

Human-AI coevolution

Dino Pedreschi^{a,*}, Luca Pappalardo^{b,c,*}, Emanuele Ferragina^{d,*},
 Ricardo Baeza-Yates^e, Albert-László Barabási^e, Frank Dignum^f, Virginia Dignum^f,
 Tina Eliassi-Rad^e, Fosca Giannotti^c, János Kertész^g, Alistair Knott^h,
 Yannis Ioannidisⁱ, Paul Lukowicz^j, Andrea Passarella^b, Alex Sandy Pentland^k,
 John Shawe-Taylor^l, Alessandro Vespignani^e

^a University of Pisa, Italy

^b Consiglio Nazionale delle Ricerche (CNR), Pisa, Italy

^c Scuola Normale Superiore, Pisa, Italy

^d Sciences Po, Paris, France

^e Northeastern University, Boston, USA

^f Umeå University, Sweden

^g Central European University (CEU), Vienna, Austria

^h Victoria University Wellington, New Zealand

ⁱ University of Athens, Greece

^j DFKI and University of Kaiserslautern, Germany

^k Massachusetts Institute of Technology (MIT), Boston, USA

^l University College London, UK

ARTICLE INFO

Keywords:

Artificial intelligence
 Complex systems
 Computational social science
 Human-AI coevolution

ABSTRACT

Human-AI coevolution, defined as a process in which humans and AI algorithms continuously influence each other, increasingly characterises our society, but is understudied in artificial intelligence and complexity science literature. Recommender systems and assistants play a prominent role in human-AI coevolution, as they permeate many facets of daily life and influence human choices through online platforms. The interaction between users and AI results in a potentially endless feedback loop, wherein users' choices generate data to train AI models, which, in turn, shape subsequent user preferences. This human-AI feedback loop has peculiar characteristics compared to traditional human-machine interaction and gives rise to complex and often "unintended" systemic outcomes. This paper introduces human-AI coevolution as the cornerstone for a new field of study at the intersection between AI and complexity science focused on the theoretical, empirical, and mathematical investigation of the human-AI feedback loop. In doing so, we: (i) outline the pros and cons of existing methodologies and highlight shortcomings and potential ways for capturing feedback loop mechanisms; (ii) propose a reflection at the intersection between complexity science, AI and society; (iii) provide real-world examples for different human-AI ecosystems; and (iv) illustrate challenges to the creation of such a field of study, conceptualising them at increasing levels of abstraction, i.e., scientific, legal and socio-political.

* Corresponding authors.

E-mail addresses: dino.pedreschi@unipi.it (D. Pedreschi), luca.pappalardo@isti.cnr.it (L. Pappalardo), emanuele.ferragina@sciencespo.fr (E. Ferragina).

<https://doi.org/10.1016/j.artint.2024.104244>

Received 7 July 2023; Received in revised form 24 September 2024; Accepted 2 November 2024

Available online 13 November 2024

0004-3702/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

It is change, continuing change, inevitable change, that is the dominant factor in society today. No sensible decision can be made any longer without taking into account not only the world as it is, but the world as it will be

[Isaac Asimov, *Asimov on Science Fiction*, 1981 [1]]

The history of humankind is a history of coevolution: between humans and other species; between humans and industrial machines; between humans and digital technologies; and, today, between humans and Artificial Intelligence (AI) [2,3]. Human-AI coevolution is a perpetual, iterative process wherein both humans and learning algorithms evolve in tandem, each influencing the evolution of the other over time. This generates complex effects on human-AI ecosystems and, therefore, on society [3–6].

In this context, recommendation systems and assistants (in short, *recommenders*) – AI-based algorithms that suggest items or content based on users’ preferences or specific requests [7–9] – play a prominent role. Recommenders on online platforms cater for a large audience of users, and therefore can be considered as “mainstream recommenders”. Mainstream recommenders¹ are of collective importance because they feed “very large online platforms” (VLOPs, as defined by the Digital Services Act) that are widely considered as highly impactful for society. Mainstream recommenders instantly shape the decisions and interactions of large segments of the population at the same time. Personalised suggestions on social media guide our content consumption and social connections, online retail recommenders propose products for consumption, navigation services suggest routes to reach our destinations, and generative AI creates content in response to users’ wishes. Other types of AI-powered assistants, such as medical diagnosis tools, self-driving cars, and educational tutors, are mainly designed for one-to-one interaction with individual users and do not have yet the same widespread influence as those on public online platforms. However, these systems can scale up and exhibit feedback loop effects. Our definition of recommender is general enough to encompass these prospective ecosystems. In this paper, we focus for the sake of concreteness on today’s VLOPs’ recommenders because they constitute a vantage point to analyse the coevolution between humans and AI machines.

Since recommenders are based on AI, and machine learning in particular, their interactions with users always give rise to a *feedback loop* [10–13], which is the core of human-AI coevolution. One could describe the feedback loop as a process: users’ choices determine the datasets on which recommenders are trained; the trained recommenders then exert an influence on users’ subsequent choices, which in turn affect the next round of training, initiating a potentially never-ending cycle (see Fig. 1). The human-AI feedback loop may also lead to “unintended” social consequences [14–17]. Personalised recommendations on social media help users deal with information overload, but may artificially amplify echo chambers, filter bubbles, and processes of radicalisation [14,15,18–20]. Profiling and targeted advertising may increase inequality and monopolies, perpetuating and accruing biases, discriminations, and the “tragedy of the commons” [21–24]. Navigation services suggest directions that make sense from an individual perspective, but may create chaos if too many drivers are sent to the same roads [17,25–33]. Chatbots based on generative AI may deteriorate their quality as increasing synthetic data is used for retraining and fine-tuning [127–135].

While coevolution between humans and technology is not new, the human-AI feedback loop imbues it with forms that are unprecedented. Throughout history, technology and society have constantly coevolved, e.g., as testified by the advent of the press, the radio and the TV [2,3]. Moreover, recommendations have always influenced human choices, whether it was the role of a librarian picking up a book, a friend’s advice on a music album to listen to, or the influence of TV advertisements on viewers. However, at least five fundamental aspects have been *simultaneously* magnified by AI-based recommenders: pervasiveness, persuasiveness, traceability, speed, and complexity.

Recommenders are pervasive in all online platforms, from social media to online retail and mapping services. This integration, fuelled by sophisticated algorithms, abundant data, and widespread user adoption, has established recommenders as customary elements of online interactions. The large availability of data portraying individual choices empowers recommenders to deliver highly personalised suggestions. This increases accuracy in capturing users’ preferences, making recommenders highly persuasive. Unlike previous technology, human-AI ecosystems leave an indelible trace of recommenders’ suggestions and related human choices (what is commonly called big data). As a consequence, AI-based recommenders possess a comprehensive outlook of individual choices and an unprecedented capacity to shape human behaviour at scale. Within this context, human-machine coevolution is faster than ever before, given that AI can be re-trained with little or no human oversight and provide suggestions at an unparalleled speed. Human-AI ecosystems also foster an extraordinary volume of interactions between the huge spaces of users and products, escalating system complexity.

The design of recommenders and the analysis of their impact are typically studied by two disciplines that gained prominence over the last two decades: artificial intelligence and complexity science. However, working in isolation, these disciplines cannot address the challenges of understanding human-AI coevolution.

AI has achieved human-like performance in many challenging tasks [34–39] and is becoming increasingly explainable and human-centric [40–46]. However, it is still embedded in methodological individualism [47], where machines are studied as solitary agents and not from a coevolution perspective. We know little about the impact of the feedback loop on human-AI ecosystems; therefore, we need to disentangle whether recommenders amplify undesired collective outcomes or mitigate them.

Complexity science has shown that networks of social interactions are heterogeneous, resulting in deeply connected hubs and modular structures [48–54]. This structural heterogeneity interacts with social effects like inequalities and segregation, impacting on network processes such as the spreading of epidemics, information and opinions [14,55,56]; the success of products, ideas, and people

¹ From now on, we use the term recommenders to signify mainstream recommenders.

THE HUMAN-AI FEEDBACK LOOP

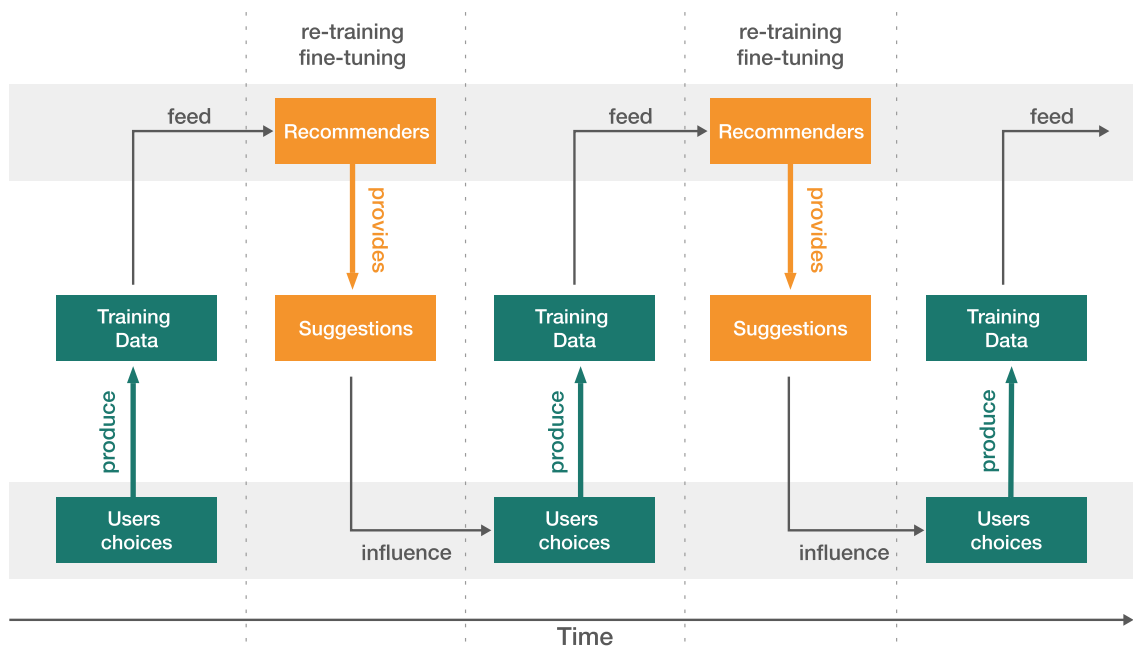


Fig. 1. Users' choices on online platforms generate data used to train recommenders. These recommenders then offer suggestions to users, influencing their choices, which in turn generate more data for re-training recommenders. This iterative process creates a potentially endless feedback loop.

