Designing and Implementing Socially Beneficial Recommender Systems: An Interdisciplinary Approach

Bernard Mallia[1,2,3]

[1] Equinox Group, Malta
[2] Institute for Research and Improvement in Social Sciences, Malta
[3] Mediterranean Institute for Innovation, Communications and Technology, Malta

Abstract

This paper studies the complexities inherent in designing recommender systems that focus on social impact. In doing so, it looks at various schools of thought in assessing social outcomes and makes the case for a paradigm shift from prioritising user engagement to promoting 'positive social outcomes,' a term that proves challenging to define universally due to divergent stakeholder perspectives and interests, but which is nevertheless crucial for such a paradigm shift to occur. It also explores the divergences between commercial objectives and ethical imperatives, such as information diversity and user privacy. The paper proposes an interdisciplinary approach, incorporating machine learning, ethics and social sciences, to establish 'appropriate' norms and values that can be embedded into recommender systems, and concludes that while defining 'positive social outcomes' is complex, their technical implementation, once agreed upon, should be more straightforward. This development is posited as an interdisciplinary, collaborative endeavour requiring the use of both technological innovation and societal wisdom.

Keywords

Recommender systems, Recommender systems norms and values, Social welfare, socially-beneficial recommender systems, User engagement, Philosophy, Ethics.

1. Introduction

Recommender systems, or recommendation engines as they are also known, serve as essential conduits in the information ecosystem of digital platforms, shaping user interactions and content discovery, while greatly affecting consumption patterns. Recommender systems employ sophisticated algorithms and generate personalised recommendations that in practice, today, primarily aim to maximise user engagement and profit, as well as to hog the biggest possible chunk of the user's time, at least for the specific category of services offered by the recommender system.

Given this state of the art, as the influence of recommender systems expands, a highly critical question, at least from a social welfare point of view, emerges, namely "should the optimisation of user engagement supersede societal considerations and align fully with the objectives of the owner of the recommender system, or can society envisage recommender systems that can strike a balance between platform usage and the promotion of social welfare as broadly defined by a particular society? While this might seem to be a simple question to which there is a relatively simple answer, it is, in fact, anything but simple, being fraught with conflicting perspectives, diverging schools of thought, economic interests and cultural differences that make even broad definitions by one society differ from those of another. Indeed, one of the prevailing teleological schools of thought that are premised on a utilitarian-libertarian foundation would argue that social welfare is maximised when so-called "utility", made up, in the neoclassical economics tradition, of producer and consumer surplus, is at a maximum when recommender systems are solely designed to maximise user engagement and subsequently increase the profit margins for

NORMalize 2023: The First Workshop on the Normative Design and Evaluation of Recommender Systems, September 19, 2023, co-located with the ACM Conference on Recommender Systems 2023 (RecSys 2023), Singapore
bernard.mallia@equinoxadvisory.com (B. Mallia)

the platform owners inasmuch as users are free to choose whether to make use of such services or not. From this standpoint, the optimisation of user engagement, personalised user experiences and the commercial value generated, all serve as markers for the system's efficiency and success, in turn contributing to overall social welfare.

According to this neoclassical economics view, the cumulative benefits accrued by both the producers (here, the platform owners) and the consumers (the users of the platform) are the determining factor. It posits that when recommender systems efficiently provide users with relevant content, thus maximising their engagement and satisfaction, and in turn generate higher revenues for platform owners, a state of optimal utility and welfare is achieved.

However, in other contexts, this perspective has faced criticism for its narrow and reductionist approach to welfare, as well as its empirical invalidity [13]. Critics argue that it tends to overlook the nuanced socio-cultural dimensions of recommender systems and their implications. The maximisation of user engagement and profit, critics contend, can sometimes lead to the creation of digital echo chambers, addictive behaviour, user radicalisation, societal polarisation, user privacy threats and bias perpetuation.

Even if the utilitarian perspective had, despite its many flaws, to be taken to be the undisputed framework, there are still important ancillary questions that have arisen in neoclassical welfare economics theory that ponder whether the elusive concept of utility is subject to diminishing returns at the individual level [13]. If, as is customary in neoclassical economics, a money-metric measure of utility had to be taken, therefore, there would still be issues with the adoption of the utilitarian framework in arriving at a clear result of whether the state of the art of recommender systems yields optimal or sub-optimal social outcomes and what adaptations recommender systems would need to become more socially adequate, unless one is willing to make strong assumptions that end up determining the outcome, in which case one might as well assume an outcome directly.

