The Dilemma of Human-Like Collective Systems

Sofia Ceppi

The Edinburgh University, UK
sceppi@inf.ed.ac.uk

Abstract. Researchers that study Human-Like computing mainly aim to understand how systems can emulate human cognitive performance. However, when Human-Like systems are designed for sharing economy applications in which humans have to collaborate in order to achieve a desired task, there are several problems that needs to be addressed before asking how cognitive performance of a single human can be emulated. In this paper, we highlight these problems, provide examples related to the ridesharing scenario, and introduce how we approach these problems and which techniques we are using to tackle them.

One of the main aims of research in the Human-Like Computing area is to understand how a machine can emulate human cognitive performance. However, in sharing economy applications, i.e., in scenarios where humans have to collaborate in order to achieve a desired task, the first problem a Human-Like intelligent system needs to solve is to understand whose human performance it should emulate and toward which criteria it should aim to perform optimally.

Consider for example the ridesharing scenario in which the aim of an intelligent system is to group in rides drivers and commuters with matching requirements, while also working toward some global objective, e.g. increasing average occupancy of passenger vehicles. For each user, there may exist several feasible alternatives over which she or he has preferences. Given this, the system cannot emulate the performance of each single user because their preferences may be in conflict and collective goals, like maximising the number of rides, may not be achievable if the system is optimal for single users. Thus, there is the need to understand which principles would drive the decision making process of a human that has to divide her or his peers into groups.

In doing this, there are two critical factors that we should consider. The first one is that users may have different opinions about the collective goal the system should try to achieve. For example, assume that the system want to achieve fairness among users as collective goal. If passengers have different opinions about what would make fair rides, which criteria the system should try to optimise in order to achieve fairness? The second factor is that, even if every drivers and commuters in the ridesharing system aims to find peers to share a ride with, each of them may have different preferences about rides and thus they would try to obtain the best ride for themselves. For example, a user may prefer rides whose pick-up point is closer to his/her location, another user may be more flexible about pick-up location and care more about the departure and arrival time, and another user main concern may be to get the ride with the lowest price.

Furthermore, if we can safely assume that the system wants to be flexible in order to account for user diversity. the first step to do in order to achieve this goal, is to provide users with options among which each of them is free to choose his/her preferred one. However, due to the combinatorial nature of solutions of sharing economy applications and the fact that each user should be able to choose the option he/she prefers independently of other users' choices, it is possible that the choices, collectively, leads to infeasible or undesirable solutions. For example, consider the following ride sharing scenario characterized by three passengers, p_1, p_2, and p_3. The system provides each of them with three options: to share the ride with one of the other passengers or to ride alone. Now assume that passengers p_1 chooses the ride with p_2, p_2 chooses the ride with p_3, and p_3 chooses the ride with p_1. The collective of these choices lead to an undesirable solution because none of the passengers gets a ride even if there exist solutions, like p_1 and p_2 share a ride while p_3 does not share it, that would allow all passengers to get one. To avoid or mitigate this problem, the system should provide guidance to users during the selection process. Essentially, it is desirable for a human-like intelligent system to be able to support users coordination such that good solutions are achieved without explicitly involving the users in the decision and coordination process.

The ridesharing example points out some of the open questions that arise when Human-Like Computing systems are introduced in scenarios like the one of shared service where humans need to cooperate in order to achieve a common goal even if each of them may have different requirements and preferences. These questions are the ones we are studying and addressing as part of the SmartSociety project where particular attention has been given to the elicitation of users preferences, the generation of a set of solutions that take into account both collective and individual users goals, and the way in which these solution are recommended to users. In our current work (whose preliminary results are presented in [1]), we aim to design a recommender system that models human diversity and uncertainty of human behaviour while addressing the problems highlighted before.

In particular, we identify two macro problems. The first one is the multi-criteria optimization problem [3] due to the fact that the system needs to account for user diversity and thus needs to balance multiple, often contrasting, objectives. In particular, we focus our attention on two categories of contrasting objectives: the objective of a system that aims to optimise for the collective of users (e.g., maximising the number of users who assigned to a ride and their overall satisfaction with it), and the objective of a system that aims to be fair with all users (e.g., maximising the satisfaction of the unhappiest user or to equally satisfy the users). The approach we are using to address this problem consist in designing a flexible and general system that allows the designer to *(i)* explicitly control the trade-off between different objectives and *(ii)* change the objectives without the need to modify the all system.

The second problem we aim to address in our current work is the one of provide guidance to users during the selection process. This, can be seen as a coalition formation problem [5] and incentives can be provided to users who otherwise would reject a suggested coalition. In particular, incentives are designed with the purpose to modify user satisfaction for the proposed options such that all users prefer a collectively good solution. It is possible to obtain this effect with techniques like intervention [4] and the promise of future rewards or discount (i.e., techniques that provide explicit incentives) or with techniques and mechanisms from the game theory literature like cost of stability [2] and taxation [7] (i.e., techniques that provide implicit incentives).

In our current work we focus on the latter type of techniques. More in details, we compute the optimal taxation, i.e., the minimum amount of taxation needed such that each user (strictly or with a chosen probability) prefers a collectively good solution to the other offered options. Note that the computation of such optimal taxation depends on the user response model, i.e., how a user selects an option among several ones, and in our work we consider the three main models studied in the literature: noiseless model, constant noise model, and logit model [6]. The empirical analysis done so far shows that our approach improves the optimization of conflicting objectives and the implicit user coordination with respect to a benchmark that focuses only on the optimization of one objective and a benchmark that does not offer options to users who, in turn, reject the proposed solution if their satisfaction is not above a given threshold.

References

[1] Pavlos Andreadis, Sofia Ceppi, Michael Rovatsos, Subramanian Ramamoorthy: Diversity-Aware Recommendation for Human Collectives. In Proceedings of the Diversity workshop, (2016).

[2] Yoram Bachrach, Edith Elkind, Reshef Meir, Dmitrii Pasechnik, Michael Zuckerman, J org Rothe, and Jeffrey S. Rosenschein: The Cost of Stability in Coalitional Games. 122134, Springer Berlin Heidelberg, (2009).

[3] R. L. Keeney and H. Raiffa: Decisions with Multiple Objectives: Preferences and Value Trade-Offs. Cambridge University Press, (1993).

[4] Avi Segal, Ya'akov (Kobi) Gal, Ece Kamar, Eric Horvitz, Alex Bowyer, Grant Miller: Intervention Strategies for Increasing Engagement in Crowdsourcing: Platform, Predictions, and Experiments. IJCAI 2016: 3861-3867. (2016).

[5] Onn Shehory and Sarit Kraus: Methods for task allocation via agent coalition formation. Artificial Intelligence, 101(1), 165 200, (1998).

[6] Paolo Viappiani and Craig Boutilier: Optimal bayesian recommendation sets and myopically optimal choice query sets. In Advances in Neural Information Processing Systems, pp. 23522360, (2010).

[7] Yair Zick, Maria Polukarov, and Nicholas R. Jennings: Taxation and stability in cooperative games. In Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, pp. 523530, (2013).