[57–59]; and urban dynamics [60–69]. The human-AI feedback loop affects these network processes differently than before the advent of AI. A coevolutionary approach may help unveil the laws governing the complex interplay between humans and recommenders [70–72], going beyond existing approaches [73–75]. We know little about the parameters that govern the network dynamics of human-AI coevolution and, therefore, how to make predictions for such complex systems.

This paper introduces human-AI coevolution as the cornerstone for a new field of study at the intersection between AI and complexity science, and focuses on the theoretical, empirical, and mathematical investigation of the human-AI feedback loop. Following in the footsteps of the literary imagination of Isaac Asimov, “no sensible decision can be made any longer without taking into account not only the world as it is, but the world as it will be” [1]. A detailed understanding of human-AI coevolution should be achieved by developing new methodologies and analytical frameworks. The study of human-AI coevolution is key for understanding how feedback loop mechanisms might impact societal dynamics, as the interaction between humans and recommenders might amplify or mitigate social phenomena. This is salient because we are immersed in a political economy that mainly privileges individual utility over collective goods, and where the means of production and recommendation are concentrated in a few hands.

The study of human-AI coevolution sits at the intersection of two prominent AI debates. First, it expands the debate on “machine behaviour”, i.e., the study of behaviour exhibited by intelligent machines [76], illustrating that the feedback loop is the key process shaping human-AI coevolution. Second, it enriches current philosophical perspectives about AI, i.e., technology-centred AI, human-centred AI, and collective intelligence [77], with a new dimension that can be defined as society-centred AI. This perspective departs from technology-centred AI because it considers that the impacts of the feedback loop cannot be managed solely with additional AI technology. Moreover, it embraces seminal insights from human-centred AI and collective intelligence. On the one hand, from a human-centred AI perspective, it hypothesises that the feedback loop might curtail human well-being. On the other, from a collective intelligence perspective, it theorises how human-recommender interactions can drive coevolution towards desirable outcomes. Society-centred AI brings three additional elements to the debate (see Fig. 2): (i) the feedback loop impacts human well-being not only at an individual, but also at the societal level; (ii) managing the feedback loop requires the development of new scientific approaches; and (iii) the issues related to human-AI coevolution cannot be solved without legal and political interventions.

The remainder of the paper proceeds as follows. We review the methodologies employed in the literature to examine the interaction between humans and recommenders (Section 2). We then illustrate the outcomes of this interaction in four key ecosystems, i.e., social media, online retail, urban mapping, and content generation ecosystems (Section 3) and speculate on the social impact of human-AI coevolution (in Section 4). Finally, we outline open challenges in the study of human-AI coevolution, providing new avenues for future research (Section 5).

PHILOSOPHICAL DEBATES OF AI

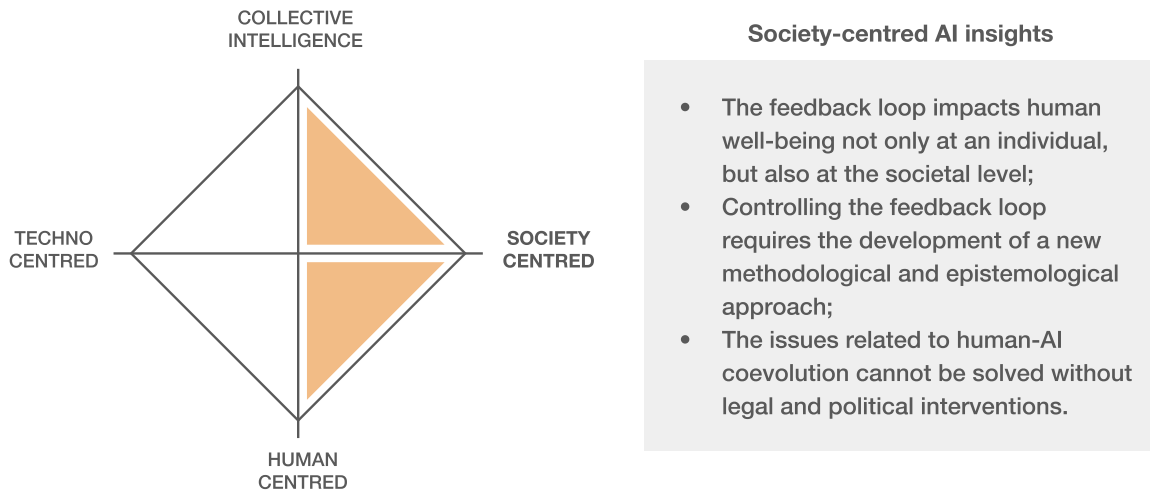


Fig. 2. We enrich current philosophical perspectives about AI, i.e., technology-centred AI, human-centred AI, and collective intelligence, with a new dimension that can be defined as society-centred AI. Society-centred AI brings three additional elements to the debate.

2. Methods for the study of human-AI coevolution

This section provides a critical perspective on existing studies about the interaction between humans and recommenders and illustrates that only a few of them investigate the role of the human-AI feedback loop. In what follows, we analyse the pros and cons of each methodology employed and highlight shortcomings and potential ways for capturing feedback loop mechanisms.

Empirical vs simulation studies

Existing studies mostly employ either an empirical or a simulation approach. Empirical studies are based on data generated by users' behaviour on real platforms, which emerge from dynamic interactions between humans and recommenders. These studies provide evidence about specific phenomena and an empirical basis for their understanding. When based on large samples or diverse datasets, empirical studies also enable researchers to make generalisations about broader populations. Nonetheless, the possibility to draw general conclusions is limited by the specific time frames and conditions under which these studies are conducted. In addition, as data are primarily owned by big tech platforms and rarely made public, these studies are hardly reproducible.

Simulation studies are based on data generated by mechanistic, AI-, or digital-twin-based models. Simulation studies provide a cost-effective alternative to empirical studies, especially when dealing with large-scale ecosystems or when data is not readily available. They can be reproduced with the same initial conditions, allowing for the verification and validation of results. Scholars can manipulate parameters to observe the effects on the human-AI ecosystem, helping the research community understand the intricate relationships between variables. However, as they are based on heavy assumptions, simulations do not necessarily reflect real-world dynamics and are, therefore, limited in unveiling unexpected or unintended outcomes. The settings of specific parameters may close the door to unforeseen results.

Controlled vs observational studies

Both empirical and simulation approaches can be either observational or controlled. Controlled studies, called experiments in social sciences [78,79], include quasi-experiments, randomised controlled trials and A/B tests. These studies split the sample into a control group and one or more experimental groups, each subjected to different recommendations [78]. The influence of recommendations is measured by comparing the outcomes obtained in the two (or more) groups. Controlled studies allow researchers to control for various factors and conditions, facilitating the isolation of the effect produced by a specific intervening variable. Their main advantage is to establish causal relationships and attribute observed effects to the recommendation. Moreover, sample randomization reduces selection biases, ensuring that participants in both groups have an equal chance of receiving the recommendation. However, controlled studies also have important shortcomings: the inclusion and exclusion criteria of the controlled settings might limit the generalisability of findings; and there is limited flexibility in adapting to changes intercurring during the experiments. Moreover, they are hard to design because they require direct access to platforms' users and recommenders [80].

Some examples of controlled studies can help clarify the characteristics of these methodologies. Huszár et al. [20] document an experiment that investigates the effect of Twitter personalised timeline on the diffusion of political content from different parties. The platform selected a group of users that were exposed to tweets in reverse-chronological order from accounts they followed, while a treatment group of users was exposed to the AI-based personalised recommender. The personalised recommendations resulted in a substantial amplification of political messages, with mainstream right-wing parties benefiting more from algorithmic personalisation than their left-wing counterparts. Cornacchia et al. [28] used a traffic simulator to estimate the impact of real-time navigation services

on the urban environment. The study found that navigation services may increase travel time and CO2 emissions when the adoption rate crosses a threshold.

Observational studies assume a single recommendation principle without any control [81]. Examples include the analysis of Facebook users' behaviour [82], Google Maps' suggestions to drivers [83], data gathered from browser loggers [84,85], platforms' APIs [82,86], bots that simulate human behaviour [87–89], and experiments that ask volunteers to behave in specific ways [90–93]. A strength of observational studies (either based on empirical data or simulations) is that, when data is big and representative enough, they allow researchers to make generalisations about a broader population, enhancing the external validity of findings and highlighting potential biases towards different population segments. A significant limitation, however, is that establishing a causal relationship is challenging; additional evidence is required to support causal claims. Moreover, findings of observational studies may suffer from selection biases, measurement errors, or the presence of confounding variables; this may compromise their accuracy and reliability.

To clarify the characteristics of observational studies, we rely on two examples. Cho et al. [93] analysed factors leading to political polarisation on YouTube, finding that a recommender can contribute to this polarisation, but mainly through the channel of a user's preference. Fleder et al. [94] investigated the influence of recommenders on users' purchases on a music streaming platform. They found that users exposed to recommendations purchase more items and are more similar to one another in the variety of what they buy.