An alternative deontological school of thought takes a diametrically opposed view, and hence calls for a broader, more inclusive understanding of welfare, one that also incorporates a set of desirable socio-ethical considerations [18]. Throughout this paper I shall call this the "Socio-Ethical Pragmatist Perspective". Advocates of this perspective[1] would posit that recommender systems should not just focus on maximising user engagement and operator profits, but also promote a diverse range of content, respect user privacy, provide equal or at least equitable representation for all users and work actively towards reducing societal polarisation. In essence, they argue for the harmonious balance between user engagement and social welfare, where social welfare is not merely defined in terms of maximising user engagement and profit. Instead, they propose a paradigm where recommender systems are designed and operated not just for the users and the platform owners, but for the amelioration of society as a whole, while taking the criteria of amelioration to be both given and commonly agreed to, even though this might not be – and might never be – the case in the absence of a political, Coasean-style bargain [5] being struck.

Several good attempts have already been made at coming up with ways of tweaking recommender systems to make them more useful in a socially-beneficial sense [22], [23]. In the latter, recommender systems were modified for various human values such as diversity, fairness, well-being, time well spent, and factual accuracy, while in the former a number of metric models were proposed for assessing and potentially tweaking the behaviour of news recommender systems. The aim of this paper is more foundational in nature as it tries to propose a socio-political framework through which to be able to premise and postulate 'appropriate' social outcomes.

A shift in recommender system design philosophy from engagement-centric to welfare-promoting recommender systems is a complex, but possible undertaking. It necessitates grappling with the intricate task of defining what constitutes "positive social outcomes". These

---

[1] In the realm of technology ethics, scholars like Helen Nissenbaum, who has written extensively on contextual integrity and privacy, could be considered to be aligned with such a perspective. Additionally, scholars focusing on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML) also advocate for ethical considerations in algorithmic systems, which would include recommender systems.

outcomes are a product of the sociocultural milieu, reflecting myriad factors including collective values, prevalent norms, societal goals, macroeconomic objectives and policies (such as those enshrined in EU competition policy or US antitrust legislation), as well as individual liberty and well-being. For instance, does a positive social outcome entail broadening user exposure to diverse viewpoints, fostering community cohesion, promoting accurate information, a combination of these or something else entirely? Navigating these nuances is crucial for the development of welfare-centric recommender systems.

2. Recommender Systems

Recommender Systems are an intricate subset of relational information retrieval, analysis and filtering systems. They play a pivotal role in tailoring the digital landscape to meet individual user needs as they meticulously curate and propose bespoke suggestions for a wide variety of items ranging from physical products to digital content. Proposals made by recommender systems are meticulously crafted, based on a user's historical preferences, other users' preferences, explicit and implicit interests, behavioural patterns either while using the platform or while using the Internet more broadly [2], as well as contextual information.

The functionality of recommender systems can be categorised into several established paradigms, each with their unique approaches to generating personalised recommendations [3]:

Content-based Filtering: This method operates by constructing a user-item interaction profile, which encapsulates the types of items that the same user has demonstrated a preference for in the past. It subsequently leverages this historical information to identify items that mirror those preferences. This paradigm incorporates a deep analysis of the user's historical behaviour to identify correlations to suggest items exhibiting similar attributes. For instance, if a user regularly consumes science fiction literature, the system may propose novels from the same genre or authors known for science fiction work.

Collaborative Filtering: The core philosophy behind this approach is predicated on the notion that users who have exhibited similar tastes and behaviours in the past are likely to share interests in the future. The system identifies patterns and correlations among a pool of users with similar statistical characteristics and makes item suggestions to the user based on the preferences and behaviours of similar users. For example, if a user often agrees with another user's movie reviews, the system might suggest movies that other similar users have watched or reviewed positively but that the user hasn't yet watched.

Hybrid Recommender Systems: These systems harmoniously amalgamate the strengths of content-based and collaborative filtering techniques to produce a robust and comprehensive recommendation engine that relies on both the user's own history and the history of similar users. Hybrid systems seek to exploit the synergies of the individual paradigms, with the intention of mitigating the limitations associated with each method when used in isolation.

In order to construct, fine-tune, and execute their proposals, recommender systems harness the power of an array of algorithms and methods derived from fields such as Machine Learning (ML), data mining and Natural Language Processing (NLP). The application of these sophisticated techniques enables the extraction and interpretation of patterns from large and complex big datasets, facilitating the generation of tailored recommendations at scale.

