Multi-Agent Systems for Enhancing Collective Intelligence

VITO TRIANNI, EUGENIA POLIZZI DI SORRENTINO, ANDREA GUIDO and GIULIA ANDRIGHETTO, Institute of Cognitive Sciences and Technologies (ISTC), National Research Council (CNR), Rome, Italy

ALBERTOANTONIONIandANGELSÁNCHEZ,UniversidadCarlosIIIdeMadrid(UC3M),Spain

1. INTRODUCTION

Collective intelligence (CI) can largely benefit from social interactions, which however bring about social biases—like herding—that jeopardise the advantages of group-work. Indeed, given a collective decision-making problem whereby a group must identify the correct answer among several alternatives, extensive argumentation and deliberation may not lead to high accuracy, let alone consensus within the group [Lorenz et al. 2011; Bahrami et al. 2012; Koriat 2015]. Groupthink, polarisation and balkanisation are the most evident examples of harmful effects in social systems [Granovskiy et al. 2015; Bang and Frith 2017]. Likewise, blindly pooling information from a set of individuals (the “wisdom of the crowd”) may be suboptimal, as the distribution of responses can be strongly biased [Jayles et al. 2017]. To enhance collective intelligence, social information should be carefully exploited.

When people participate to group decisions, they make explicit judgements on their own confidence and on the competence of others, which have a bearing on the way social information is evaluated and integrated. Such judgements do not come without error, as they are influenced by several factors pertaining the individual and the social sphere (e.g., wrong beliefs, conformism). This leads to a series of individual and social biases that can impact the ability of groups to make better decisions than isolated individuals [Bang and Frith 2017]. Experimental studies on the effects of social influence on collective decision making have provided contradictory results: in some cases, it enhances decision abilities [Navajas et al. 2018; Becker et al. 2017]; in other cases, it undermines the expected beneficial effects [Lorenz et al. 2011; Koriat 2015]. Crucially, subtle differences in the decision-making protocol seem to be decisive [Lorenz et al. 2011; Granovskiy et al. 2015], indicating that information flow design is key to harness collective intelligence.

Thanks to current advances in artificial intelligence and human-computer interaction, novel social technologies are available to counteract existing biases in group decision-making [Masthoff 2015; Castro et al. 2018; Sato et al. 2011;

Ganzer-Ripoll et al. 2017], leading to a profitable exploitation of social information and to more accurate decisions. In this work, we envisage a technology that implements targeted interventions during the decision-making process, exploiting an intelligent multi-agent system (MAS) that provides the interface to social information and mediates interactions within the group, with the goal of channeling the decision process toward optimal decision dynamics. To this end, the principled understanding of the group decision-making dynamics resulting from theoretical models can be exploited to implement the MAS [Reina et al. 2015; Reina et al. 2017]. Our hypothesis is that, by enacting the dynamics described by models into mixed human-agent groups, the group decision process can be channeled to follow bias-free dynamics. We performed online experiments with groups of 5 participants facing two alternative decision problems, characterised by different required abilities and cognitive loads. Our experimental results demonstrate that the MAS-mediated decisions are not affected by biases existing into social interactions, notwithstanding the type of problem faced by the group. These results open up a fruitful research direction and multiple potential applications in several domains, where the expertise and abilities of humans are matched with the efficiency and reliability of intelligent agents to provide higher levels of collective intelligence, leading to new socio-technical systems with a huge impact on future connected societies.

2. RESEARCH OBJECTIVES

In this study, we test MASs as mediating technologies, whereby agents take an active role in gathering and aggregating information from decision-makers. To this end, we carry out an experimental activity comparing both individual and group decisions, the latter being carried out either allowing participants to directly see information shared by others, or exploiting a MAS as mediator of the social information. The MAS is designed after theoretical models of collective decision making [Reina et al. 2015; Reina et al. 2017], and parametrised to counteract herding effects and to promote the achievement of consensus within the group.

We focus on discrete-choice problems where the best alternative must be chosen among several options. We consider problems entailing costly sampling—leading to noisy quality estimation—across multiple rounds. A classic scenario is the multi-armed bandit (MAB) problem, where exploration is necessary to understand the profitability of the dif- ferent options available [Toyokawa et al. 2018]. An alternative scenario involves a perceptual decision (PD) in which an image is presented for a short time and some noisy features must be estimated [Kurvers et al. 2015]. Here, images associated to different alternatives can be sampled—at some cost—in order to gain additional information about their features. Sampling introduces an exploration-exploitation dilemma at the individual and collective level. Exploration allows to evaluate options quality. Exploitation allows to select (and promote) one option basing on available knowl- edge. Individually, an explore/exploit trade-off must be found when exploration is costly. Collectively, individuals may free-ride exploration costs taking advantage of information available from others, potentially leading to herding phenomena [Raafat et al. 2009]. Under severe cognitive pressures (e.g., too many alternatives, time pressure), indi- vidual/social biases are expected to strike. For instance, accuracy loss was reported in a MAB problem when social information was available [Toyokawa et al. 2014]. The experiment tests the following hypotheses:

H1: Problem complexity affects collective decision accuracy, due to difficulties in reliably providing and processing social information.