Modelling the feedback loop

There is a handful of studies that move beyond the static investigation of the influence of recommenders on users' behaviour, analysing feedback loop mechanisms theoretically and/or empirically [11–13,95,96]. Some of these works introduce mathematical models to provide insights on recommenders' influence based on their parameters. Jiang et al. [11] theoretically investigated whether feedback loop mechanisms lead to a degeneration of users' interest: an oracle recommender (with perfect accuracy) induces a quick degeneration, while injecting randomness in users' choices and enlarging the pool of items slows this process down. Therefore, this study provides insights and potential remedies against the degeneration of the feedback loop.

There is also research that combines a theoretical and an empirical appraisal of the feedback loop. Ensign et al. [95] investigated the impact of a predictive policing recommender, based on historical crime data, on distributing police resources among urban districts. They simulated a scenario where the following feedback loop takes place: officers are deployed daily to districts with the highest predicted crime rate, crimes discovered by these officers are reported, and then data about these reported crimes are fed back into the recommender. The process continues iteratively. Given this feedback loop, the recommender repeatedly redirects police attention to districts where more crimes are reported. As more crimes are likely to be discovered in these districts due to the increased presence of officers, in the long run, the simulation leads to a distribution of crimes that is unrealistic if compared with observed historical crime data. The authors proposed a correction mechanism where the likelihood of police deployment to a district decreases as discovered crime data are incorporated into the recommender. Simple urn models were employed to model the feedback loop and the correction mechanism. One may interpret this evidence as suggesting that small changes in the recommender can make a real difference in the human-AI ecosystem.

Although these seminal studies provide intriguing insights into feedback loop mechanisms, there are important avenues for further improvement in the analysis of human-AI coevolution. At the empirical level, data employed in these studies allow drawing only an incomplete picture of the interaction between humans and recommenders. Typically, they only describe users' choices at time t_i , without considering which recommendations are provided to influence users' choices at t_{i-1} . Moreover, we have no information about how often platforms re-train recommenders on users' choices to update the knowledge of their preferences. For these reasons, existing studies employ available data only to validate theoretical approaches that model feedback loop mechanisms. To overcome this limitation, we need empirical studies based on longitudinal data describing, at each iteration of the feedback loop: (i) the recommendations provided to users; (ii) the reaction of users to these recommendations; and (iii) the process of re-training based on users' choices influenced by previous recommendations. Studies of this kind would permit looking at causation bi-directionally and not only unidirectionally.

3. Outcomes of human-AI coevolution

An outcome can be defined as the by-product of the interaction between recommenders and users within each step of the feedback loop. Outcomes can be measured at different levels, i.e., individual, item, model, and systemic. Individual outcomes refer to the influence of recommenders on users, e.g., sellers and buyers in the online retail ecosystem, and drivers and passengers in the urban mapping ecosystem. Item outcomes refer to the influence of recommenders and users' choices on the characteristics of specific objects; examples include posts on social media, products on online retail platforms, rides in urban services, and text or images on the content generation ecosystem. Model outcomes pertain to the influence of users' choices on the characteristics of the recommender. This encompasses whether the recommender alters its behaviour and the nature of its recommendations in response to users' choices. Systemic outcomes refer to the collective impact of the interaction between humans and recommenders.

Reviewing the literature, we detect several systemic outcomes, such as polarization, echo chambers, inequality, concentration, and segregation. Polarization is a sharp separation of users or items into groups based on some attributes (opinions or beliefs) [14, 19,88,93,97–110]. Echo chambers are environments in which opinions or item choices within a group are confirmed and reinforced [19,85,103,111]. Inequality indicates an uneven distribution of resources among members of a group, while concentration is a close gathering of users or items [86,112–116]. Concentration is typically identified as congestion in the urban context [17,28,31,32,68,

83,117–121]. Segregation is a situation in which groups of users are set apart from each other [122,123]. Some outcomes emerge at the individual or model level only. Examples are filter bubbles (individual level), which is a conformation of contents with a user's own beliefs [43,124–126], and model collapse (model level), i.e., a deterioration of the recommender's performance as it continues to coevolve with users [127–135]. Finally, outcomes emerging at all levels are changes in volume (i.e., a quantity measuring some users' or items' attribute) and changes in diversity, which can be defined as the variety of items and users' behaviours [136–144].

All these outcomes emerge, often unintendedly, as the result of the human-AI coevolution. A notable example is the emergence of concentration, conformism and diversity loss. In what follows, we briefly discuss other examples gathered from four largely studied human-AI ecosystems, i.e., social media, online retail, urban mapping, and content generation ecosystems. For a more comprehensive discussion of outcomes in these four ecosystems, refer to [145].

Social media. Recommenders have achieved considerable success in the realm of social media, where they are deployed to suggest new posts and users to follow. This coevolution between users and recommenders gives rise to two interconnected feedback loops. First, previous interactions between users and posts shape actual recommendations; these recommendations influence subsequent interactions between users and future posts. Second, the accounts a user follows shape actual user recommendations; these recommendations influence subsequent interactions between users and their followers. Given that social media disseminate opinions, this feedback loop can lead to various outcomes, at individual (filter bubble, homophily) and systemic levels (polarisation, fragmentation, echo chamber). While recommenders assist users in accessing content and connecting with like-minded individuals, the underlying algorithms may confine them within a filter bubble. This confinement can contribute to significant polarisation of opinions and users, fostering the potential for radicalization of ideas. For example, as previously discussed, personalised recommendations on Twitter overexpose users to certain political content [20].

Online retail. Recommenders also play a pivotal role in the success of e-commerce and streaming giants like Amazon, eBay, and Netflix. The coevolution of these recommenders with consumers may give rise to intricate feedback loop mechanisms. The recommended items (e.g., consumer goods, songs, movies) depend on previous purchases, which in turn were influenced by previous recommendations. A crucial distinction in this human-AI ecosystem lies between collaborative filtering and personalised recommenders. Collaborative filtering operates on the principle of “who-buys-this-also-buys-that”, relying on collective user behaviour; while personalised recommendations tailor suggestions to individual user taste. For example, collaborative filtering may increase sales volume and individual consumption variety while, at the same time, may decrease overall aggregate consumption diversity, amplifying the success of popular products [146,147]. On the one hand, recommenders help users better navigate the large space of product choice, reducing effort and choice overload, allowing a rapid allocation of needed goods and boosting platform revenues. On the other, at the aggregate level, they might reduce the variety of purchased products (i.e., creating filter bubbles around the user) and increase concentration, favouring certain brands and reducing competition.

Urban mapping. Navigation services recommend a route to a destination, considering changing traffic conditions and assisting users in exploring unfamiliar areas. Therefore, users with the same origin and destination receive similar recommendations. The impact of navigation services on the city is unclear: being designed to optimise individual travel times, they may also cause congestion and lead to longer travel times and higher CO2 emissions in the environment [17,25,26,28,31–33]. For example, in 2017, Google Maps, Waze and Apple Maps re-routed drivers from congested highways to the narrow and hilly streets of Leonia (a small town in New Jersey), creating such congestion that people could not get out of their driveways [17,25,27]. These issues are exacerbated by the coevolution between drivers and algorithmic updates, generating a feedback loop: travel times, which shape actual recommendations, also depend on drivers' route choices that were influenced by previous recommendations. Drivers' behaviour may change travel times in return, shaping subsequent recommendations. In this context, if too many drivers choose the same “eco-friendly” route, this route will cease to be eco-friendly.

Content Generation. Recently emerging large language models (LLMs), deployed by big tech companies (e.g., Amazon, Baidu, Google, Meta, Microsoft, OpenAI) and/or open source projects (e.g., Bloom, Cerebras-GPT, Dolly, Falcon, Mistral, Zephyr), are rapidly permeating various domains, e.g., education, politics, work. It has been observed that LLMs may potentially compress diversity, thus standardising the use of language in generated text [127–135]. Recent studies show that when LLM-generated content is used to fine-tune LLMs themselves, a “regression to the mean” can take place with a loss of linguistic diversity in the form and substance of the generated text. This process of autophagy will be increasingly widespread as content on the web, which is the fuel that trains and fine-tunes LLMs will be more likely to be generated by machines rather than humans. At an analytical level, we need to understand this coevolution to develop recommenders that balance the feedback loop's potential positive (standardisation of language) and negative impacts (compression of linguistic diversity).²

4. Social impact of human-AI coevolution

With our concept of human-AI coevolution, we do not simply propose a reflection on the interconnection between complex systems and AI, but also on their relationship with society. Some of the outcomes discussed above have a long history in social science. Notable examples are polarization (in political science [148,149]), inequality (in economics [150], in sociology [151]), and segregation (in economics [152], in urban studies [153]).

² Note that, in this ecosystem, feedback loops can occur at various levels of abstraction. For example, LLMs can be improved through human feedback within a Reinforcement Learning from Human Feedback (RLHF) framework. This approach leverages human insights to refine AI models, ensuring they evolve in ways that align with human needs and ethical standards. As a result, LLMs will be frequently and continuously modified based on human feedback. This creates a distinct feedback loop, different from autophagy, where human input shapes the next generation of LLMs, which in turn influences future human feedback.

The fact that recommenders have substantial and often “unintended” consequences on these social phenomena is central to human-AI coevolution, as they might amplify trends already in motion in society. Studies in different social science fields (e.g., political science, sociology, political economy, economics and psychology) have shown that the increase of individualism [154] and inequalities [155,156] and the retrenchment of public and social policies [157–162] are tangible realities across several countries. In this respect, human-AI coevolution as a field of study should also integrate insights from social sciences to evaluate potential societal consequences of feedback loop mechanisms. For reasons of brevity, we will put the spotlight on three examples of such an intersection, but our considerations extend well beyond: (i) individual utility vs common goods; (ii) ownership of “the means of recommendation”; and (iii) phenomena of inequality and concentration.

