Optimized Human-AI Decision Making: A Personal Perspective

Alex Pentland

Massachusetts Institute of Technology
pentland@mit.edu

ABSTRACT

AI is turning up everywhere, but when people try to use it to as a tool help make better decisions it often stumbles. . .users reject it, rely too much on it, and so forth. Not only is this a problem for AI-as-tool applications, it is increasingly clear that AI without human oversight is prone to bad mistakes, typically because the AI has such a narrow view of the world and can't tell when it is violating norms or when the context has changed. As a consequence, AI-automation is getting serious pushback from citizens and lawmakers. What are we to do in order to integrate AI tools into human group decision making?

CCS CONCEPTS

• **Computer systems organization**; • **Human computer systems**; • **System reliability**;

KEYWORDS

Decision systems, Human-AI systems, Cognitive Science

ACM Reference Format:

Alex Pentland. 2021. Optimized Human-AI Decision Making: A Personal Perspective. In *Proceedings of the 2021 International Conference on Multimodal Interaction (ICMI '21)*, October 18–22, 2021, Montréal, QC, Canada. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3462244.3479880>

1 INTRODUCTION

Integrating AI with human group decision making is a critical problem since most decisions are either made in formal teams or heavily influenced by social networks. There is huge interest in the idea of extending human intelligence by use of AI tools, however along with this hunger to make human decision making better, there are numerous examples that illustrate the problems encountered when trying to integrate human decision making with AI. For instance, IBM's Watson experienced real problems providing advice in the medical domain, and expert systems in 1980s-1990s often failed to achieve real-world adoption because of social and psychological issues (see, for instance, Wikipedia for more details).

2 STATE OF THE ART

There is an immense literature about human social decision making and about computer tools to help decision making, but much of it is limited or even contradictory because humans act differently in different contexts [1]. Moreover, there is relatively little in the

way of a general, computational account of how the information flow within group decision making affects outcomes. Most theory is about the effects of specific features of discussions in general, and so applies equally well to individual, group, and crowd decision making of all sorts, but as a consequence gives only limited specific guidance for optimizing social / group decision making.

As a consequence, efforts to improve group decision making have focused on improving specific discussion features. For instance, there is a substantial literature on improving decision making by “nudging” people to guide attention towards more careful evaluation, toward considering argument and discussion structures, or providing missing evidence. These sorts of interventions have shown positive results using interaction formats ranging from wearable “advice agents”, to on-line small-group feedback, to feedback within large-scale social media. However, they address only specific types of failures in human cognition rather than offering a comprehensive theory of group and social decision making, they cannot offer a general account of the “best” way for a group decision making discussion to proceed.

2.1 Optimal Decision Making

In contrast, there are some powerful mathematical results concerning combining evidence from different agents to achieve optimal, minimum regret decision making. The criteria of minimum regret means that an agent or group of agents make the best decision possible at each time period given the information and previous experience available at the time.

Classic minimum-regret decision making is called Thompson sampling and the literature often refers to this as a *bandit problem*, because of the formal equivalence to the question of which slot machine (aka “one arm bandit”) to try in a gambling casino (see https://en.wikipedia.org/wiki/Thompson_sampling). In the last decade the mathematical solution to such problems have been extended to distributed agents, e.g., a gambler observes the payouts of other casino patrons and combines those observations with their personal knowledge to decide which slot machine to try next.

The importance of these mathematical results is that they provide a strategy that a group of agents may use to form an optimal (minimum regret) policy for action. The core of the distributed Thompson sampling strategy is for each decision making agent to use the experience of other agents to form an estimate of the prior probability for each potential action, and then multiply this prior by each actions' likelihood distribution as determined by their personal knowledge. This produces a posterior distribution of what reward each action is expected to produce.

The Thompson sampling strategy for optimal (minimum regret) decision making has demonstrated excellent performance in many domains (see Wikipedia), enough so that it is a standard approach in domains such as signal processing, medical decision making, and finance. It shows very fast convergence to optimal policies, good

This work is licensed under a Creative Commons Attribution International 4.0 License.

ICMI '21, October 18–22, 2021, Montréal, QC, Canada
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8481-0/21/10.
<https://doi.org/10.1145/3462244.3479880>

generalization to new and changing situations [2], and the ability to work with very noisy and ill-conditioned data inputs.