Moreover, recommender systems can be supplemented with additional information to enhance their accuracy and reliability. Examples of such augmentations include the incorporation of user ratings, text-based reviews and social media interactions that can be mined using sentiment analysis techniques, all of which can provide a richer and more nuanced understanding of a user's preferences and behaviour. Recommender systems today epitomise the effective

[2] This is in several instances illegal under the EU's GDPR but is still in use at least in some Recommender Systems.
[3] For further and more elaborate details please refer to Ricci, F., Rokach, L., & Shapira, B. (2022). Recommender Systems Handbook (3rd ed.). Springer. ISBN 978-1-0716-2196-7.

application of data-driven personalisation and remain a vital component in the realm of online user experiences, especially in a digital world increasingly characterised by information overload.

3. The Social Costs and Benefits of Recommender Systems

Recommender systems can and do raise ethical concerns, such as privacy, radicalisation and bias, and this is over and above the concerns that they raise in terms of digital services markets, competition policies and more generally the very democratic foundations on which modern Western societies are premised. Users may be concerned about the collection and use of their personal data, and the systems may inadvertently perpetuate biases and stereotypes if they are not designed with fairness and diversity in mind. Therefore, it is essential to design and implement recommender systems that are transparent, explainable and inclusive, and which do not radicalise user prejudices that are already present by consistently reinforcing the user's belief system.

Recommender systems are becoming increasingly influential in shaping the content that users are exposed to and the priority ordering with which such content is served, and in that sense they can effectively act as 'gatekeepers' [4]. From a social welfare point of view, it is thus crucial that they are designed in a way that promotes positive social and macroeconomic policy outcomes rather than trying to maximise the time the user spends using the platform on top of which the recommender system has been set up. This is, however, a much more complex question to deal with than first meets the eye and begs the fundamental question of what constitutes positive social outcomes worth pursuing and promoting. While with respect to existing regulations, such as those making up competition policy and digital services, this is quite easy to establish as the respective regulatory framework is already in existence, when it comes to the unregulated space, like the one for norms and values, it is everything other than easy to determine. To achieve this, it is important to first and foremost consider the norms and values that are appropriate for the domain in which the recommender system is operating. After that, it is also necessary to consider the norms and values of the social context of society as a whole. The latter inform (and are interlinked with) the former, which thus makes them, to some extent, interdependent. They also change across space and time such that:

(1) the norms and values in one geographical area at one point in time might not be the same for the same geographical area at another point in time; and

(2) norms and values might be different in different geographical areas at the same point in time.

Thus far, there are therefore two issues with norms and values. Firstly, that they can differ, sometimes drastically, between different content areas and geographical regions, and secondly that they can change over time. These can clearly be handled by underlying data structures but they also need to be updated either dynamically or at regular intervals if they are to remain valid. This also means that metrics for norms and values may not generalise across different domains or geographical areas, and may need to be tailored to the specific needs and values of each domain and geographical area, in addition to the existing categories.

Measuring norms and values can be a complex task, as they are often abstract and context-dependent despite there being a number of approaches that can be used in a recommender

---

[4] A gatekeeper, within the context of the framework of EU regulations, refers to large online platforms that have a significant influence over the digital single market due to their size, user base and control over data and access to their platforms. These platforms act as intermediaries between content creators and consumers and have the power to determine which content is visible and accessible to which users as well as the priority ordering at which such content is made visible. Due to this significant control, they can essentially 'gatekeep' the flow of information, services and products to significant proportions of users.
The concept of a gatekeeper has become particularly salient with the EU's Digital Markets Act (DMA). The DMA aims to ensure fair and open digital markets by addressing the challenges presented by gatekeeper platforms. Under the DMA, certain criteria determine if a platform is a gatekeeper, such as its annual turnover, active user base and entrenched position.
The role of gatekeepers in the context of recommender systems becomes critical as these systems inherently influence what users see, do and ultimately think. By controlling the algorithms of these systems, gatekeepers have a significant say in shaping user experiences, online choices and even social narratives. The European Commission's concern is that without proper regulations, these gatekeepers might abuse their position, leading to reduced competition, limited user choice and potential biases of multiple forms in the digital space.

system context. User surveys and feedback, for example, can be used to understand user preferences and values, while data analytics can be used to identify patterns and trends in user behaviour. Additionally, metrics such as diversity, novelty and serendipity can be used to measure the effectiveness of recommender systems in exposing users to a range of viewpoints and perspectives.