H2: MAS-mediated interactions reduce the effects of biases. If we introduce a MAS that mediates interactions among decision-makers, higher decision accuracy is expected.

3. EXPERIMENTAL DESIGN

We consider multiple-choice problems (number of choices: M = 5) administered to groups (number of participants per group: N = 5). In each game, participants have a starting endowment (E = 80 points). A game consists in multiple rounds (number of rounds R 2 [40, 50], stochastic), wherein participants can make costly explorations (cost per exploration: C = 2 points), can rate alternatives on a 5-stars scale, and interact with others. A limited number of costless exploration rounds R_e = 10 is granted to each participant at the beginning to promote exploration. Anytime, participants may select one option as their "current best", possibly revising previous judgments. The group choice is made by majority voting with a quorum Q = 75%: if correct, participants—whatever their individual choice—retain their residual endowment, otherwise everybody loses everything. We vary task type (MAB vs PD) and availability of social feedback. When participants play individually (S condition), we obtain a baseline of the individual performance in the task. In the social conditions, we test the following treatments:

- Condition C: participants see which option is chosen by others (to control for baseline social information)

- Condition R: participants see also the rating of options chosen by others (to control for possibly-biased social information)

- Condition M: participants interact with the MAS designed after a collective decision model (to test the mediating effect, see [Reina et al. 2015; Reina et al. 2017]) The MAS in condition C evaluates the likelihood that each option is correct on the sole basis of the choices made by participants, without using any rating information. It is designed to remain initially unbiased and to reinforce social information only towards the end of the experiment, trying to converge on the current majority displayed by the users.

In addition to these conditions, we consider a virtual treatment build ex post by grouping together participants that played in the S condition. This treatment—referred to as V—represents a situation in which a group has to decide but without any possibility of interaction. It has been conceived as a baseline to test the effects of the harsh quorum requirements for the social conditions, allowing a fairer comparison with cases in which people do not interact.

Before playing the game, we include a risk aversion test and a training session where the participant can familiarise with the MAB and PD problems, and where we can gather data on the participants confidence calibration [Moore and Healy 2008]. At the end of the game, we administer a questionnaire to record the appreciation and difficulties encountered during the experiment. Each participant plays only one combination of the possible treatments, being assigned a game type (MAS or PD) and one of the four treatments (S, C, R, M). This choice has been made to constrain the length of the experiment within one hour. At the end of the experiment, participants receive a payment associated to their performance within the training and the experimental session, plus a show-up fee. The average payment is about 10e.

4. RESULTS AND CONCLUSIONS

We performed several experimental sessions between May 2019 and February 2020. In total, 1152 different partici- pants enrolled into the experiment. Considering that experiments were performed online, participants could drop out at any time. Hence, out of the initial pool, 1084 participants successfully terminated the training phase, and 986 com- pleted the main experiment. For the latter, we could obtain 175 individuals playing in the S treatment (75 for MAB and 100 for PD), and 165 valid groups in the social treatments (78 for MAB and 87 for PD). A group is considered valid if it terminates with at least 4 active individuals (e.g., only one dropout is tolerated during the experimental phase).

During the training sessions, we observed that individuals could solve a binary decision problem with an average success rate of 60%. At the same time, the subjective confidence was in general higher, with an average of 80%, without substantial difference between MAB and PD problems. This suggests that, in the binary decision problem, over-confidence is very likely, and this may bring biases in the group decisions pulling down the group performance.

Looking at the results of the experimental sessions, we found that the problem type (MAB vs. PD) has a significant effect on the way in which social information is used. Problems that suffer from possibly misleading information (like MAB) are badly affected by the rating information, which may be biased in favour of incorrect options. Problems that entail a high cognitive load (like PD) may instead suffer from herding effects, as participants tend to follow other's opinion without acquiring sufficient information individually. Interestingly, the purposely designed mediating technology enables to retrieve good performance in both the MAB and PD conditions, showing that it is possible to design the way social information is distributed to counteract biases in group decisions.

This experiment demonstrates, once again, that group decision making can be strongly affected by individual and social biases. However, it is possible to devise decision protocols that are robust to biases, and that can correct detri- mental decision dynamics resulting from social feedback. In particular, we demonstrated that subjective information like ratings not always provides advantages for group decisions, especially in the presence of misleading information and overconfident individuals. The usage of subjective judgements can have a strong impact on the final outcome, yet it is not possible to control the quality of the information conveyed through ratings when these are costless. The choices made by participants provide a more reliable information because they honestly represent the participant's commitment to one of the available solutions, which cannot be inflated by subjective opinions. We also demonstrated that complex problems—e.g., problems that require a considerable effort in terms of cognitive load—may lead to herding when social information is readily available at no additional cost. In this cases, social information can be useful only if it is either modulated by honest signalling of uncertainty (e.g., through rating) or delayed until sufficient individual evidence has been gathered. In our experiment, a mediating technology that uses only the information about the choices made by different users—hence avoiding subjective judgements—and at the same time limits the herding effects proved to be the best possible choice. This finding can be useful to practitioners in collective intelligence, who can therefore design their interaction protocol on the basis of these suggestions, placing particular care to identify whether the problems they are confronted with may be affected by misleading information, herding effects or both.