Research on recommenders is mostly designed for the maximisation of individual utility and profit for companies (e.g., [122,163] for the urban mapping ecosystem, [24,137,140,164,165] for the online retail ecosystem), and this generates a lack of consideration for common goods. A good starting point to understand the classical limitations of this approach is the debate on rational choice theory and methodological individualism [166,167]. Rational choice theory is grounded in the idea that individual agents act in society by following rational reasoning, and therefore maximising their utility. Methodological individualism is a sort of corollary and explains societal outcomes as the sum of the relational behaviours of all agents within society. Recommenders designed on the basis of methodological individualism might not take into account collective utility.

History offers plenty of examples of societal damage generated by the prevalence of individualistic over collective behaviours. One could go back to Thomas More and “the drama of enclosures” to conceptualise the potential paradoxes that the coevolution between humans and recommenders might generate. The enclosures were a vast movement, which privatised land that had been used collectively. Similarly, employing individual utility as the main driver to set up recommenders can reduce the potential collective utility provided by the development of AI, and therefore, reinforce polarisation mechanisms already existent in society.

Human-AI coevolution is connected to the functioning of capitalism: the dramatic technological developments related to recommenders and AI are taking place in a period dominated by rational choice theory, methodological individualism and ultimately neoliberal economics [168]. Our reflection extends Kean Birch’s notion of “automated neoliberalism” [169], which posits that digital platforms shape markets, personal data accumulation transforms individual lives, and algorithms have the potential to automate social relations.

The potential for improving collective utility might be reduced because the coevolution between AI and society takes place in a context where there is no balance between those who own mainstream recommenders and those who use them uncritically. The old question about who owns the means of production strikingly applies today and takes a new form in an environment of human-AI coevolution. In this respect, owning the platforms that generate the largest number of interactions through mainstream recommenders – what we call “means of recommendations” – might be a game changer in economics and society. Readapting Marx’s reasoning [170], who owns the means of recommendations? How do these recommendations affect subsequent interactions between humans and AI? How could this reinforce outcomes such as inequalities and polarisation?

It is not the scope of this perspective article to delineate a political economy of human-AI coevolution [171], intended as a systematic analysis of the interconnection between asset ownership and the influence of recommenders. However, this aspect is a fundamental contextual element to discuss, at least theoretically, the potential unintended consequences of the coevolution between humans and recommenders. If recommenders have a powerful influence in shaping individual choices, and these individual choices shape collective outcomes, we have to reflect on how these choices are taken. Moreover, the impact of coevolution can further strengthen outcomes and increasingly harm collective utility. The opposite reasoning is also valid, and a different socio-economic context, where recommenders are geared towards shared collective objectives, can foster positive societal outcomes through coevolution. This point underlines that mere techno-solutionism [172], the idea that technology is the answer to any challenge we face, is a risky intellectual posture, and we need more studies on human-AI coevolution to steer positive outcomes. A review of recent research shows that AI tendentially increases the social divide, especially for historically marginalised groups (e.g., research about the US points to racial and gender effects) [173]. These patterns are even stronger in low/middle income countries [174]. One could reasonably hypothesise that the uneven impacts of AI are underestimated or overestimated because we miss research measuring human-AI coevolution and feedback loop mechanisms.

Human-AI coevolution might, therefore, foster inequality and concentration across different ecosystems. Within the social media ecosystem, recommenders may intensify exposure inequalities at the individual level, exacerbating the rich-get-richer effect regardless of user attributes or network characteristics [114]. Concerning the online retail ecosystem, purchase-based collaborative filtering might increase the number of items purchased and the variety of products considered. However, this might also push users to buy the same products, contributing to concentration at the systemic level [165]. In the urban mapping ecosystem, ride-hailing recommenders (like Uber and Lyft) may mislead low ridership in poor and black neighbourhoods as a reflection of low users’ demand, reinforcing in this way existing racial and socio-economic inequalities [175,176]. The autophagy process in a content generation ecosystem may lead to diversity loss in the generated content [127].

5. Open challenges of human-AI coevolution

The study of human-AI coevolution presents important challenges for the future that can be conceptualised at increasing levels of abstraction, i.e., scientific, legal and socio-political (see Fig. 3).

From a scientific perspective, we need a method to continually measure the impact of the feedback loop on the behaviour of humans and recommenders. Such a method could be developed by tracking step-wise how the measured outcomes change every time the recommender is re-trained. For example, one can measure the variety of products purchased by users at time t_0 , assessing the

OPEN CHALLENGES OF COEVOLUTION AI

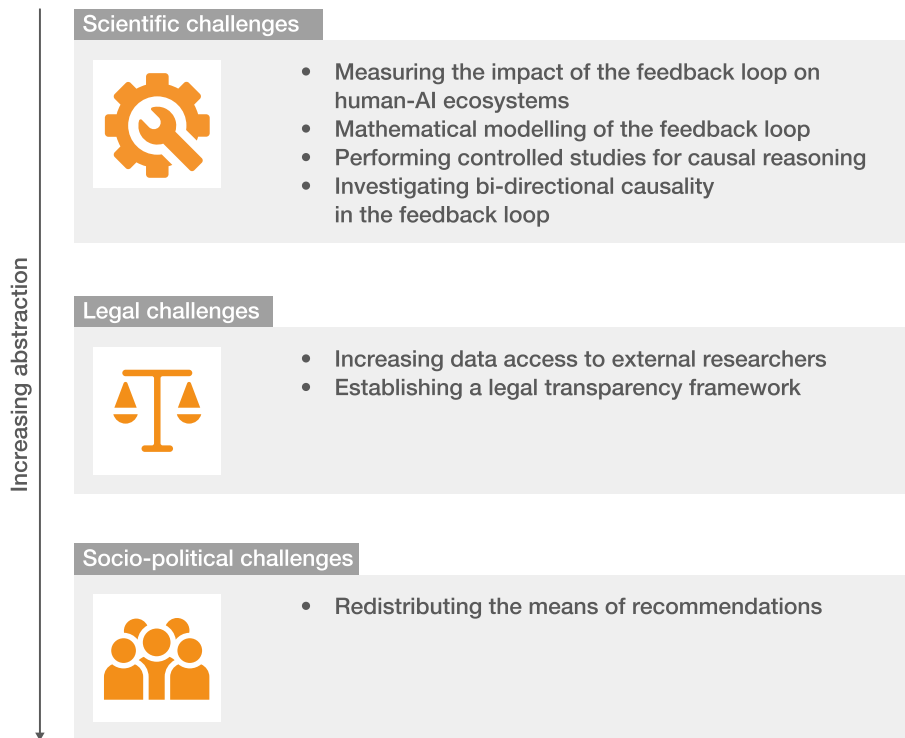


Fig. 3. Human-AI coevolution presents important challenges for the future, that can be conceptualised at increasing levels of abstraction: scientific, legal and socio-political.

extent to which feedback loop mechanisms might amplify or reduce this variety with successive steps, e.g., t_1, t_2, \dots, t_n . How many successive feedback loop iterations might be required to alter this variety of purchases? In a similar fashion, we could track how the recommender changes after subsequent re-training rounds: how long it takes a generative AI model to collapse and lose linguistic diversity?

An important technical aspect of studying human-AI coevolution is the occurrence of biases at different stages of the feedback loop [177]: data collection from users may introduce biases such as user-selection and exposure biases; learning of recommenders from collected data can result in inductive bias; and returning recommendation results to users can lead to popularity bias and unfairness. These biases as well as other issues like data contamination, can shape feedback loop mechanisms and individual, item, model, and systemic outcomes. The interactions among biases at the different stages of the feedback loop are not sufficiently measured in the current literature.

The conformism/diversity balance is yet another fundamental aspect related to the feedback loop that deserves more attention. At a general level, we also need a conscious effort in mathematical modelling to capture feedback loop mechanisms and their impact on human-AI ecosystems. At the epistemological level, understanding the causal interplay between humans and recommenders is key; however, much of the existing literature overlooks coevolution. Some steps could be undertaken to mitigate this issue, such as the development of controlled studies accounting for feedback loop mechanisms. Moreover, we must move beyond a unidirectional view of causality and explore it bi-directionally: humans and recommenders exert continuous influence on each other, necessitating a holistic study of their coevolutionary dynamics.

Beyond scientific challenges, other barriers can prevent our capacity to study the emerging phenomenon of human-AI coevolution, e.g., the limited access to data for researchers that are external to platforms and the lack of transparency on how recommenders are designed and employed within different platforms [15,80]. This severely impairs the reproducibility and replicability of potential studies about human-AI coevolution. Initiatives like EU's Digital Services Act might mitigate this barrier, but it remains unclear how vetted researchers will be allowed to access privately-owned platforms. Besides a new legal transparency framework, an intriguing way to overtake these barriers might involve the development of specialized APIs. APIs can allow external researchers to interact with platforms, and conduct empirical controlled experiments by changing recommenders' parameters or building experimental and control groups of users. Governments should foster a culture of impact evaluation of feedback loop mechanisms among platforms. This might follow the model already in place to reduce and compensate for other forms of negative externalities, e.g., pollution and drug undesired effects.

Increasing transparency also requires dealing with other fundamental challenges at the socio-political level. The concentration of what we defined as "the means of recommendations" is crucial in this perspective. In a context where big tech companies enjoy a

situation of oligopoly, recommenders are calibrated to generate high profits for the few. In this respect, a socially desirable development of human-AI coevolution as a field of study might be hindered by the lack of political intervention to redistribute the means of recommendation across markets with many smaller players and, more broadly, society. Such a configuration could be conducive to developing more transparent rules in data access and a fairer distribution of the means of recommendation. In the long term, small changes in the functioning of recommenders or humans' behaviour might lead to a significant impact on social outcomes, be they positive or negative. This could create a sort of “butterfly effect” of the feedback loop [178] that we have to study and understand.