Designing experiments that measure norms and values can also be challenging, as it requires a careful balance between controlling for confounding factors and ensuring that the experiment is representative of real-world scenarios. One approach is to use split testing (also referred to as A/B testing), where different versions of a recommender system are tested with different user groups, and the impact on user behaviour and outcomes is measured. With the application of Artificial Intelligence (AI) algorithms, this can also be done dynamically and in an automated manner. Additionally, user studies and surveys can be used to gather qualitative feedback and insights from users, which can help to inform the design of recommender systems that align with an explicit statement of norms and values.

Social outcomes, particularly those associated with digital platforms, are thus inherently complex to define and measure. To navigate the multifaceted constructs of social outcomes associated with digital platforms, it is crucial to apply multiple ethical and philosophical lenses. This approach enables us to examine and gauge these outcomes from different viewpoints, ensuring that a wide range of values and norms are respected. Consequently, various ethical theories, each offering unique insights into what constitutes positive social outcomes, come into play.

Deontological ethics, teleological perspectives, consequentialism, virtue ethics, and social contract theory are among the ethical paradigms we can use to assess and subsequently shape the influence of recommender systems. Each of these theories offers a different perspective on the determination and quantification of social outcomes. They provide various yardsticks for evaluating the effects of these systems on individual users and of different societies. They enable us to understand whether the practices used by recommender systems align with the values and norms upheld by different societal and user groups, as well as how these practices contribute to broader social welfare.

Examining recommender systems through the lens of these ethical theories, we move from considering only the immediate goals such as maximising user engagement and platform profit, to a broader view of societal impact. This comprehensive examination allows us to better understand the ethical implications of recommender systems, which can then inform the design and implementation of such systems, guiding them towards not only optimising user experience and business profit, but also towards promoting positive social outcomes.

Whether we're considering the deontological perspective of adhering to principles such as data privacy and user consent, or the consequentialist viewpoint that emphasises the maximisation of positive outcomes for the greatest number of people, it becomes clear that each ethical framework can provide valuable insights into the operation of recommender systems. For instance, the teleological perspective can help us evaluate the effectiveness of these systems in promoting user satisfaction and enhancing knowledge, while social contract theory can help us understand how well digital platforms fulfil their implicit contract with users and society at large.

It is amply clear that understanding the potential and actual social outcomes of recommender systems requires a multidimensional approach that incorporates diverse ethical and philosophical perspectives. This approach can help to ensure that these systems are designed and implemented in a way that aligns with societal norms and values, while also contributing positively to social welfare. This is, again, a complex endeavour, but one that is also essential as recommender systems become increasingly integral to our digital experiences.

The Deontological Perspective: Deontological (duty-based) ethics, posit that certain actions are intrinsically right or wrong, regardless of their consequences [12]. From a deontological perspective, positive social outcomes can be determined by the adherence to ethical duties and principles, such as respect for autonomy, fairness and privacy. This approach is particularly relevant in the context of digital ethics, where principles like data privacy and user consent have been highlighted as important considerations [15]. By way of a practical example, a recommender system that respects user data privacy and curtails misinformation could be seen as promoting positive social outcomes, regardless of its impact on user engagement or platform profitability.

The Teleological Perspective: Teleological ethics, also known as consequentialism, judges the morality of an action based on its outcomes or ends [16]. This perspective defines a goal, and then goes on to define social good as the maximisation of that goal. From this perspective, social outcomes could be measured in terms of their impact on societal or individual well-being. Recommender systems, for instance, could be evaluated based on their effectiveness in promoting user satisfaction, enhancing knowledge or reducing information overload. This approach, however, poses challenges in defining and measuring well-being, as well as balancing the well-being of different stakeholders because of its reliance on the utilitarian framework alluded to in the introduction.

Consequentialism: In addition to the teleological approach, consequentialist ethics can also be applied in the area of recommender systems. Here, the moral rightness of an action is based on the maximisation of positive outcomes for the greatest number of individuals [3]. This approach might consider factors such as the broad accessibility and utility of a recommender system.

Virtue Ethics: Virtue ethics emphasise the development of moral character and the embodiment of virtues. A virtue-ethical perspective [2] might focus on the cultivation of intellectual virtues, such as critical thinking and openness to diverse viewpoints, in determining and measuring social outcomes.