2.2 Recent Progress

Recently, several research groups (including my own) have extended the distributed Thompson sampling optimality results to cover privacy-sensitive, communication-limited social networks analogous to the networks seen in human behavior and in social media [2–4]. This optimality guarantee has also been extended to adversarial situations and situations where different agents have different goals by having each agent compare the choices they would make to the choices others, and from that comparison estimate the utility function of the other agents [3]. This allows agents to select the best subset of agents with which to trade experiences, and identify agents that are acting in an adversarial manner or appear to be reacting to "fake news".

In more plain English, these results provide a statistical estimation framework that allows groups or networks of agents with different experiences and different goals to trade experiences without privacy compromising details in order to maximize their personal set of goals [3]. The result is a guaranteed minimum regret decision sequence for each agent, and indeed there is an analogous optimality result for the group as a whole.

As a consequence, these mathematical optimality results are fairly closely analogous to the sorts of decision processes that we would like to have in human groups and networks, and it is natural to ask if the strong optimality guarantees provided by these mathematical discoveries can be applied to better understand and improve human decision making and perhaps even to serve as a framework for designing human-AI tools. To explore this question, let us first look at animal behavior.

2.3 The Exploration-Exploitation Dilemma

Thompson sampling is recognized as a good description of the group foraging behavior of many social animals (see Wikipedia). Group foraging may be viewed as a "portfolio strategy" where the animals don't just choose the maximum likelihood action (e.g., the action that has been yielding the most food), instead they make the frequency of different actions proportional to the likelihood. This provides a solution to the exploration-exploitation dilemma, where animals must allocate some effort to actions that have been reliably producing good results, while at the same time exploring for new resources, with the consequence that their portfolio of frequent actions continuously evolves over time.

Studies of human financial decision making, as well as human mobility and shopping behavior over many days show exactly this sort of exploration-exploitation characteristic [4, 5, 7]. An important function of the exploration actions is to avoid behavioral rigidity, where only a limited number of very familiar actions are chosen within a group. The phenomenon of insufficient exploration, resulting in a static portfolio of actions, is familiar in human discussions as echo chambers and group-think. Exploration of novel actions is therefore critical for avoiding unforeseen risks, finding new opportunities, and adapting to changing conditions. For instance, financial experts use observations about the actions of other experts to control their portfolio risk [5]. One of the important

strengths of the distributed Thompson sampling framework is that it gives you a formal method of determining when there is insufficient exploration and thus risk of group-think, as well as a formula for de-risking decisions by accounting for rare outcomes or unusual opinion sampling error.

3 HUMAN GROUP DECISION MAKING

Inspired by these mathematical decision systems results, my research group has worked on the cognitive science version of this decision making literature and feel that we have made interesting progress toward relating human decision making to Thompson sampling style optimal decision making. This work began with our 2010 Science paper [6] which showed that the collective intelligence of small human groups is different from, and often more effective than, individual human intelligence and that it is heavily dependent on entropy in the patterns of communication, and our 2013 Science Advances paper [7] which showed that day-to-day behavior of human populations exhibit the same exploration-exploitation behavior seen in social animal species. These results were extended by our 2020 Proceedings of the National Academy of Science paper [1] which showed that group performance depends on selecting agents with similar utility functions as the prior probability for decision making.

More recently, our 2021 Cognition paper [8] showed that commonly observed individual-level social heuristics closely approximate minimum regret group decision making (e.g., distributed Thompson sampling) and accurately model human small-group behavior in consumer financial markets. Our 2021 Entropy paper [5] showed that financial experts show the same distributed Thompson sampling behavior, and in particular use the social information of others to control portfolio risk.

We have also seen that the Thompson sampling framework provides a good model of some important examples of multimodal learning and decision making. For instance, in our 2002 Cognitive Science paper on infant language learning, we showed that visual cues from adults provided critical cues allowing infants to locate and disambiguate nouns within the audio signal [9]. In more recent work we have shown that the distributed Thompson sampling framework provides a useful model for humans to create and learn category boundaries in visual stimuli [10], and for highly efficient domain generalization in image classification [2].

4 ACHIEVING REAL-WORLD IMPACT

As a consequence of this convergence of decision theory and cognitive science, I think that the mathematical progress and cognitive science research together form a foundation for building human-AI systems that have solid optimality properties, at least in domains that have explicit learning outcomes analogous to financial decision making. The general idea is that AI's should be seen as just part of the social network that humans naturally use to make decisions, and the distributed Thompson sampling framework gives us a recipe for combining AI forecasts with human opinions in order to select actions. In our experiments with providing computer-processed feedback human experts in the financial domain, this distributed Thompson sampling decision framework seems to work quite well [5].