ACKNOWLEDGMENTS

This work has been supported by the 2019 Collective Intelligence Grants from Nesta, the innovation foundation (https: //www.nesta.org.uk).

REFERENCES

Bahador Bahrami, Karsten Olsen, Dan Bang, Andreas Roepstorff, Geraint Rees, and Chris Frith. 2012. What failure in collective decision-making tells us about metacognition. *Philos. Trans. Royal Soc. B* 367, 1594 (2012), 1350–1365.

Dan Bang and Chris Frith. 2017. Making better decisions in groups. *Royal Society Open Science* 4, 8 (2017), 170193–22.Joshua Becker, Devon Brackbill, and Damon Centola. 2017. Network dynamics of social influence in the wisdom of crowds. *PNAS* 114 (2017),

5070–5076.J Castro, J Lu, G Zhang, Y Dong, and L Mart ́ınez. 2018. Opinion Dynamics-Based Group Recommender Systems. *IEEE Transactions on Systems,*

Man, and Cybernetics: Systems 48, 12 (2018), 2394–2406.Jordi Ganzer-Ripoll, Maite Lo ́pez-Sa ́nchez, and Juan Antonio Rodriguez-Aguilar. 2017. A Multi-agent Argumentation Framework to Support

Collective Reasoning. In *COREDEMA Conflict Resolution in Decision Making*, Reyhan Aydog ̆an, Tim Baarslag, Enrico Gerding, Catholijn M

Jonker, Vicente Julian, and Victor Sanchez-Anguix (Eds.). Springer, 100–117.Boris Granovskiy, Jason M Gold, David J T Sumpter, and Robert L Goldstone. 2015. Integration of Social Information by Human Groups. *Top.*

Cogn. Sci. 7, 3 (2015), 469–493.Bertrand Jayles, Hye-rin Kim, Ramo ́n Escobedo, Ste ́phane Cezera, Adrien Blanchet, Tatsuya Kameda, Cle ́ment Sire, and Guy Theraulaz. 2017.

How social information can improve estimation accuracy in human groups. *PNAS* 114, 47 (2017), 12620–12625.Asher Koriat. 2015. When two heads are better than one and when they can be worse: The amplification hypothesis. *Journal of Experimental*

Psychology: General 144, 5 (2015), 934–950.Ralf H J M Kurvers, Max Wolf, Marc Naguib, and Jens Krause. 2015. Self-organized flexible leadership promotes collective intelligence in human

groups. *Royal Society Open Science* 2, 12 (2015), 150222–10.Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *PNAS* 108

(2011), 9020–9025.Judith Masthoff. 2015. Group Recommender Systems: Aggregation, Satisfaction and Group Attributes. In *Recommender Systems Handbook*,

Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer, 743–776.Don A Moore and Paul J Healy. 2008. The trouble with overconfidence. *Psychological Review* 115, 2 (2008), 502–517.Joaquin Navajas, Tamara Niella, Gerry Garbulsky, Bahador Bahrami, and Mariano Sigman. 2018. Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour* 2, 2 (2018), 126–132.Ramsey M Raafat, Nick Chater, and Chris Frith. 2009. Herding in humans. *Trends in Cognitive Sciences* 13, 10 (2009), 420–428.Andreagiovanni Reina, James A R Marshall, Vito Trianni, and Thomas Bose. 2017. Model of the best-of-N nest-site selection process in honeybees.

Phys Rew E 95 (2017), 52411–5.AndreagiovanniReina,GabrieleValentini,CristianFerna´ndez-Oto,MarcoDorigo,andVitoTrianni.2015. ADesignPatternforDecentralised

Decision Making. *PLoS ONE* 10, 10 (2015), e0140950–18. GilsonYukioSato,HiltonJose´SilvadeAzevedo,andJean-PaulABarthe`s.2011. Agentandmulti-agentapplicationstosupportdistributed

communities of practice: a short review. *Autonomous Agents and Multi-Agent Systems* 25, 1 (2011), 87–129.Wataru Toyokawa, Hye-rin Kim, and Tatsuya Kameda. 2014. Human Collective Intelligence under Dual Exploration-Exploitation Dilemmas. *PLoS*

ONE 9, 4 (2014), e95789–9.Wataru Toyokawa, Andrew Whalen, and Kevin N Laland. 2018. Social learning strategies regulate the wisdom and madness of interactive crowds.

bioRxiv (2018), 326637.