The discussed challenges will persist even in the case of potential deep transformations in future online platforms, such as the emergence of decentralised platform architectures, user ownership of data and AI agents, and novel models for platform oversight and governance. These challenges not only concern scholars and their capacity to do research, but also go well beyond and bear upon the social and political realm. Only by accurately measuring and understanding the influence of recommenders on human behaviour we will be able to inform policymakers on how to make sensible decisions “taking into account not only the world as it is, but the world as it will be” [1]. This understanding will be key to designing adequate policies to avoid the potential negative externalities of an uncontrolled coevolution between humans and recommenders. We aim at a future society-centric AI that is part of the solution to long-standing societal problems, instead of being part of the problem.

CRediT authorship contribution statement

Dino Pedreschi: Writing – review & editing, Writing – original draft, Validation, Conceptualization. **Luca Pappalardo:** Writing – review & editing, Writing – original draft, Validation, Supervision, Project administration, Conceptualization. **Emanuele Ferragina:** Writing – review & editing, Validation, Supervision, Conceptualization. **Ricardo Baeza-Yates:** Validation. **Albert-László Barabási:** Validation. **Frank Dignum:** Validation. **Virginia Dignum:** Validation. **Tina Eliassi-Rad:** Writing – original draft, Validation. **Fosca Giannotti:** Validation, Funding acquisition. **János Kertész:** Writing – review & editing, Writing – original draft, Validation. **Alistair Knott:** Writing – review & editing, Writing – original draft, Validation. **Yannis Ioannidis:** Validation. **Paul Lukowicz:** Validation, Funding acquisition. **Andrea Passarella:** Validation. **Alex Sandy Pentland:** Validation. **John Shawe-Taylor:** Validation. **Alessandro Vespignani:** Validation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work has been partially supported by:

- PNRR - M4C2 - Investimento 1.3, Partenariato Esteso (grant No. PE00000013) - “FAIR - Future Artificial Intelligence Research” - Spoke 1 “Human-centered AI”, funded by the European Commission under the Next Generation EU programme;
- EU H2020 projects HumaneAI-net G.A. 952026 and SoBigData++ G.A. 871042;
- European Research Council ERC-2018-ADG 834756 “XAI: Science and technology for the eXplanation of AI decision making”;
- CHIST-ERA grant CHIST-ERA-19-XAI-010, by MUR (grant No. not available), FWF (grant No. I 5205), EPSRC (grant No. EP/V055712/1), NCN (grant No. 2020/02/Y/ST6/00064), ETAG (grant No. SLTAT21096), BNSF (grant No. KPI-06-AOO2/5);
- PNRR (Piano Nazionale di Ripresa e Resilienza) in the context of the research program 20224CZ5X4_PE6_PRIN 2022 “URBAI – Urban Artificial Intelligence” (CUP B53D23012770006), funded by European Union – Next Generation EU;
- Emanuele Ferragina has been partially supported by a public grant overseen by the French National Research Agency (ANR) as part of the ‘Investissements d’Avenir’ program LIEPP (ANR-11-LABX-0091, ANR-11-IDEX-0005-02) and the Université de Paris IdEx (ANR-18-IDEX-0001).

We thank Vincenzo Vivarini for the inspiration about the importance of the feedback loop in complex social systems. There is always a human agent beyond any tactical system. We also thank Daniele Fadda for the support on making the figures. Finally, we thank all members of KDD-Lab for insightful discussions about human-AI coevolution.

Data availability

No data was used for the research described in the article.

References

- [1] I. Asimov, *Asimov on Science Fiction*, Doubleday, New York City, US, 1981.
- [2] E. Russell, *Coevolutionary history*, *Am. Hist. Rev.* 119 (5) (2014) 1514–1528.
- [3] E.A. Lee, *The Coevolution: The Entwined Futures of Humans and Machines*, MIT Press, Cambridge, US, 2020.
- [4] F. Geels, *Co-evolution of technology and society: the transition in water supply and personal hygiene in the Netherlands (1850–1930)—a case study in multi-level perspective*, *Technol. Soc.* 27 (3) (2005) 363–397.

- [5] J. Moky, *The Fourth Discontinuity: The Co-Evolution of Humans and Machines*, 1995, JSTOR.
- [6] A. Santosuosso, About coevolution of humans and intelligent machines: preliminary notes, *BioLaw J. (Rivista di BioDiritto)* 15 (2021) 445–454.
- [7] F. Ricci, L. Rokach, B. Shapira, *Recommender Systems: Introduction and Challenges*, *Recommender Systems Handbook*, 2015, pp. 1–34.
- [8] Y. Li, K. Liu, R. Satapathy, S. Wang, E. Cambria, Recent developments in recommender systems: a survey, arXiv:2306.12680, 2023.
- [9] J. Lu, D. Wu, M. Mao, W. Wang, G. Zhang, *Recommender system application developments: a survey*, *Decis. Support Syst.* 74 (2015) 12–32.
- [10] C. Wagner, M. Strohmaier, A. Olteanu, E. Kiciman, N. Contractor, T. Eliassi-Rad, Measuring algorithmically infused societies, *Nature* 595 (7866) (2021) 197–204.
- [11] R. Jiang, S. Chiappa, T. Lattimore, A. György, P. Kohli, Degenerate feedback loops in recommender systems, in: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, Association for Computing Machinery, New York, NY, USA, 2019, pp. 383–390.
- [12] W. Sun, S. Khenissi, O. Nasraoui, P. Shafto, Debiasing the human-recommender system feedback loop in collaborative filtering, in: *Companion Proceedings of the 2019 World Wide Web Conference, WWW '19*, Association for Computing Machinery, New York, NY, USA, 2019, pp. 645–651.
- [13] M. Mansoury, H. Abdollahpour, M. Pechenizkiy, B. Mobasher, R. Burke, Feedback loop and bias amplification in recommender systems, in: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 2145–2148.
- [14] A. Sirbu, D. Pedreschi, F. Giannotti, J. Kertész, Algorithmic bias amplifies opinion fragmentation and polarization: a bounded confidence model, *PLoS ONE* 14 (3) (2019) 0213246.
- [15] Alistair Knott, Kate Hannah, Dino Pedreschi, Tapabrata Chakraborti, Sanjana Hattotuwa, Andrew Trotman, Ricardo Baeza-Yates, Rituparna Roy, David Eysers, Virginia Morini, Valentina Pansanella, *Responsible AI for social media guidance: a proposed collaborative method for studying the effects of social media recommender systems on users*, Technical report, Global Partnership on AI (GPAI), 2021.
- [16] F.O. Isinkaye, Y.O. Fojajimi, B.A. Ojokoh, Recommendation systems: principles, methods and evaluation, *Egypt. Inform. J.* 16 (3) (2015) 261–273.
- [17] B. Kulynych, R. Overdorf, C. Troncoso, S. Gürses, Pots: protective optimization technologies, in: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 177–188.
- [18] M. Del Vicario, G. Vivaldo, A. Bessi, F. Zollo, A. Scala, G. Caldarelli, W. Quattrociocchi, Echo chambers: emotional contagion and group polarization on Facebook, *Sci. Rep.* 6 (2016) 37825.
- [19] N. Perra, L.E.C. Rocha, Modelling opinion dynamics in the age of algorithmic personalisation, *Sci. Rep.* 9 (1) (2019) 7261.
- [20] F. Huszár, S.I. Ktena, C. O'Brien, L. Belli, A. Schlaikjer, M. Hardt, Algorithmic amplification of politics on Twitter, *Proc. Natl. Acad. Sci.* 119 (1) (2022).
- [21] D. Pedreschi, S. Ruggieri, F. Turini, Discrimination-aware data mining, in: *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD 08*, 2008, pp. 560–568.
- [22] J. Kleinberg, J. Ludwig, S. Mullainathan, C.R. Sunstein, Algorithms as discrimination detectors, *Proc. Natl. Acad. Sci.* 117 (48) (2020) 30096–30100.
- [23] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, A. Galstyan, A survey on bias and fairness in machine learning, *ACM Comput. Surv.* 54 (6) (2021) 1–35.
- [24] K. Hosanagar, D. Fleder, D. Lee, A. Buja, Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation, *Manag. Sci.* 60 (4) (2014) 805–823.
- [25] J. Macfarlane, Your navigation app is making traffic unmanageable, *IEEE Spectr.* (2019) 22–27.
- [26] S. Siuhi, J. Mwakalonge, Opportunities and challenges of smart mobile applications in transportation, *J. Traffic Transp. Eng.* 3 (6) (2016) 582–592.
- [27] L.W. Foderaro, Navigation Apps Are Turning Quiet Neighborhoods into Traffic Nightmares, *The New York Times*, 2017.
- [28] G. Cornacchia, M. Böhm, G. Mauro, M. Nanni, D. Pedreschi, L. Pappalardo, How routing strategies impact urban emissions, in: *Proceedings of the 30th International Conference on Advances in Geographic Information Systems, Association for Computing Machinery, New York, NY, USA, 2022*, pp. 1–4.
- [29] A. Lima, R. Stanojevic, D. Papagiannaki, P. Rodriguez, M.C. González, Understanding individual routing behaviour, *J. R. Soc. Interface* 13 (116) (2016).
- [30] S. Çolak, A. Lima, M.C. González, Understanding congested travel in urban areas, *Nat. Commun.* 7 (1) (2016) 1–8.
- [31] G. Cornacchia, M. Nanni, L. Pappalardo, One-shot traffic assignment with forward-looking penalization, in: *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems, Association for Computing Machinery, New York, NY, USA, 2023*, pp. 1–10.
- [32] G. Cornacchia, M. Nanni, D. Pedreschi, L. Pappalardo, Navigation services and urban sustainability, *Fluct. Noise Lett.* 2450016 (2023).
- [33] G. Cornacchia, M. Nanni, D. Pedreschi, L. Pappalardo, Navigation services amplify concentration of traffic and emissions in our cities, <https://arxiv.org/abs/2407.20004>, 2024.
- [34] T.J. Sejnowski, *The Deep Learning Revolution*, MIT Press, Cambridge, US, 2018.
- [35] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444.
- [36] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. Drissi, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis, Mastering the game of go with deep neural networks and tree search, *Nature* 529 (7587) (2016) 484–489.
- [37] J.J. Titano, M. Badgeley, J. Schefflein, M. Pain, A. Su, M. Cai, N. Swinburne, J. Zech, J. Kim, J. Bederson, J. Mocco, B. Drayer, J. Lehar, S. Cho, A. Costa, E.K. Oermann, Automated deep-neural-network surveillance of cranial images for acute neurologic events, *Nat. Med.* 24 (9) (2018) 1337–1341.
- [38] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, *Commun. ACM* 60 (6) (2017) 84–90.
- [39] Y. Wu, M. Schuster, Z. Chen, Q.V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, J. Dean, Google's neural machine translation system: bridging the gap between human and machine translation, arXiv:1609.08144 [cs.CL], 2016.
- [40] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, D. Pedreschi, A survey of methods for explaining black box models, *ACM Comput. Surv.* 51 (5) (2019) 1–42.
- [41] N. Burkart, M.F. Huber, A survey on the explainability of supervised machine learning, *J. Artif. Intell. Res.* 70 (2021) 245–317.
- [42] J.M. Hofman, D.J. Watts, S. Athey, F. Garip, T.L. Griffiths, J. Kleinberg, H. Margetts, S. Mullainathan, M.J. Salganik, S. Vazire, A. Vespignani, T. Yarkoni, Integrating explanation and prediction in computational social science, *Nature* 595 (7866) (2021) 181–188.
- [43] A. Abdul, J. Vermeulen, D. Wang, B.Y. Lim, M. Kankanhalli, Trends and trajectories for explainable, accountable and intelligible systems, in: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–18.
- [44] P. Lukowicz, The challenge of human centric AI, *Digit. Welt* 3 (4) (2019) 9–10.
- [45] D. Horvatić, T. Lipic, Human-centric AI: the symbiosis of human and artificial intelligence, *Entropy* 23 (3) (2021).
- [46] B. Lepri, N. Oliver, A. Pentland, Ethical machines: the human-centric use of artificial intelligence, *iScience* 24 (3) (2021) 102249.
- [47] V. Dignum, Relational artificial intelligence, arXiv:2202.07446 [cs.CY], 2022.
- [48] D.J. Watts, S.H. Strogatz, Collective dynamics of 'small-world' networks, *Nature* 393 (6684) (1998) 440–442.
- [49] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, *Science* 286 (5439) (1999) 509–512.
- [50] J.M. Kleinberg, Navigation in a small world, *Nature* 406 (6798) (2000) 845.
- [51] R.Z. Albert, Statistical mechanics of complex networks, *Rev. Mod. Phys.* 74 (87) (2002).
- [52] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, A.-L. Barabási, Structure and tie strengths in mobile communication networks, *Proc. Natl. Acad. Sci.* 104 (18) (2007) 7332–7336.
- [53] M. Newman, *Networks: An Introduction*, Oxford University Press, Oxford, UK, 2018.
- [54] A.-L. Barabási, *Network Science*, Cambridge University Press, Cambridge, UK, 2016.
- [55] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, A. Vespignani, Epidemic processes in complex networks, *Rev. Mod. Phys.* 87 (3) (2015) 925–979.