Some of the most successful human-AI decision making interfaces seem to follow this framework. For instance, Google search results (at least in their original format) was a nice example of what may be required: they listed the items that humans linked to most centrally. The perception of the user was that this was a picture of the distribution of most common human opinions (although processed by machine), and there were many different opinions listed.

This sort of “voting behavior” interface seems to be typically perceived as a probability distribution that is very much like the sort of prior probability used in Thompson sampling [7], and people adapt to using visualization of this “range of opinions” information quite easily. In our experiments with financial experts we used a similar approach, and showed financial experts the distribution of other expert’s opinions, and asked if they would like to alter their decision. Those experts who used the “other” opinions in line with the Thompson sampling framework suffered fewer big losses than ones that didn’t [5]. Decision quality in these sorts of human-AI systems seems to depend on the provided results being seen as representing the distribution of opinions, and so it is critical that the user interface design maintains diversity and does not pretend to be authoritative

To explore the generality of this approach to enhancing human intelligence we need to build and test systems such as mobile phone agents, on-line meeting feedback systems, and systems for improving discussion with close friends on social media. Similarly, we need to explore content areas chosen because of their potential societal impact, the ability to easily get thousands to millions of examples of team or network decision making, the ability to have quantitative metrics of performance, and the fact that we have extensive experience in these areas and interaction platforms. Example areas include financial decision making, personal health decision making (e.g., vaccination), and validity of news articles in social media.

This suggests a series of challenge contests, with common data covering domains such as consumer financial investment, on-line education, and large-scale social media discussions of health and politics. Ideally, code would also be made openly available.

Finally, there should also be close attention to safety and ethical dimensions. In addition to “immediate” problems such as bias, fairness, and unintended outcomes, we need to consider longer term potential impact. For instance, will people lose certain skills and knowledge when they become reliant on machines for a certain tasks? What are the consequences, both ethical and practical, of adopting systems that we cannot fully comprehend?

ACKNOWLEDGMENTS

This research is the thesis work of two “generations” of my graduate students, who have been supported by my Human Dynamics group within the MIT Media Lab’s Industry Consortium and my Trust Data Consortium within MIT’s Institute for Data Systems and Society.

REFERENCES

- [1] Abdullah Almaatouq, Alejandro Noriega-Campero, Abdurahman Alotaibi, Peter M Krafft, M Moussaid, Alex Pentland, 2020, daptive social networks promote the wisdom of crowds, *Proceedings of the National Academy of Sciences*, 117 (21) 11379–11386; <https://doi.org/10.1073/pnas.1917687117>
- [2] Abhi Dubey, V. Ramanathan, Alex Pentland, Dhruv. Mahajan, 2021, Adaptive Methods for Real-World Domain Generalization, , *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 14340–14349
- [3] Abhi Dubey, Alex Pentland, 2020, Private and Byzantine-Proof Cooperative Decision-Making, *20th Conf. Autonomous Agents and Multiagent Systems*, 2020, p.357–365, <http://proceedings.mlr.press/v119/dubey20a/dubey20a.pdf>
- [4] Abhi Dubey, Alex Pentland, 2020, Differentially-Private Federated Linear Bandits, *Proc. Advances in Neural Information Processing Systems*, <https://arxiv.org/pdf/2010.11425.pdf>
- [5] Dhaval Adjudah, Yan. Leng, S.K. Chong, Peter. M. Krafft, Eestban. Moro, Alex. Pentland, 2021, Accuracy-Risk Trade-Off Due to Social Learning in Crowd-Sourced Financial Predictions, *Entropy* 2021, 23(7), 801; <https://doi.org/10.3390/e23070801>
- [6] Anita W Woolley, Christopher F Chabris, Alex Pentland, Nadia Hashmi, Tom W Malone, 2010, Evidence for a collective intelligence factor in the performance of human groups, *Science* 330 (6004), 686–688
- [7] Coco Krumme, Alejandro Lorente, Manuel Cebrian, Estaban Moro, Alex Pentland, 2013, The predictability of consumer visitation patterns, *Scientific Reports*, Vol.3, 1645
- [8] Peter M Krafft, Erez Shmueli, TomL Griffiths, J oshB Tenenbaum, Alex. Pentland 2021, Bayesian collective learning emerges from heuristic social learning, *Cognition* 212, 104469
- [9] Deb K Roy, Alex P Pentland, 2002, Learning words from sights and sounds: A computational model, *Cognitive science* 26 (1), 113–146
- [10] Ziv Espstien, Mat Groh, Abhi Dubey, Alex Pentland, 2021, An Experimental Study of Social Influence and Information Design, *ACM CSCW*