

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/115977>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

A Vision for Socially Incentivised Recommendations

Phillip Taylor¹, Lina Barakat², Nathan Griffiths¹, and Simon Miles³

¹ The University of Warwick, UK

phillip.taylor@warwick.ac.uk, nathan.griffiths@warwick.ac.uk

² University of Essex, UK

lina.barakat@essex.ac.uk

³ King's College London, UK

simon.miles@kcl.ac.uk

Abstract. Typically, recommender systems focus solely on individual preferences of users or small groups of users, but recommendations can have effects on the wider social structure. Social considerations are therefore necessary in recommendation generation. In this paper, we identify gaps in literature relevant to socially responsible recommendation systems, and present a number of challenges. Finally, we present a vision for an architecture capable of generating socially responsible recommendations and encouraging their acceptance via incentives and rationales.

Keywords: Recommender systems, Norms, Incentivisation, Fairness Accountability and Transparency

1 Introduction

Existing research in recommender systems has typically focussed on the preferences of individuals and small groups of users, and do not account for social constraints and policies at other levels within the community that are essential to improve social welfare. For example, in a transportation recommender system, a user may own both a bicycle and a car, but have particular goals, preferences, or constraints for a journey. They may need to arrive within 30 minutes (a hard constraint), but they may also prefer to minimise physical activity (a soft constraint). If both cycling and driving enables the user to arrive within the time limit, driving is the best choice when considering individual constraints and preferences. From a wider perspective, however, the destination neighbourhood may have an issue with noise pollution, leading to an emergent norm of reducing car use after certain hours. Moreover, the local authority may aim to improve air quality, and discourage car use. Accounting for such goals of the wider community would entail recommending cycling.

A community may have multiple layers (e.g. national and local governments) where members in one layer can overrule decisions made in another, causing complex relationships with conflicting interests and norms of differing granularities. For example, one locality may benefit by encouraging traffic to areas

outside it, but to the detriment of others. Another challenge is encouraging recipients to adopt such socially beneficial recommendations, which may conflict with their individual preferences. In this paper we elaborate on these challenges, and present a vision for an architecture capable of producing socially responsible recommendations, which comply with multi-level rules and norms, and also encourage their adoption.

The remainder of this paper is as follows. Section 2 provides a gap analysis of the literature, leading to the challenges outlined in Section 3. The proposed architecture is presented in Section 4, and Section 5 concludes this paper.

2 Literature Review and Gap Analysis

In this section, we discuss selected state-of-the-art in areas related to socially responsible recommender systems, namely aims, constraints, and norms in Section 2.1, recommender systems in Section 2.2, and incentivisation in Section 2.3. In each case we identify gaps to be filled to achieve a socially responsible recommender system.

2.1 Social aims and constraints

Constraints and accepted behaviour patterns within a social structure are referred to as norms. They help regulate societies, guide the behaviour of individuals, and facilitate coordination and interactions among them [5]. Recommendations should typically comply with existing norms, unless the aim is to modify such norms. Norms can be classified into obligations, prohibitions, social commitments, and social codes [26]. Obligations and prohibitions are created by authorised agents (not the addressees) in a society, and are usually met with punishment if violated. By contrast, social commitments are created as a result of agreements between addressee agents, and may disappear after their fulfilment and their fulfilment is often rewarded. Finally, social codes are general societal principles that have neither rewards nor punishments, but compliance rests on empathy or sympathy towards norm beneficiaries. For example, recycling typically has no direct reward or punishment associated with it, but a social code may lead to widespread participation [2].

To formally represent norms and reason about them, logics and formal languages can be used [23, 8, 26]. Formal representation enables agents to reason about which norms to comply with, by considering normative goals (what should be done or avoided), addressees (who should comply), beneficiaries (who benefits), context (where the norm applicable), exceptions (when non-compliance is not punished), rewards (for compliance), and punishments (for non-compliance). Frameworks such as that developed by King et al. [10] enable institutions in one level of a hierarchy (e.g. national government) to check norm compliance at a lower level (e.g. local council).

While norms exist to benefit a society, autonomous entities may choose to ignore them if they conflict with their personal goals. To encourage compliance, norm enforcement mechanisms can be used, for example, using education,

promotion by influential peers, or introducing punishments for compliance and rewards for non-compliance (administered by either peers or dedicated agents). Posner and Rasmusen [1] discuss different types of sanctions that can help enforce norms, including automatic sanctions (where reduced coordination cause harm to the violator), informational sanctions (where undesirable characteristics of the violator that may damage reputation are revealed), bilateral or multilateral costly sanction (where a punishment is applied by other members of the society), guilt, and shame. With respect to punishment, adaptive mechanisms have been studied in computational multi-agent systems and proved to be effective, including both a simple escalating punishment model [6], and more sophisticated dynamic punishment models [13, 24], where the level of punishment is increased or decreased in response to agent reaction to that punishment.