- [56] V. Colizza, A. Barrat, M. Barthelemy, A.-J. Valleron, A. Vespignani, Modeling the worldwide spread of pandemic influenza: baseline case and containment interventions, *PLoS Med.* 4 (1) (2007) 13.
- [57] X. Wang, B. Yuceysoy, O. Varol, T. Eliassi-Rad, A.-L. Barabási, Success in books: predicting book sales before publication, *EPJ Data Sci.* 8 (1) (2019).
- [58] S.P. Fraiberger, R. Sinatra, M. Resch, C. Riedl, A.-L. Barabási, Quantifying reputation and success in art, *Science* 362 (6416) (2018) 825–829.
- [59] L. Pappalardo, P. Cintia, Quantifying the relation between performance and success in soccer, *Adv. Complex Syst.* 21 (03n04) (2018) 1750014.
- [60] D. Brockmann, L. Hufnagel, T. Geisel, The scaling laws of human travel, *Nature* 439 (7075) (2006) 462–465.
- [61] M.C. González, C.A. Hidalgo, A.-L. Barabási, Understanding individual human mobility patterns, *Nature* 458 (7235) (2008) 779–782.
- [62] L. Pappalardo, F. Simini, S. Rinzivillo, D. Pedreschi, F. Giannotti, A.-L. Barabási, Returners and explorers dichotomy in human mobility, *Nat. Commun.* 6 (2015) 8166.
- [63] L. Alessandretti, U. Aslak, S. Lehmann, The scales of human mobility, *Nature* 587 (7834) (2020) 402–407.
- [64] M. Böhm, M. Nanni, L. Pappalardo, Gross polluters and vehicle emissions reduction, *Nat. Sustain.* (2022) 1–9.
- [65] M. Luca, G. Barlacchi, B. Lepri, L. Pappalardo, A survey on deep learning for human mobility, *ACM Comput. Surv.* 55 (1) (2021) 1–44.
- [66] H. Barbosa, M. Barthelemy, G. Ghoshal, C.R. James, M. Lenormand, T. Louail, R. Menezes, J.J. Ramasco, F. Simini, M. Tomasini, Human mobility: models and applications, *Phys. Rep.* 734 (2018) 1–74.
- [67] F. Simini, G. Barlacchi, M. Luca, L. Pappalardo, A deep gravity model for mobility flows generation, *Nat. Commun.* 12 (1) (2021) 6576.
- [68] R. Hanna, G. Kreindler, B.A. Olken, Citywide effects of high-occupancy vehicle restrictions: evidence from “three-in-one” in Jakarta, *Science* 357 (6346) (2017) 89–93.
- [69] L. Pappalardo, E. Manley, V. Sekara, L. Alessandretti, Future directions in human mobility science, *Nat. Comput. Sci.* 3 (7) (2023) 588–600.
- [70] Q. Yang, A. Steinfeld, C. Rosé, J. Zimmerman, Re-examining whether, why, and how human-AI interaction is uniquely difficult to design, in: *CHI '20: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems 2020*, 2020, pp. 1–13.
- [71] S. Thurner, R. Hanel, P. Klimek, *Introduction to the Theory of Complex Systems*, Oxford University Press, Oxford, UK, 2018.
- [72] M.W. Eysenck, C. Eysenck, *AI Vs Humans*, Routledge, London, UK, 2021.
- [73] W.F. Lawless, D.A. Sofge, D. Lofaro, R. Mittu, Editorial: interdisciplinary approaches to the structure and performance of interdependent autonomous human machine teams and systems, *Front. Phys.* 11 (2023) 1.
- [74] J. Piao, J. Liu, F. Zhang, J. Su, Y. Li, Human-AI adaptive dynamics drives the emergence of information cocoons, *Nat. Mach. Intell.* 5 (2023) 1214–1224.
- [75] P. Contucci, J. Kertész, G. Osabutey, Human-AI ecosystem with abrupt changes as a function of the composition, *PLoS ONE* 17 (5) (2022) 0267310.
- [76] I. Rahwan, M. Cebrian, N. Obradovich, J. Bongard, J.-F. Bonnefon, C. Breazeal, J.W. Crandall, N.A. Christakis, I.D. Couzin, M.O. Jackson, et al., Machine behaviour, *Nature* 568 (7753) (2019) 477–486.
- [77] M.M. Peeters, J. Diggelen, K. Van Den Bosch, A. Bronkhorst, M.A. Neerinx, J.M. Schraagen, S. Raaijmakers, Hybrid collective intelligence in a human-AI society, *AI Soc.* 36 (2021) 217–238.
- [78] A. Deaton, N. Cartwright, Understanding and misunderstanding randomized controlled trials, *Soc. Sci. Med.* 210 (2018) 2–21.
- [79] E. Hariton, J.J. Locascio, Randomised controlled trials - the gold standard for effectiveness research, *BJOG* 125 (13) (2018) 1716.
- [80] Alistair Knott, Dino Pedreschi, Transparency mechanisms for social media recommender algorithms: from proposals to action, Technical report, Global Partnership on AI (GPAI), 2022.
- [81] A. Pera, G.d.F. Morales, L.M. Aiello, Measuring behavior change with observational studies: a review, arXiv:2310.19951, 2023.
- [82] E. Bakshy, S. Messing, L.A. Adamic, Exposure to ideologically diverse news and opinion on Facebook, *Science* 348 (6239) (2015) 1130–1132.
- [83] N. Arora, T. Cabannes, S. Ganapathy, Y. Li, P. McAfee, M. Nunkesser, C. Osorio, A. Tomkins, I. Tsogsuren, Quantifying the sustainability impact of Google maps: a case study of Salt Lake City, arXiv:2111.03426 [physics.soc-ph], 2021.
- [84] S. Flaxman, S. Goel, J.M. Rao, Filter bubbles, echo chambers, and online news consumption, *Public Opin. Q.* 80 (S1) (2016) 298–320.
- [85] H. Hosseinmardi, A. Ghasemian, A. Clauset, et al., Evaluating the scale, growth, and origins of right-wing echo chambers on YouTube, arXiv:2111.03426, 2020.
- [86] M. Ledwich, A. Zaitsev, Algorithmic extremism: examining YouTube’s rabbit hole of radicalization, arXiv preprint, arXiv:1912.11211, 2019.
- [87] J. Whittaker, S. Looney, A. Reed, F. Votta, Recommender systems and the amplification of extremist content, *Int. Policy Rev.* 10 (2) (2021).
- [88] M.H. Ribeiro, R. Ottoni, R. West, V.A.F. Almeida, W. Meira, Auditing radicalization pathways on YouTube, in: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 131–141.
- [89] K. Papadamou, S. Zannettou, J. Blackburn, E. De Cristofaro, G. Stringhini, M. Sirivianos, “How over is it?” Understanding the incel community on YouTube, in: *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, Association for Computing Machinery, New York City, USA, 2021, pp. 1–25.
- [90] H. Allcott, M. Gentzkow, C. Yu, Trends in the diffusion of misinformation on social media, *Res. Polit.* 6 (2) (2019).
- [91] N. Asimovic, J. Nagler, R. Bonneau, J.A. Tucker, Testing the effects of Facebook usage in an ethnically polarized setting, *Proc. Natl. Acad. Sci.* 118 (25) (2021).
- [92] R. Levy, Social media, news consumption, and polarization: evidence from a field experiment, *Am. Econ. Rev.* 111 (3) (2021) 831–870.
- [93] J. Cho, S. Ahmed, M. Hilbert, B. Liu, J. Luu, Do search algorithms endanger democracy? An experimental investigation of algorithm effects on political polarization, *J. Broadcast. Electron. Media* 64 (2) (2020) 150–172.
- [94] D. Fleder, K. Hosanagar, A. Buja, Recommender systems and their effects on consumers, in: *Mediterranean Conference on Information Systems (MCIS)*, 2011.
- [95] D. Ensign, S.A. Friedler, S. Neville, C. Scheidegger, S. Venkatasubramanian, Runaway feedback loops in predictive policing, in: *Conference on Fairness, Accountability and Transparency*, PMLR, 2018, pp. 160–171.
- [96] T.T. Nguyen, P.-M. Hui, F.M. Harper, L. Terveen, J.A. Konstan, Exploring the filter bubble: the effect of using recommender systems on content diversity, in: *Proceedings of the 23rd International Conference on World Wide Web*, 2014, pp. 677–686.
- [97] C.M. Valensise, M. Cinelli, W. Quattrociocchi, The drivers of online polarization: fitting models to data, *Inf. Sci.* 642 (2023) 119152.
- [98] V. Pansanella, G. Rossetti, L. Milli, From mean-field to complex topologies: network effects on the algorithmic bias model, in: R.M. Benito, C. Cherifi, H. Cherifi, E. Moro, L.M. Rocha, M. Sales-Pardo (Eds.), *Complex Networks & Their Applications X*, Springer, Cham, 2022, pp. 329–340.
- [99] A.F. Peralta, M. Neri, J. Kertész, G. Iniguez, Effect of algorithmic bias and network structure on coexistence, consensus, and polarization of opinions, *Phys. Rev. E* 104 (4–1) (2021) 044312.
- [100] A.F. Peralta, J. Kertész, G. Iniguez, Opinion formation on social networks with algorithmic bias: dynamics and bias imbalance, *J. Phys. Complex.* 2 (4) (2021).
- [101] V. Pansanella, A. Stirbu, J. Kertesz, G. Rossetti, Mass media impact on opinion evolution in biased digital environments: a bounded confidence model, *Sci. Rep.* 13 (1) (2023) 14600.
- [102] N. Liu, M.A. Baum, A.J. Berinsky, A.J. Chaney, J. Benedictis-Kessner, A. Guess, D. Knox, C. Lucas, R. Mariman, B.M. Stewart, Algorithmic recommendations have limited effects on polarization: a naturalistic experiment on YouTube, Technical report, Harvard Kennedy School, 2023.
- [103] U. Chitra, C. Musco, Analyzing the impact of filter bubbles on social network polarization, in: *Proceedings of the 13th International Conference on Web Search and Data Mining*, 2020, pp. 115–123.
- [104] M. Haroon, M. Wojcieszak, A. Chhabra, X. Liu, P. Mohapatra, Z. Shafiq, Auditing YouTube’s recommendation system for ideologically congenial, extreme, and problematic recommendations, *Proc. Natl. Acad. Sci.* 120 (50) (2023) 2213020120.
- [105] P. Bouchaud, D. Chavalarias, M. Panahi, Crowdsourced audit of Twitter’s recommender systems, *Sci. Rep.* 13 (1) (2023) 16815.
- [106] A.M. Guess, N. Malhotra, J. Pan, P. Barberá, H. Allcott, T. Brown, A. Crespo-Tenorio, D. Dimmery, D. Freelon, M. Gentzkow, et al., How do social media feed algorithms affect attitudes and behavior in an election campaign?, *Science* 381 (6656) (2023) 398–404.
- [107] M.S. Alvim, B. Amorim, S. Knight, S. Quintero, F. Valencia, A formal model for polarization under confirmation bias in social networks, *Log. Methods Comput. Sci.* 19 (2023).

