

Categorizing Algorithmic Recommendations: A Matter of System, Agent, Patient*

Matteo Fabbri¹, Jesus Salgado²

¹ IMT School for Advanced Studies, Lucca, Italy

² Universidad Politécnica de Madrid, Madrid, Spain

Abstract

In the contemporary digital age, recommender systems (RSs) shape the way in which people interact online and offline: from social media to music streaming, from e-commerce to news websites, suggested contents and products have the spotlight on platforms' interfaces and influence individuals' interests and priorities. RSs have recently been addressed by European regulations such as the Digital Services Act, whose impact on the design and management of online platforms can already be observed. While algorithmic recommendations, as the output of RSs, are aimed at improving user's experience by reducing the information overload, they can give rise to ethical concerns related to privacy, autonomy and fairness, and generate risks such as misinformation, filter bubbles and epistemic fragmentation. RSs have even been featured in legal cases involving the endangerment of minors through social media challenges and the recruitment of terrorists: this evidence underlines their deep impact on society. However, the concept of recommendation lacks a unified understanding due to the variety of domains in which the corresponding term is used. In fact, if the context of use is not specified, what is referred to as a recommendation includes not only the output of RSs, which may influence users without constraining their freedom, but also the outcomes of decision support systems (DSSs) or automated decision-making systems (ADMSs), whose impact on individuals is direct and often does not depend on their choice. This contribution aims to propose a framework for the ontological differentiation between the different concepts of algorithmic recommendation. The differentiation is based on the identification of the subject who has the responsibility and autonomy to decide whether to follow the recommendation. The adoption of this framework has the potential to improve the ethical scrutiny and auditing of AI technologies, which are required by European regulations like the DSA, for what concerns RSs, and the AI Act, regarding DSSs and ADMSs.

Keywords

Algorithmic Recommendations, Ontological Differentiation, AI Regulation

1. Introduction

In the contemporary digital age, recommender systems (RSs) shape the way in which people interact online and offline: from social media to music streaming, from e-commerce to news websites, suggested contents and products have the spotlight on platforms' interfaces and influence individuals' interests and priorities. Because of the risks related to their nudging potential, RSs deployed by online platforms are now subject to the transparency requirements of the Digital Services Act [1], whose impact on the design and management of digital environments can already be observed [2]. While algorithmic recommendations, as the output of RSs, are aimed at improving user's experience by reducing the information overload, they can give rise to ethical concerns related to privacy, autonomy and fairness [3], and generate risks such as misinformation, filter bubbles and epistemic fragmentation [4]. RSs have even been featured in legal cases involving the endangerment of minors through social media challenges [5] and the recruitment of terrorists [6]: this evidence underlines their deep impact on society.

EWAF'24: European Workshop on Algorithmic Fairness. July 1-3, 2024. Mainz, Germany.

✉ matteo.fabbri@sns.it (M. Fabbri); jesus.salgado@upm.es (J. Salgado)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

However, the concept of recommendation lacks a unified understanding due to the variety of domains in which the corresponding term is used. In fact, if the context of use is not specified, what is referred to as a recommendation includes not only the output of RSs, which may influence users without constraining their freedom, but also the outcomes of decision support systems (DSSs) or automated decision-making systems (ADMSs)², whose impact on individuals is direct and often does not depend on their choice. The conceptual boundary between RSs and DSSs has not been clearly established, considering that “there is still no accepted definition of DSS” [8] in computer science. [9] observes that, whilst DSSs are “devoted to performing a content-specific task that supports human decision making (although human decisions often tend to be determined rather than supported by it)”, RSs “are not content- but context-specific: the content of their output can vary widely depending on the user, but they are directed by a defined aim within a particular context, i.e. maximizing user engagement in a social media platform”. Following this argument, if a recommendation always falls under a specific topic within a wider domain (e.g., personalized therapy for lung cancer), then it should represent the output of a DSS. Otherwise, if a recommendation deals with various topics in the same domain (e.g., miscellaneous daily news based on a user’s profile), then it can be considered the output of a RS.

However, this argument does not provide a defined boundary that allows to distinguish precisely, from the recipient’s perspective, whether a recommendation is the output of DSSs or RSs, as it does not clarify whether the person that directly faces the implications of the decision can choose whether to follow the suggestion of the system. For example, in the healthcare domain, a recommendation about keeping the appropriate heartbeat will have very different implications if it is produced by a runner’s wearable device or by a Holter monitor worn by a patient under anaesthesia in the operating room: although the content of the recommendation is the same, in the former case it is “consumed” by the person directly concerned by it (i.e. the runner whose heartbeat is being measured), while in the latter case it is “consumed” by a third person who decides whether it will impact the person directly involved (i.e. the surgeon).