In the context of a socially responsible recommender system, various norms (e.g. societal obligations and conventions), soft and hard constraints, and preferences of various stakeholders, at multiple competing levels of granularity, need to be considered. With regards to this, the gaps identified are as follows.

- How to formally represent and reason about norms in a multi-level system with conflicting goals across all levels?
- How to ensure multi-level norm enforcement using influential agents, adaptive punishments and rewards?

2.2 Social responsibility in recommendation generation

Recommender systems are used in several domains [12] [18], and are typically based on content (what a user has liked previously), collaborative filtering (users with similar preferences), demography (clusters of similar users), knowledge (specific needs of a user), community (a user’s friends), or a hybrid of these [18]. In general, recommender systems aim to increase the number of interactions (for example sales in commerce), or increase user satisfaction. With the exception of knowledge-based systems and other contextual recommendation systems, there is little research on generating recommendations within constraints, and in particular recommendations that aim to achieve social goals.

Group recommender systems, which generate recommendations to satisfy a group of users rather than individuals [14], have similarities with social recommendation generation. In generating recommendations for a group of users, the system must consider all their preferences to either maximize fairness or minimize overall misery [20] [3] [25], which is a kind of social aim. However, group recommender systems do not consider individuals outside the group but in the same society, thus ignoring the wider values. Moreover, in the case of generating recommendations that aim to achieve social goals, interactions between recommendations and their side effects must also be considered. In traffic management, for example, recommendations that motivate users to all take the same route home may interact and cause a traffic jam. Redirecting some of this traffic to reduce congestion in one area may inconvenience some individuals to the benefit of the majority, but this diversion may have a side effect and conflict with

a higher level goal of reducing the total number of cars on the road (i.e. lower congestion levels may encourage more people to drive). A socially responsible recommender system should consider as many interactions and side effects of recommendations as possible.

Another aspect of a socially responsible recommender is system accountability, for example when giving recommendations that may affect an individual's security [12] [21]. Similarly, is it responsible to recommend users' content similar to what they have previously read, potentially creating echo chambers? Nguyen et al. [15] suggest that the creation of echo chambers is natural, and recommender systems based on collaborative filtering may reduce their likelihood, but this is unlikely to be sufficient in all domains and in the long term. It is also important to encourage users to follow recommendations (see Section 2.3) as well as increasing recommendation diversity [15]. An accountable recommender system should go further than this, however, and provide some guarantees that its recommendations will assist in achieving a social aim. Recommendations should be auditable to identify how and why they were generated, and justifications should be presented to users detailing intended and possible unintended outcomes [11].

Gaps identified regarding social responsible recommendations are as follows.

- How to account for multi-level hard and soft constraints, with multiple stakeholders?
- How to reason about multi-level side effects of concurrent and consecutive recommendations?
- How to account for, audit, and justify recommendations?

2.3 Encouraging acceptance of recommendations

Besides recommendation accuracy, user interaction aspects are important to improve effectiveness and acceptance of recommendations [9]. Transparency and persuasion are particularly important factors for user acceptance, and *explanation* techniques have been proposed to contribute to both, via exposing the reasoning behind a recommendation [22] [16], e.g. 'use X because ...'. Transparency and persuasion are of vital importance in socially responsible recommendations, but become more challenging for a number of reasons. In particular, an explanation should communicate the logic and benefits behind a recommendation, while reflecting the complex, possibly conflicting, social constraints and policies at various hierarchical levels. Moreover, since the recommendation may appear sub-optimal from the perspective of an individual (but beneficial socially), further encouragement for adoption beyond conventional explanations are also required.

Various economic and psychological studies emphasise the importance of *incentives* in directing an entity towards a desired behaviour [4] [7]. Such studies reveal that there is typically a mixture of motives driving an entity to undertake a particular task, which could be intrinsic (e.g. out of personal interest in the task) or extrinsic (e.g. to gain financial reward or social approval), and may differ among entities. Targeting such motives with relevant incentives would thus

allow pushing the entity’s behaviour in the desired direction. The applicability of such incentivisation in computational systems has been investigated for multi-agent systems, where researchers have considered the effect of punishment (in the form of fines) [13], credibility scores [17], violation alerts [24], deferred reciprocity [19], and reputation [27]. We believe that investigating similar incentives, and their effect on user behaviour, in the context of socially responsible recommender systems is a promising direction to improve user acceptance.

Gaps identified regarding encouraging recommendations are as follows.