- [108] I. Rahaman, P. Hosein, A model for optimizing article recommendation for reducing polarization, in: Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM '21, Association for Computing Machinery, New York, NY, USA, 2022, pp. 107–110.
- [109] C. Yang, X. Xu, B.P. Nunes, S.W.M. Siqueira, Bubbles bursting: investigating and measuring the personalisation of social media searches, *Telemat. Inform.* 82 (2023) 101999.
- [110] V. Pansanella, G. Rossetti, L. Milli, Modeling algorithmic bias: simplicial complexes and evolving network topologies, *Appl. Netw. Sci.* 7 (1) (2022) 57.
- [111] F. Cinus, M. Minici, C. Monti, F. Bonchi, The effect of people recommenders on echo chambers and polarization, in: Proceedings of the International AAAI Conference on Web and Social Media, vol. 16, Association for the Advancement of Artificial Intelligence, Washington, DC USA, 2022, pp. 90–101.
- [112] M. Boeker, A. Urman, An empirical investigation of personalization factors on TikTok, in: Proceedings of the ACM Web Conference 2022, 2022, pp. 2298–2309.
- [113] N. Bartley, A. Abeliuk, E. Ferrara, K. Lerman, Auditing algorithmic bias on Twitter, in: Proceedings of the 13th ACM Web Science Conference 2021, 2021, pp. 65–73.
- [114] F. Fabbri, M.L. Croci, F. Bonchi, C. Castillo, Exposure inequality in people recommender systems: the long-term effects, in: Proceedings of the International AAAI Conference on Web and Social Media, vol. 16, Association for the Advancement of Artificial Intelligence, Washington, DC USA, 2022, pp. 194–204.
- [115] B. Kirdemir, J. Kready, E. Mead, M.N. Hussain, N. Agarwal, D. Adjero, Assessing bias in YouTube's video recommendation algorithm in a cross-lingual and cross-topical context, in: Social, Cultural, and Behavioral Modeling: 14th International Conference, SBP-BRiMS 2021, 2021, Virtual Event, July 6–9, 2021, in: Proceedings, vol. 14, Springer, 2021, pp. 71–80.
- [116] Y.M.M. Ng, K. Hoffmann Pham, M. Luengo-Oroz, Exploring YouTube's recommendation system in the context of Covid-19 vaccines: computational and comparative analysis of video trajectories, *J. Med. Internet Res.* 25 (2023) 49061.
- [117] G.D. Erhardt, S. Roy, D. Cooper, B. Sana, M. Chen, J. Castiglione, Do transportation network companies decrease or increase congestion?, *Sci. Adv.* 5 (5) (2019) 2670.
- [118] A.M. Falek, A. Gallais, C. Pelsser, S. Julien, F. Theoleyre, To re-route, or not to re-route: impact of real-time re-routing in urban road networks, *J. Intell. Transp. Syst.* 26 (2) (2022) 198–212.
- [119] P. Santi, G. Resta, M. Szell, S. Sobolevsky, S.H. Strogatz, C. Ratti, Quantifying the benefits of vehicle pooling with shareability networks, *Proc. Natl. Acad. Sci.* 111 (37) (2014) 13290–13294.
- [120] D.J. Fagnant, K.M. Kockelman, Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in Austin, Texas, *Transportation* 45 (2018) 143–158.
- [121] R. Jalali, S. Koohi-Fayegh, K. El-Khatib, D. Hoornweg, H. Li, Investigating the potential of ridesharing to reduce vehicle emissions, *Urban Plann.* 2 (2) (2017) 26–40.
- [122] S. Zhang, N. Mehta, P.V. Singh, K. Srinivasan, Frontiers: can an artificial intelligence algorithm mitigate racial economic inequality? An analysis in the context of Airbnb, *Mark. Sci.* 40 (5) (2021) 813–820.
- [123] V. Koh, W. Li, G. Livan, L. Capra, Offline biases in online platforms: a study of diversity and homophily in Airbnb, *EPJ Data Sci.* 8 (1) (2019) 11.
- [124] I. Srba, R. Moro, M. Tomlein, B. Pecher, J. Simko, E. Stefancova, M. Kompan, A. Hrcckova, J. Podrouzek, A. Gavornik, et al., Auditing YouTube's recommendation algorithm for misinformation filter bubbles, *ACM Trans. Recommend. Syst.* 1 (1) (2023) 1–33.
- [125] M. Tomlein, B. Pecher, J. Simko, I. Srba, R. Moro, E. Stefancova, M. Kompan, A. Hrcckova, J. Podrouzek, M. Bielikova, An audit of misinformation filter bubbles on YouTube: bubble bursting and recent behavior changes, in: Proceedings of the 15th ACM Conference on Recommender Systems, 2021, pp. 1–11.
- [126] Q. Grossetti, C. Du Mouza, N. Travers, Community-based recommendations on Twitter: avoiding the filter bubble, in: Web Information Systems Engineering–WISE 2019: 20th International Conference, Hong Kong, China, January 19–22, 2020, in: Proceedings, vol. 20, Springer, 2019, pp. 212–227.
- [127] I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, Y. Gal, AI models collapse when trained on recursively generated data, *Nature* 631 (8022) (2024) 755–759.
- [128] Y. Guo, G. Shang, M. Vazirgiannis, C. Clavel, The curious decline of linguistic diversity: training language models on synthetic text, arXiv:2311.09807, 2023.
- [129] M. Briesch, D. Sobania, F. Rothlauf, Large language models suffer from their own output: an analysis of the self-consuming training loop, arXiv:2311.16822, 2023.
- [130] E. Dohmatob, Y. Feng, P. Yang, F. Charton, J. Kempe, A tale of tails: model collapse as a change of scaling laws, arXiv:2402.07043, 2024.
- [131] S. Alemohammad, J. Casco-Rodriguez, L. Luzi, A.I. Humayun, H. Babaei, D. LeJeune, A. Siahkoohi, R.G. Baraniuk, Self-consuming generative models go mad, arXiv:2307.01850, 2023.
- [132] G. Martínez, L. Watson, P. Reviriego, J.A. Hernández, M. Juarez, R. Sarkar, Towards understanding the interplay of generative artificial intelligence and the Internet, arXiv:2306.06130, 2023.
- [133] E. Dohmatob, Y. Feng, J. Kempe, Model collapse demystified: the case of regression, arXiv:2402.07712, 2024.
- [134] M. Bohacek, H. Farid, Nepotistically trained generative-AI models collapse, arXiv:2311.12202, 2023.
- [135] R. Hataya, H. Bao, H. Arai, Will large-scale generative models corrupt future datasets?, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 20555–20565.
- [136] B. Pathak, R. Garfinkel, R.D. Gopal, R. Venkatesan, F. Yin, Empirical analysis of the impact of recommender systems on sales, *J. Manag. Inf. Syst.* 27 (2) (2010) 159–188.
- [137] P.-Y. Chen, S.-y. Wu, J. Yoon, The impact of online recommendations and consumer feedback on sales, in: International Conference on Information Systems (ICIS), 2004.
- [138] R. Donnelly, A. Kanodia, I. Morozov, The long tail effect of personalized rankings, Available at SSRN 3649342, 2021.
- [139] S. Chen, H. Qiu, S. Zhao, Y. Han, W. He, M. Siponen, J. Mou, H. Xiao, When more is less: the other side of artificial intelligence recommendation, *J. Manag. Sci. Eng.* 7 (2) (2022) 213–232.
- [140] S. Yi, D. Kim, J. Ju, Recommendation technologies and consumption diversity: an experimental study on product recommendations, consumer search, and sales diversity, *Technol. Forecast. Soc. Change* 178 (2022) 121486.
- [141] C. Matt, T. Hess, C. Weiß, The differences between recommender technologies in their impact on sales diversity, in: International Conference on Interaction Sciences, 2013.
- [142] G. Alves, D. Jannach, R.F. Souza, D. Damian, M.G. Manzato, Digitally nudging users to explore off-profile recommendations: here be dragons, *User Model. User-Adapt. Interact.* (2023) 1–41.
- [143] E. Noordeh, R. Levin, R. Jiang, H. Shadmany, Echo chambers in collaborative filtering based recommendation systems, arXiv:2011.03890, 2020.
- [144] G. Aridor, D. Goncalves, S. Sikdar, Deconstructing the filter bubble: user decision-making and recommender systems, in: Proceedings of the 14th ACM Conference on Recommender Systems, 2020, pp. 82–91.
- [145] L. Pappalardo, E. Ferragina, S. Citraro, G. Cornacchia, M. Nanni, G. Rossetti, G. Gezici, F. Giannotti, M. Lalli, D. Gambetta, G. Mauro, V. Morini, V. Pansanella, D. Pedreschi, A survey on the impact of AI-based recommenders on human behaviours: methodologies, outcomes and future directions, <https://arxiv.org/abs/2407.01630>, 2024.
- [146] D. Fleder, K. Hosanagar, Blockbuster culture's next rise or fall: the impact of recommender systems on sales diversity, *Manag. Sci.* 55 (5) (2009) 697–712.
- [147] D. Lee, K. Hosanagar, How do recommender systems affect sales diversity? A cross-category investigation via randomized field experiment, *Inf. Syst. Res.* 30 (1) (2019) 239–259.
- [148] M.J. Hetherington, Putting polarization in perspective, *Br. J. Polit. Sci.* 39 (2) (2009) 413–448.
- [149] A.I. Abramowitz, K.L. Saunders, Is polarization a myth?, *J. Polit.* 70 (2) (2008) 542–555.