This example highlights a guiding question for the ontological differentiation between the different concepts of algorithmic recommendation: who has the responsibility and autonomy to decide whether to follow the recommendation? To attempt an answer, we consider three subjects: the system (S), the agent (A) and the patient (P): S is the technology produces the recommendation; A evaluates the recommendation and decides whether to follow it; P directly bears the consequences of following the recommendation. The relationship between these subjects can determine a taxonomy that allows to distinguish between RSs, DSSs and ADMSs:

- If $A = P$ and $S \neq A$, we have a RS. The subject who receives a recommendation from the AI system and can choose whether to follow it is the same who directly bears the consequences of following it. Therefore, the recommendation can influence but cannot constrain the choices of the subject who is exposed to it (e.g.: a YouTube user sees a list of recommended videos, clicks on one of them and watches it).
- If $S \neq A$, $S \neq P$ and $A \neq P$, we have a DSS. The subject who receives a recommendation from the AI system decides which impact it will have on another person who is not actively involved in the decision-making process but bears its consequences. This is typical of domains where a specific expertise is required, like medicine or law (e.g.: a judge decides for how many years a culprit should be convicted based on his recidivism score).
- If $S = A$ and $A \neq P$, we have an ADMS. The recommendation coming from the AI system is directly enforced on the subject who must bear its consequences without any human-in-the-

² In fact, the AI Act [7] lists recommendations as a type of output of an AI system, alongside “content”, “predictions” and “decisions” (art. 3.1).

loop intervention. The recommendation *de facto* becomes an automated decision (e.g.: in the UK in 2020, when final high-school exams could not be taken due to the pandemic, A-level grades, which determine admission to university, were assigned by the so-called Ofqual algorithm without any mediation by teachers or schools [10]; as the algorithm turned out to be biased, the artificially estimated results were not considered).

The system-agent-patient (SAP) framework would contribute to establishing whether an algorithmic recommendation comes from a RS, DSS or ADMS from the perspective of its human recipient, thereby bringing conceptual clarity on the distribution of responsibility for the output of these AI systems, each of which has different implications on society. In fact, while RSs influence individuals indirectly through nudging strategies [11], DSSs and, even more so, ADMSs constrain the autonomy and freedom of the subjects who bear the consequences of following the recommendation but are not responsible for choosing whether to follow it. Therefore, the ontological differentiation based on the SAP framework has the potential to improve the ethical scrutiny and auditing of AI technologies, which are required by European regulations such as the DSA, for what concerns RSs, and the AI Act, as regards DSSs and ADMSs.

References

- [1] REGULATION (EU) 2022/2065 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act). URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32022R2065>
- [2] TikTok Newsroom, An update on fulfilling our commitments under the Digital Services Act, 2023. URL: <https://newsroom.tiktok.com/en-eu/fulfilling-commitments-dsa-update>
- [3] S. Milano, M. Taddeo, L. Floridi, Recommender systems and their ethical challenges, *AI & Society*, 35(4), 2020, pp. 957-967.
- [4] S. Milano, B. Mittelstadt, S. Wachter, C. Russell, Epistemic fragmentation poses a threat to the governance of online targeting. *Nature Machine Intelligence*, 3(6), 2021, pp. 466-472.
- [5] J. Edwards, Mother sues TikTok after 10-year-old died trying “Blackout Challenge”, *The Washington Post*, 2022. URL: <https://www.washingtonpost.com/nation/2022/05/17/tiktok-blackout-challenge-lawsuit/>
- [6] Centre for Democracy and Technology, CDT and Technologists File SCOTUS Brief Urging Court To Hold that Section 230 Applies to Recommendations of Content, 2023. URL: <https://cdt.org/insights/cdt-and-technologists-file-scotus-brief-urging-court-to-hold-that-section-230-applies-to-recommendations-of-content/>
- [7] Proposal for a regulation of the European Parliament and of the Council laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021PC0206&from=EN>
- [8] M. C. Er, Decision support systems: a summary, problems, and future trends, *Decision support systems*, 4(3), 1988, pp. 355-363.
- [9] M. Fabbri, Self-determination through explanation: an ethical perspective on the implementation of the transparency requirements for recommender systems set by the Digital Services Act of the European Union, *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 2023, pp. 653-661. URL: <https://doi.org/10.1145/3600211.3604717>
- [10] D. Kolkman, “F**k the algorithm”? What the world can learn from the UK’s A-level grading fiasco. *LSE Blogs*, 2020. URL: <https://blogs.lse.ac.uk/impactofsocialsciences/2020/08/26/fk-the-algorithm-what-the-world-can-learn-from-the-uks-a-level-grading-fiasco/>

- [11] M. Jesse, D. Jannach, Digital nudging with recommender systems: Survey and future directions. *Computers in Human Behavior Reports*, 3, 100052, 2021. URL: <https://doi.org/10.1016/j.chbr.2020.100052>