Heuristics for Real-time Internet Control of Situated Human Agents

NICCOLO PESCETELLI, MANUEL CEBRIAN, IYAD RAHWAN, MIT Media Lab

1. INTRODUCTION

The collective action problem defines a situation in which individuals would be better off cooperating but fail to do so because of coordination failures affecting large groups or simply due conflicting personal preferences preventing joint action. From an evolutionary perspective, social norms and institutions are considered solutions to this problem, facilitating the suppression of selfish instincts in favor of greater group outcomes [Richerson and Henrich 2009; Henrich et al. 2010]. Yet, while individuals can respond to changes in their environment within seconds, the time-scale at which these institutions operate or change can often range from years to decades. This time scale limits the ability of large collectives to act as ”one mind and one body” to flexibly react to situated problems. To bypass the limitations of large group size, groups often use hierarchies to quickly reach decisions, which allow fewer individuals holding decision power to act as representatives of the collective will [Turchin et al. 2013; Bowles 2009; Hallpike 2008]. Whether collective situated action can be achieved using a fast decentralized decision system remains to be understood.

In this submission, we describe a platform that we created as a playful attempt to experiment with new creative solutions to collective action [BBC Technology 2018]. The platform, called BeeMe, is accessible at https://beeme.online and can be demoed with a live or online audience. We have elsewhere described the differences between BeeMe and other forms of audience participation in entertainment and art [Pescetelli et al. 2019]. BeeMe allows a scalable number of online users to observe and collectively control the actions of one individual (the agent) operating in physical space during scheduled events. The paper describes data obtained during one of such events, which went live on Halloween night 2018, when thousands of users played simultaneously. Online users had to coordinate their suggestions to the agent (using an in-built chat) in order to solve a sci-fi mystery. The agent had to be guided through a series of missions that unlocked new tasks and developed the narrative further. Although the narrative gave a structure to the game, users were free to choose whatever action they preferred [Phillips 2012; Harrigan and Wardrip-Fruin 2010]. BeeMe represents a new frontier in collective systems for gaming and art, alongside Reddit Place and Twitch Plays Pokemon [Aleta and Moreno 2019; Müller and Winters 2018]. BeeMe is designed to algorithmically maintain a running tally of most requested actions, by allowing users to suggest commands for the agent as well as to vote on other users' suggestions. Another individual (the operator) relays requested actions to the agent via an earpiece. Due to the large number of users in the 2018 event, the voting feature was temporarily unavailable, letting the operator in charge of aggregating the crowd's will via monitoring chat messages into an actionable series of actions that ultimately achieved the group's goal (*i.e.*completing the mission). This unexpected silver lining allowed us to investigate (1) how a human interpreter aggregated online discussions in real time and (2) whether an algorithm can achieve the same or better performance than a human in interpreting the ”will of the crowd”.

Results suggest that human operators were often biased and did not follow the majority when interpreting the crowd's will. For example, they tended to relay commands that were only mentioned once and thus not representative of the majority opinion (here called for simplicity *singletons*). For this reason, simple algorithms trained on raw human data had difficulties fitting more than a third of the operators' issued commands. However, after filtering out such arbitrary segments, algorithmic performance greatly improved. Several algorithms were designed to read in human discussions (in forms of chat logs) and map them into relevant action commands representing the democratic view in the crowd. We find that machine-generated commands often loosely match human-generated commands. When the two differ, machine-generated commands are still sensible (*i.e.*, actionable) given the recent past chat discussion. We discuss how these findings inform the debate on Internet mediated collective intelligence, democratic representation and collective action.

The platform

The platform consists of (1) a central video-stream where online users can observe the surroundings of the agent; (2) an experimenter text box, where the experimenter could send messages to all users to give the crowd new missions or advance the game's narrative; (3) a chat system that allows users to coordinate and socialize; (4) a command suggestion box, where users can suggest their own commands (like "turn left" or "run away"); (5) a voting mechanism that users can use to up-vote and down-vote others' commands. In the 2018 event, online users were split in two teams each controlling a different character in the story (*i.e.*, a different agent) who was freely roaming the MIT campus. The game consisted of a series of 11 missions that had to be carried out under time pressure to defeat the characters' foe.

2. RESULTS

We created a simple algorithm that is able to read in a continuous stream of users' chat messages and dynamically updates the most wanted command. The model used by the real-time aggregator is trained on a heuristic matching function fitted to the human operator's behavior. We segmented the chat data based on the commands issued by the operators. Chat messages and operators' commands were transformed into bi-grams (verb+noun). We then designed heuristics to maximize the proportion of matches between human generated bi-gram commands Y and the machine generated bigram commands y as a function of the segment chat messages:

$$\frac{1}{N}(\sum_{i=1}^{N} f_k(X_i) = Y_i) \tag{1}$$

where X_i is the chat data relative to a segment i and $f_k(X_i)$ is an algorithm k operating on the segment's chat input data. We experimented with different matching functions of increasing complexity and found a *decay match* function to better represent the human data, achieving 84% performance on the testing set. This algorithm ignores old chat messages (above a threshold t) and computes a weighted sum for each remainng bi-gram, where weights are proportional to a square rooted harmonic function: $\alpha/sqrt(d+\beta)$. Here, α is a constant ($\alpha = 1000$), d is the time difference between the bi-gram occurrence i and the issued action, and β is a dampening constant that controls the bias towards instances that are very close to end of the segment (here, $\beta = 5000$).