Social Contract Theory: The social contract theory, originally articulated by thinkers such as Thomas Hobbes [10], John Locke [15] and Jean-Jacques Rousseau [20], refers to an implicit agreement among the members of a society to cooperate for social benefits. In essence, it denotes the understanding that individual self-interest and societal wellbeing are interdependent and, to a certain extent, the health of the latter shapes the potential of the former.

In the context of recommender systems, the social contract would comprise an implicit agreement between the platform providers and their systems on the one hand, and the users on the other. On the users' end, they provide their data and attention, and in return, they expect a system that serves their information needs, respects their privacy, upholds fairness, and contributes to their overall wellbeing.

This social contract also extends to the broader society beyond the individual users. Platforms, through their recommender systems, have a responsibility not to propagate harmful content or behaviours, foster polarisation, or intensify societal divisions. Furthermore, they have a social responsibility to promote diverse content, prevent the amplification of extreme or harmful viewpoints, and counteract the formation of “filter bubbles” [5] and “echo chambers” [6].

From a social contract perspective, recommender systems’ positive social outcomes could be those that uphold this theoretical contract. It would mean creating systems that do not exploit user data, are designed to avoid undue influence or manipulation, and aim to provide a broad and diverse range of information, thus contributing to a well-informed public. It could also involve systems that actively promote social cohesion, foster constructive discourse and uphold human-rights-based and democratic values [7]. The challenge, however, is in the operationalisation of this perspective as the precise terms of the social contract can be challenging to define – even with respect to a single society – and also differ across various cultural and societal contexts.

---

[5] A filter bubble is a state of intellectual isolation that can occur when websites use algorithms to selectively serve up content to users based on information about them, such as their location, past click-behaviour, and search history. This leads to the users being separated from information that disagrees with their viewpoints or that is outside of their interest areas, effectively isolating them in a “bubble” of reinforcing content.

[6] An echo chamber is a situation in which an individual is exposed primarily to opinions and beliefs similar to their own, without much exposure to differing viewpoints. This can occur on digital platforms when users, either through their own choices or the platform's recommender system algorithms, mainly interact with like-minded individuals or consume content that aligns with their preexisting beliefs. As a result, their ideas and beliefs get amplified or echoed back to them, often leading to reinforcement of pre-existing views, radicalisation and polarisation.

[7] This being, of course, a normative value deriving from an author who comes from a democratic country and who values democracy. A Chinese author might very well have placed autocratic, regime-preserving values here, and that is perfectly fine insofar as that would represent the state of norms and values in China.

To fully grasp the potential and actual social outcomes of recommender systems, it is paramount to adopt a multidimensional approach that incorporates various ethical and philosophical perspectives. The deontological, teleological, consequentialist, virtue-ethical, and social contract viewpoints are not mutually-exclusive and can contribute distinct layers of understanding in evaluating these systems.

Despite the complexities involved in defining and measuring social outcomes, it is an endeavour worth undertaking, as only by considering the full spectrum of ethical perspectives can we develop systems that truly serve the users and society at large, rather than simply focusing on immediate goals like user engagement and platform profitability.

In pursuing this goal, in a democratic, pluralistic society, answering these questions should not rest on a select few, or a specific class of experts, stakeholders or institutional gatekeepers. Instead, it should be a collective, collaborative endeavour that involves a broad array of stakeholders, including but not limited to platform owners, users, content creators, academics, policymakers, social activists, civil society and ethicists. Together, these stakeholders can engage in a participatory dialogue, in which a pluralism of views is respected, valued and incorporated into the decision-making process to shape the direction of recommender systems and ensure they contribute positively to our societies, reinforcing a digital social contract that respects user autonomy, fosters diversity and inclusion and ultimately serves the public good. In this way, recommender systems can also become trusted aids in our exploration of the digital world, rather than forces that unduly influence our experiences and decisions in pursuit of some hidden profit-maximising formula that fails to take any consider of unintended consequences resulting in the accretion of social externalities that have to be borne by society. Only after having gone through this process can we really claim to be able to have the blueprint for a class of recommender systems that is aligned with social objectives. This is a process that entails not only a thoughtful articulation of values and norms, but also a collective negotiation of the underlying assumptions, uncertainties and trade-offs.