- How to generate transparent rationale or explanation for recommendations accounting for complex multi-level social goals and constraints?
- What are appropriate intrinsic and extrinsic incentives to influence behaviour?
- How to track recommendation acceptance in complex environments?

3 Challenges

The overarching challenge in creating a socially aware recommender system is to combine the diverse state-of-the-art for social norms, recommender systems, and incentivisation, through the lens of fairness, accountability, and transparency. In addition to addressing the gaps identified above, there are four challenges that must be overcome.

Social goals, norms, and preferences: A socially responsible system requires social aims and constraints to be specified. Moreover, because multiple parties may have a social stake in the system, the aims must be structured to reflect this and represent their differing perspectives, personal goals, and preferences. In particular, one community may depend on another, e.g. the services provided in a city comprise the services provided in its boroughs, and the social goals at the city level may compete with those of individual boroughs. The challenge is therefore to define a language for specifying a socially responsible system in terms of soft and hard constraints, preferences and norms at multiple competing levels of granularity.

Socially responsible recommendation mechanism: Actions in the best interests of an individual may conflict with social aims, system constraints, or the interests of others. Recommendations considering only the personal interests of an agent, rather than multi-level system aims and constraints, are therefore not socially responsible. The challenge is to develop a recommender system that processes the goals, constraints, and preferences to deliver recommendations aimed at maximising acceptability across all levels. Algorithms for identifying and reasoning with social constraints must be identified and developed, to ensure the generation of socially valid recommendations that satisfy constraints. A ranking mechanism is then required, to prioritise the recommendations with the greatest chances of acceptance and successful implementation. Thus, measures of acceptance and success are required, both for recommendation ranking and system evaluation.




Fig. 1. A preliminary architecture for socially responsible recommendations

Incentivisation generation Self-interested entities may ignore socially beneficial recommendations if they conflict with their individual goals and appear to be sub-optimal. The challenge is then to encourage acceptance of choices that yield social benefits. To achieve this, it is important to analyse the requirements of stakeholders, to elicit their goals and perspectives, and generate relevant incentives. Moreover, since motivational needs may vary with time, context, and across individuals, the incentive generator must be equipped with a learning component and provide personalisation and adaptation to user behaviours and preferences. For example, while some users may be motivated by immediate financial incentives (e.g. vouchers), others may respond more to how their behaviour affects others (e.g. traffic congestion).

Rationale generation To increase the likelihood of individually sub-optimal recommendations being accepted, it is useful to provide their rationales alongside other incentives. Rationales should consist of a description of the model and how its output accounts for context, preferences, and social constraints. In addition, rationales should illustrate the intended outcome of the recommendation, and highlight possible individual and social benefits and drawbacks. The challenge therefore is to aid incentivisation by generating rationales that go further than simply explaining model outputs. The rationale engine must process social aims and constraints, and describe how adoption of the recommendation can improve society.

4 Proposed Architecture

We envision a socially responsible recommendation framework, as outlined in Figure 1, which includes a component for each challenge identified in Section 3. Instantiations of components are to use existing techniques supplemented by new methods to address the gaps identified in Section 2. Individuals and groups, with possibly conflicting preferences, goals, and norms at different granularities, generate a set of hard and soft constraints, social aims, and personal preferences. The socially responsible recommendation engine will consider all of these factors

by analysing historical feedback and context data to provide recommendations that are socially beneficial. To encourage adoption of recommendations, they are accompanied by generated rationales and incentives. Both recommendations and incentives must be personalised to the changing preferences and contexts of individuals, thus, this framework allows learning and reasoning about the effect of recommendations, rationales, and incentives on a individuals' behaviour.

5 Conclusion

We have considered the state-of-the-art of social norms, recommender systems, and incentivisation in the context of socially responsible recommendation. We then identified several gaps that must be filled in order to achieve a socially responsible recommender system, notably the consideration of complex multi-level hard and soft social constraints and potentially conflicting aims. We then outlined four challenges, namely (i) defining social goals, norms, and preferences, (ii) generating socially responsible recommendations, (iii) incentivisation, and (iv) rationale generation. Finally, we proposed an architecture vision for a socially responsible recommender system, which models a system on multiple levels and generates recommendations, along with incentives and rationales.