- [150] A.B. Atkinson, et al., On the measurement of inequality, *J. Econ. Theory* 2 (3) (1970) 244–263.
- [151] U. Beck, Beyond class and nation: reframing social inequalities in a globalizing world 1, *Br. J. Sociol.* 58 (4) (2007) 679–705.
- [152] T.C. Schelling, Models of segregation, *Am. Econ. Rev.* 59 (2) (1969) 488–493.
- [153] J. Jacobs, *The Death and Life of Great American Cities*, Random House, New York, 1961.
- [154] H.C. Santos, M.E. Varnum, I. Grossmann, Global increases in individualism, *Psychol. Sci.* 28 (9) (2017) 1228–1239.
- [155] B. Milanovic, *Global Inequality: A New Approach for the Age of Globalization*, Harvard University Press, Cambridge, USA, 2016.
- [156] T. Piketty, *Capital in the Twenty-First Century*, Harvard University Press, Cambridge, USA, 2014.
- [157] E. Ferragina, Welfare state change as a double movement: four decades of retrenchment and expansion in compensatory and employment-oriented policies across 21 high-income countries, *Soc. Policy Adm.* 56 (5) (2022) 705–725.
- [158] E. Ferragina, The political economy of family policy expansion: fostering neoliberal capitalism or promoting gender equality supporting social reproduction?, *Rev. Int. Polit. Econ.* 26 (6) (2019) 1238–1265.
- [159] E. Ferragina, The ‘two lives’ of Esping-Andersen and the revival of a research programme: gender equality, employment and redistribution in contemporary social policy, *Soc. Policy Adm.* (2024) 1–19.
- [160] E. Ferragina, A. Arrigoni, T.F. Spreckelsen, The rising invisible majority: bringing society back into international political economy, *Rev. Int. Polit. Econ.* 29 (1) (2022) 114–151.
- [161] E. Ferragina, A. Arrigoni, Selective neoliberalism: how Italy went from dualization to liberalisation in labour market and pension reforms, *New Polit. Econ.* 26 (6) (2021) 964–984.
- [162] E. Ferragina, F.D. Filetti, Labour market protection across space and time: a revised typology and a taxonomy of countries’ trajectories of change, *J. Eur. Soc. Policy* 32 (2) (2022) 148–165.
- [163] M.-À. García-López, J. Jofre-Monseny, R. Martínez-Mazza, M. Segú, Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona, *J. Urban Econ.* 119 (2020) 103278.
- [164] M.B. Dias, D. Locher, M. Li, W. El-Deredey, P.J. Lisboa, The value of personalised recommender systems to e-business: a case study, in: *Proceedings of the 2008 ACM Conference on Recommender Systems*, 2008, pp. 291–294.
- [165] D. Lee, K. Hosanagar, Impact of recommender systems on sales volume and diversity, in: *International Conference on Interaction Sciences*, 2014.
- [166] J.A. Oppenheimer, Rational choice theory, in: *Encyclopedia of Political Theory*, vol. 3, 2008, pp. 1150–1159.
- [167] D.J. Kjosavik, Methodological individualism and rational choice in neoclassical economics: a review of institutionalist critique, in: *Forum for Development Studies*, vol. 30, Taylor & Francis, 2003, pp. 205–245.
- [168] D. Harvey, *A Brief History of Neoliberalism*, Oxford University Press, Oxford, UK, 2007.
- [169] K. Birch, Automated neoliberalism? The digital organisation of markets in technoscientific capitalism, *New Form.* 100 (100–101) (2020) 10–27.
- [170] K. Marx, *Capital, Volume I, A Critique of Political Economy*, vol. 1, Courier Corporation, North Chelmsford, MA, USA, 2011.
- [171] M. Trajtenberg, AI as the next GPT: a political-economy perspective, Technical report, National Bureau of Economic Research, 2018.
- [172] E. Morozov, To Save Everything, Click Here: The Folly of Technological Solutionism, PublicAffairs, New York City, USA, 2013.
- [173] L. Sartori, A. Theodorou, A sociotechnical perspective for the future of AI: narratives, inequalities, and human control, *Ethics Inf. Technol.* 24 (1) (2022) 4.
- [174] A. Hagerty, I. Rubinov, Global AI ethics: a review of the social impacts and ethical implications of artificial intelligence, arXiv:1907.07892, 2019.
- [175] Y. Ge, C.R. Knittel, D. MacKenzie, S. Zoepf, Racial discrimination in transportation network companies, *J. Public Econ.* 190 (2020) 104205.
- [176] A. Yan, B. Howe, Fairness-aware demand prediction for new mobility, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, Association for the Advancement of Artificial Intelligence, Washington, DC USA, 2020, pp. 1079–1087.
- [177] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, X. He, Bias and debias in recommender system: a survey and future directions, *ACM Trans. Inf. Syst.* 41 (3) (2023) 1–39.
- [178] R. Bradbury, A sound of thunder, in: *Science Fiction and Philosophy: From Time Travel to Superintelligence*, 2016, pp. 331–342.