By way of two practical examples, in the case of news recommender systems, it might be important to ensure that users are exposed to a broad range of viewpoints and perspectives. This can promote a more informed and nuanced understanding of the world, and can help combat the echo chambers and filter bubbles that can emerge when users are only exposed to content that confirms and at times radicalises their existing beliefs. To achieve this, news recommender systems can use algorithms that prioritise diversity and expose users to a range of viewpoints and perspectives, while at the same time eliminating misinformation and fake news.

In the case of travel recommender systems, sustainability aspects could be taken into account. This can involve considering factors such as the carbon footprint of different travel options, the impact of tourism on local communities and ecosystems, and the availability of sustainable transportation options. By considering these factors, travel recommender systems can help to promote more sustainable and responsible travel practices and to create a market for sustainable options in tourism, thereby aligning more closely with society’s goals.

4. Socially-Beneficial Recommender System Implementation

After designating socially-beneficial goals, the implementation of a recommender system entails a straightforward (even if technically burdensome) multistep process involving data collection and processing, model building and refinement, and system evaluation and optimisation. Each step can be shaped by the designated social goals to ensure that the recommender system promotes only the intended outcomes.

Data Collection and Processing: First, data collection is a fundamental step in building any recommender system [1], including one that is aligned with a set of social objectives. The type of data gathered typically depends on the application domain and the type of recommender system being built. For example, collaborative filtering systems require data on user-item interactions, while content-based systems require item attribute data [19]. If a socially-beneficial goal is to promote diversity, data on a wide variety of items and user preferences would need to be collected. In collecting data, considerations should also be made regarding user privacy and consent / legitimate business interest premises in line with data protection regulations [6].

Moreover, it would be essential to preprocess data to ensure it is of a good-enough quality and fit for modelling. This may involve handling missing values, eliminating noise and transforming variables. The preprocessing stage also allows for the embedding of social considerations by ensuring fairness and avoiding biases in the data that could lead to discriminatory recommendations [7].

Model Building and Refinement: The next step is the development of the recommendation algorithm. Different types of algorithms, such as collaborative filtering, content-based filtering, and hybrid methods, are suitable for different applications and objectives [4]. Socially-beneficial goals can guide the selection or modification of algorithms. For instance, if the goal is to promote exposure to diverse information, techniques that encourage diversity in recommendations, such as reranking or item-based collaborative filtering, could be used [23].

The model building stage also involves the selection of evaluation metrics. Traditional metrics like precision and recall may need to be complemented by others that align with social goals. For example, if the goal is to increase diversity, metrics like the Gini index, Intra-list Similarity, Coverage, Novelty, Unexpectedness, the Shannon Diversity Index, the Herfindahl-Hirschman Index (HHI), entropy or a composite index of the foregoing measures could be used. However, conflicts may arise between traditional recommender system performance metrics like precision and recall, and the social objectives and these need to be dealt with [8].

System Evaluation and Optimisation: Finally, the system needs to be evaluated and optimised. This involves testing the system on real or simulated data, assessing its performance, and making the necessary improvements. Socially-beneficial goals can guide the evaluation process by focusing on metrics that reflect these goals. User studies and split testing can also be used to measure the impact of the recommender system on users and evaluate whether it is achieving the intended social outcomes [14].

It is also important to continuously monitor and update the system to ensure its performance over time, considering the dynamic nature of user preferences and the digital environment [8].

5. Conclusion

A shift in design philosophy from engagement-centric to welfare-promoting recommender systems is a very complex but socially-desirable undertaking. It necessitates grappling with the intricate task of defining what constitutes “positive social outcomes”. These outcomes are a product of the sociocultural milieu, reflecting myriad factors including collective values, prevalent norms, social goals, and individual and societal well-being. What does a positive social outcome entail in practice, and is self-regulation enough to attain such an outcome? Navigating these nuances is crucial for the development of welfare-centric recommender systems.