References

1. R. A Posner, E. Rasmusen, R. Ellickson, R. Mcadams, M. McGinnis, and M. Ramseyer. Creating and enforcing norms, with special reference to sanctions. *Law and Economics 9907004*, *EconWPA*, pages 435–5806, 08 1999.
2. A. Abbott, S. Nandeibam, and L. O'Shea. Recycling: Social norms and warm-glow revisited. *Ecological Economics*, 90(Supplement C):10–18, 2013.
3. J. Castro, F. J. Quesada, I. Palomares, and L. Martínez. A consensus-driven group recommender system. *International Journal of Intelligent Systems*, 30(8):887–906, aug 2015.
4. P. B. Clark and J. Q. Wilson. Incentive systems: A theory of organizations. *Administrative Science Quarterly*, 6(2):129–166, 1961.
5. R. Conte, C. Castelfranchi, and F. Dignum. Autonomous norm acceptance. In *Intelligent Agents V: Agents Theories, Architectures, and Languages*, pages 99–112. Springer Berlin Heidelberg, 1999.
6. W. Emons. Escalating penalties for repeat offenders. *International Review of Law and Economics*, 27(2):170–178, 2007.
7. U. Gneezy, S. Meier, and P. Rey-Biel. When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, 25(4):191–210, 2011.
8. G. Governatori and A. Rotolo. Bio logical agents: Norms, beliefs, intentions in defeasible logic. *Autonomous Agents and Multi-Agent Systems*, 17(1):36–69, 2008.
9. C. He, D. Parra, and K. Verbert. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications*, 56:9–27, 2016.
10. Thomas King, Marina De Vos, Virginia Dignum, Catholijn Jonker, Tingting Li, Julian Padget, and Birna van Riemsdijk. Automated multi-level governance compliance checking. *Autonomous Agents and Multi-Agent Systems*, 31(6):1283–1343, Nov 2017.

11. B. Lepri, N. Oliver, E. Letouzé, A. Pentland, and P. Vinck. Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology*, 31(4):611–627, 2018.
12. J. Lu, D. Wu, M. Mao, W. Wang, and G. Zhang. Recommender system application developments. *Decision Support Systems*, 74(C):12–32, Jun 2015.
13. S. Mahmoud, N. Griffiths, J. Keppens, and M. Luck. Efficient norm emergence through experiential dynamic punishment. In *Proceedings of the 20th European Conference on Artificial Intelligence*, ECAI’12, pages 576–581, 2012.
14. J. Masthoff. Group recommender systems: Combining individual models. In *Recommender systems handbook*, pages 677–702. Springer, 2011.
15. T. Nguyen, PM. Hui, F. Harper, L. Terveen, and J. 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*, pages 677–686. ACM, 2014.
16. I. Nunes, S. Miles, M. Luck, S. Barbosa, and C. Lucena. Pattern-based explanation for automated decisions. In *Proceedings of the 21th European Conference on Artificial Intelligence*, ECAI’2014, pages 669–674, 2014.
17. G. T. Papaioannou and D. G. Stamoulis. A mechanism that provides incentives for truthful feedback in peer-to-peer systems. *Electronic Commerce Research*, 10(3):331–362, 2010.
18. F. Ricci, L. Rokach, and B. Shapira. *Recommender Systems Handbook*. Springer Publishing Company, Incorporated, 2nd edition, 2015.
19. M. Rodrigues and M. Luck. Cooperative interactions: An exchange values model. In *Coordination, Organizations, Institutions, and Norms in Agent Systems II*, volume 4386 of *Lecture Notes in Computer Science*, pages 356–371. 2007.
20. A. Salehi-Abari and C. Boutilier. Preference-oriented social networks: Group recommendation and inference. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys ’15, pages 35–42, New York, NY, USA, 2015. ACM.
21. C. Sibley. Responsible recommendations for irrational people. In *2017 AAAI Spring Symposium Series*, 2017.
22. N. Tintarev and J. Masthoff. Designing and evaluating explanations for recommender systems. In F. Ricci, L. Rokach, B. Shapira, and P. Kantor, editors, *Recommender Systems Handbook*, pages 479–510. Springer US, Boston, MA, 2011.
23. L. van der Torre. Contextual deontic logic: Normative agents, violations and independence. *Annals of Mathematics and Artificial Intelligence*, 37(1):33–63, 2003.
24. D. Villatoro, G. Andrighetto, J. Sabater-Mir, and R. Conte. Dynamic sanctioning for robust and cost-efficient norm compliance. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume One*, IJCAI’11, pages 414–419, 2011.
25. L. Xiao, Z. Min, Z. Yongfeng, G. Zhaoquan, L. Yiqun, and M. Shaoping. Fairness-aware group recommendation with pareto-efficiency. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, RecSys ’17, pages 107–115, New York, NY, USA, 2017. ACM.
26. F. López y López and M. Luck. Modelling norms for autonomous agents. In *ENC ’03: Proceedings of the 4th Mexican International Conference on Computer Science*, pages 238–245, Washington, DC, USA, 2003. IEEE Computer Society.
27. J. Zhang, R. Cohen, and K. Larson. Combining trust modeling and mechanism design for promoting honesty in e-marketplaces. *Computational Intelligence*, 28:549–578, 2012.