarization in online social networks. *Proc. Natl. Acad. Sci.* **118**, e2102141118 (2021).
151. Möller, J., Trilling, D., Helberger, N. & Van Es, B. Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. *Inf. Commun. Soc.* **21**, 959–977 (2018).
152. Bakker, M. A. *et al.* Fine-tuning language models to find agreement among humans with diverse preferences. Preprint at <https://doi.org/10.48550/arXiv.2211.15006> (2022).

153. Christiano, P. *et al.* Deep reinforcement learning from human preferences. Preprint at <https://doi.org/10.48550/arXiv.1706.03741> (2023).
154. Ouyang, L. *et al.* Training language models to follow instructions with human feedback. Preprint at <https://doi.org/10.48550/arXiv.2203.02155> (2022).
155. Perez, E. *et al.* Discovering Language Model Behaviors with Model-Written Evaluations. Preprint at <http://arxiv.org/abs/2212.09251> (2022).
156. Claidière, N., Scott-Phillips, T. C. & Sperber, D. How Darwinian is cultural evolution? *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 20130368 (2014).
157. Blancke, S., Van Breusegem, F., De Jaeger, G., Braeckman, J. & Van Montagu, M. Fatal attraction: the intuitive appeal of GMO opposition. *Trends Plant Sci.* **20**, 414–418 (2015).
158. Miton, H. & Mercier, H. Cognitive Obstacles to Pro-Vaccination Beliefs. *Trends Cogn. Sci.* **19**, 633–636 (2015).
159. Poulsen, V. & DeDeo, S. Cognitive Attractors and the Cultural Evolution of Religion. Preprint at <https://doi.org/10.31234/osf.io/daxyu> (2023).
160. Kirchenbauer, J. *et al.* A Watermark for Large Language Models. Preprint at <http://arxiv.org/abs/2301.10226> (2023).
161. Shumailov, I. *et al.* The Curse of Recursion: Training on Generated Data Makes Models Forget. Preprint at <http://arxiv.org/abs/2305.17493> (2023).
162. Veselovsky, V., Ribeiro, M. H. & West, R. Artificial Artificial Artificial Intelligence: Crowd Workers Widely Use Large Language Models for Text Production Tasks. Preprint at <https://doi.org/10.48550/arXiv.2306.07899> (2023).
163. Japkowicz, N. & Stephen, S. The class imbalance problem: A systematic study. *Intell. Data Anal.* **6**, 429–449 (2002).
164. Kalish, M. L., Griffiths, T. L. & Lewandowsky, S. Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychon. Bull. Rev.* **14**, 288–294 (2007).

165. Axelrod, R. The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *J. Confl. Resolut.* **41**, 203–226 (1997).
166. Touvron, H. *et al.* LLaMA: Open and Efficient Foundation Language Models. Preprint at <https://doi.org/10.48550/arXiv.2302.13971> (2023).
167. West, S. M., Whittaker, M. & Crawford, K. Discriminating systems: gender, race and power in AI. *AI Inst.*
168. Autor, D. H. Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *J. Econ. Perspect.* **29**, 3–30 (2015).
169. Ayers, J. W. *et al.* Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Intern. Med.* (2023).
170. Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C. & Althoff, T. Human–AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support. *Nat. Mach. Intell.* **5**, 46–57 (2023).
171. Perry, A. AI will never convey the essence of human empathy. *Nat. Hum. Behav.* 1–2 (2023) doi:10.1038/s41562-023-01675-w.
172. Weisz, E. & Zaki, J. Motivated empathy: a social neuroscience perspective. *Curr. Opin. Psychol.* **24**, 67–71 (2018).
173. Carroll, M., Hadfield-Menell, D., Russell, S. & Dragan, A. Estimating and Penalizing Preference Shift in Recommender Systems. in *Proceedings of the 15th ACM Conference on Recommender Systems* 661–667 (Association for Computing Machinery, 2021). doi:10.1145/3460231.3478849.
174. Bakshy, E., Messing, S. & Adamic, L. A. Exposure to ideologically diverse news and opinion on Facebook. *Science* **348**, 1130–1132 (2015).
175. Robertson, R. E. *et al.* Users choose to engage with more partisan news than they are exposed to on Google Search. *Nature* **618**, 342–348 (2023).
176. Art made by artificial intelligence is developing a style of its own. *The Economist*.
177. Obradovich, N. *et al.* Expanding the measurement of culture with a sample of two

- billion humans. *J. R. Soc. Interface* **19**, 20220085 (2022).
178. Garg, N., Schiebinger, L., Jurafsky, D. & Zou, J. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proc. Natl. Acad. Sci.* **115**, E3635–E3644 (2018).
179. Karjus, A., Solà, M. C., Ohm, T., Ahnert, S. E. & Schich, M. Compression ensembles quantify aesthetic complexity and the evolution of visual art. *EPJ Data Sci.* **12**, 21 (2023).
180. Santy, S., Liang, J. T., Bras, R. L., Reinecke, K. & Sap, M. NLPositionality: Characterizing Design Biases of Datasets and Models. Preprint at <http://arxiv.org/abs/2306.01943> (2023).
181. Awad, E. *et al.* The Moral Machine experiment. *Nature* **563**, 59–64 (2018).
182. Brandt, F., Conitzer, V. & Endriss, U. Computational Social Choice. in *Multiagent Systems* (The MIT Press, 2013).
183. Koster, R. *et al.* Human-centred mechanism design with Democratic AI. *Nat. Hum. Behav.* **6**, 1398–1407 (2022).
184. Small, C. T. *et al.* Opportunities and Risks of LLMs for Scalable Deliberation with Polis. Preprint at <http://arxiv.org/abs/2306.11932> (2023).
185. Rahwan, I. Society-in-the-loop: programming the algorithmic social contract. *Ethics Inf. Technol.* **20**, 5–14 (2018).
186. Jernite, Y. *et al.* Data Governance in the Age of Large-Scale Data-Driven Language Technology. in *2022 ACM Conference on Fairness, Accountability, and Transparency* 2206–2222 (2022). doi:10.1145/3531146.3534637.
187. Laurençon, H. *et al.* The BigScience ROOTS Corpus: A 1.6TB Composite Multilingual Dataset. Preprint at <https://doi.org/10.48550/arXiv.2303.03915> (2023).
188. Ziegler, D. M. *et al.* Fine-Tuning Language Models from Human Preferences. Preprint at <https://doi.org/10.48550/arXiv.1909.08593> (2020).
189. Bai, Y. *et al.* Constitutional AI: Harmlessness from AI Feedback. Preprint at <http://arxiv.org/abs/2212.08073> (2022).

190. Bergstrom, C. T. & Lachmann, M. The Red King effect: When the slowest runner wins the coevolutionary race. *Proc. Natl. Acad. Sci.* **100**, 593–598 (2003).
191. Bostrom, N. *Superintelligence: Paths, Dangers, Strategies*. (Oxford University Press, 2014).
192. Wilson, D. S. *et al.* Multilevel cultural evolution: From new theory to practical applications. *Proc. Natl. Acad. Sci. U. S. A.* **120**, e2218222120 (2023).
193. DALL·E: Creating Images from Text, <https://openai.com/research/dall-e> (OpenAI, 2021).
194. Lisi, E., Malekzadeh, M., Haddadi H., Lau, F.D.-H. & Flaxman, S. Modelling and forecasting art movements with CGANs *R. Soc. Open Science* **7**, 191569 (2020).