Recommender systems operate across a plethora of domains and geographical areas, each carrying unique norms, values and expectations. As such, the integration of these contextual nuances within the design and operation of recommender systems is paramount. However, the translation of abstract societal norms and values into concrete algorithmic criteria is fraught with challenges. It requires a nuanced understanding of the domain, a careful consideration of ethical implications, and the ability to navigate potential trade-offs between conflicting norms and values. Moreover, the dynamic nature of social norms implies that recommender systems must be adaptable and capable of evolving with societal changes. This process involves not just ethical theorisation, but also empirical investigation, consultation with stakeholders, as well as continuous iterative refinement. Ultimately, the integration of these norms and values into

[8] By way of illustration, maximising precision might lead to highly similar recommendations, contradicting a goal like diversity. A trade-off strategy is necessary here and can take the form of multiplicative integration of precision with diversity, to balance these potentially conflicting goals. Similarly, increasing recall might conflict with societal goals like misinformation mitigation as broadening the scope of recommendations could inadvertently amplify misleading content. One potential solution could be to integrate a misinformation detection module in the recommender system, ensuring the wider recommendations maintain a high standard of information quality.

recommender systems requires rigorous methodologies, spanning fields from machine learning and data science to philosophy, ethics and various branches of social sciences like sociology, economics, anthropology and political science.

As recommender systems continue to evolve, these considerations are instrumental in ensuring they serve as a force for positive social impact, rather than vehicles for platform engagement and profit maximisation. Only by using a range of approaches to measure norms and values, and by designing experiments that are representative of real-world scenarios, can recommender systems that have a positive impact on society be envisaged and programmed into being.

References

[1] Adomavicius, G., Tuzhilin, A., Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions, IEEE Transactions on Knowledge and Data Engineering 17(6) (2005) 734-749.

[2] Aristotle, Nicomachean Ethics, 350 BCE.

[3] Bentham, J., An Introduction to the Principles of Morals and Legislation, 1789.

[4] Burke, R., Hybrid recommender systems: Survey and experiments, User Modeling and User-Adapted Interaction 12(4) (2002) 331-370.

[5] Coase, R. H., The Problem of Social Cost, Journal of Law and Economics 3 (1960) 1-44.

[6] Custers, B., Hof, S., Schermer, B., Appleby-Arnold, S., Brockdorff, N., Informed consent in social media use – The gap between user expectations and EU personal data protection law, SCRIPTed 15(4) (2018) 435-468.

[7] Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R., Fairness through awareness, Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (2012) 214-226.

[8] Ekstrand, M. D., Kluver, D., Harper, F. M., Konstan, J. A., Letting users choose recommender algorithms: An experimental study, ACM Transactions on Interactive Intelligent Systems (TiiS) 8(3) (2018) 1-26.

[9] Hildebrandt, M., The Issue of Proxies and Choice Architectures. Why EU Law Matters for Recommender Systems, Frontiers in Artificial Intelligence 5 (2022). DOI: 10.3389/frai.2022.789076.

[10] Hobbes, T., Leviathan, 1651.

[11] Kaminskas, M., Bridge, D., Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems, ACM Transactions on Interactive Intelligent Systems (TiiS) 7(1) (2016) 1-42.

[12] Kant, I., Groundwork for the Metaphysics of Morals, 1785.

[13] Keen, S., Debunking Economics - Revised and Expanded Edition: The Naked Emperor Dethroned?, Paperback Edition, 2011.

[14] Kohavi, R., Henne, R. M., Sommerfield, D., Practical Guide to Controlled Experiments on the Web: Listen to Your Customers not to the HiPPO, Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2009) 959-967.

[15] Locke, J., Two Treatises of Government, 1689.

[16] Mill, J.S., Utilitarianism, 1863.

[17] Moor, J.H., "What is Computer Ethics?" Metaphilosophy 16(4) (1985) 266-275.

[18] Rawls, J., A Theory of Justice, Harvard University Press, 1971.

[19] Ricci, F., Rokach, L., Shapira, B., Introduction to Recommender Systems Handbook, In Recommender Systems Handbook (2011) 1-35.

[20] Rousseau, J.J., The Social Contract, 1762.

[21] Stray, J., Vendrov, I., Nixon, J., Adler, S., Hadfield-Menell, D., What are you optimizing for? Aligning Recommender Systems with Human Values, Journal Name, Volume(Issue) (2021), arXiv:2107.10939.

[22] Vrijenhoek, S., Kaya, M., Metoui, N., Möller, J., Odijk, D., Helberger, N., Recommenders with a Mission: Assessing Diversity in News Recommendations, In CHIIR '21: Proceedings of the 2021 Conference on Human Information Interaction and Retrieval, March 14–19, Canberra, Australia (2021).
[23] Ziegler, C. N., McNee, S. M., Konstan, J. A., Lausen, G., Improving recommendation lists through topic diversification, Proceedings of the 14th International Conference on World Wide Web (2005) 22-32.