We then fixed the parameters that best described human commands, and used the algorithm to dynamically update a tally of all available bi-grams derived from the chat and ranks them by relevance. The tally score S for each bi-gram b is computed on the basis of the *decay match* test as:

$$S_b = \sum_{i=1}^{M} \frac{\alpha}{\sqrt{(d_i + \beta)}} \tag{2}$$

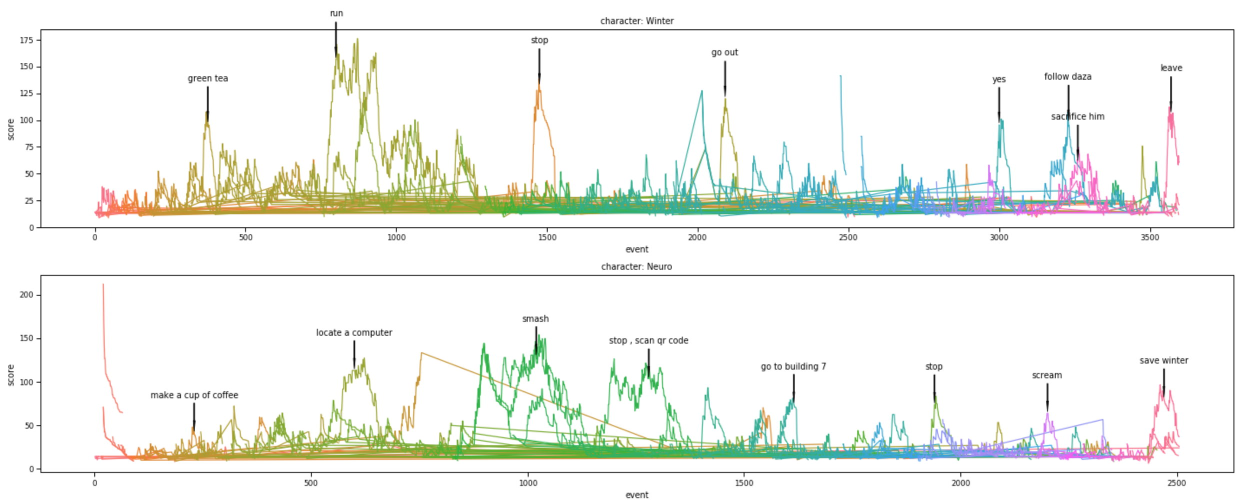

Fig. 1. Real-time decoding of aggregated human suggestions. The aggregator takes in a continuous stream of chat messages and dynamically updates the most popular bi-gram (verb+noun pair) to be communicated to the agent.

where α is a constant (here set to 1000.0), β is a decay dampener controlling the rate of decay for the relevance of each bi-gram's occurrence i (set to 5000), d is the time difference between the current time and each occurrence timestamp ($t_0 - t_i$) and M is the total number of occurrences of a bi-gram, namely all messages mapping onto the same bi-gram b.

Results (Figure 1) show the score timeseries of the top three bi-grams at each timestep over the course of the 2018 Halloween event. Color represents different bi-grams. For reference, we have annotated original chat messages corresponding to noticeable scoring peak. It can be seen that score peaks represent highly relevant commands that correspond to popular users' suggestions, most of which were issued by the human operator on the night of the event. We concluded that the online aggregator can be used in the future on the BeeMe platform to entirely replace the operator's role as interpreter of the "will of the crowd".

3. IMPLICATIONS AND FUTURE DIRECTIONS

Our results are the first attempt at describing the BeeMe ecosystem, where human users interact together online to produce in real-time a situated action stream via a surrogate (the agent) operating in physical space. BeeMe offers a playful yet powerful demonstration of real-time situated decision-making at scale, with applications in gaming, social mobilization and collective participation.

Firmly believing in the potentiality of this tool, we also notice that this work represent only a case study (*i.e.*, one single large scale event), and thus its results must be interpreted with caution. Rather than offering a precise description of collective intelligence in this environment, we tried here to capture the intuition that simple aggregation heuristics operating on real-time chat entries can achieve remarkable performance in dynamically interpreting the "will of the crowd".

Finally, we draw attention to the importance of such large scale gamified social experiments. Although unique and difficult to replicate, these events are valuable *per se* because they embody a proof of *existence*, highlighting the potential (and boundaries) of collective intelligence in the digital era. They spark the imagination of users as well as practitioners interested in improving collective intelligence design.

REFERENCES

Alberto Aleta and Yamir Moreno. 2019. The dynamics of collective social behavior in a crowd controlled game. *EPJ Data Science* 8, 1 (12 2019), 22. DOI:http://dx.doi.org/10.1140/epjds/s13688-019-0200-1

BBC Technology. 2018. MIT invites you to control a human on Halloween. (2018). https://www.bbc.com/news/business-46016696

S. Bowles. 2009. Did Warfare Among Ancestral Hunter-Gatherers Affect the Evolution of Human Social Behaviors? *Science* 324, 5932 (6 2009), 1293–1298. DOI:http://dx.doi.org/10.1126/science.1168112

C. R. Hallpike. 2008. *How We Got Here. From bows and arrows to the space age*. AuthorHouse, Central Milton Keynes.

P. Harrigan and N. Wardrip-Fruin. 2010. *Second person: Role-playing and story in games and playable media*. The MIT press, Cambridge, MA.

J. Henrich, Jean Ensminger, R. McElreath, Abigail Barr, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins Gwako, N. Henrich, C. Lesorogol, F. Marlowe, D. Tracer, and J. Ziker. 2010. Markets, Religion, Community Size, and the Evolution of Fairness and Punishment. *Science* 327, 5972 (3 2010), 1480–1484. DOI:http://dx.doi.org/10.1126/science.1182238

Thomas F. Müller and James Winters. 2018. Compression in cultural evolution: Homogeneity and structure in the emergence and evolution of a large-scale online collaborative art project. *PLOS ONE* 13, 9 (9 2018), e0202019. DOI:http://dx.doi.org/10.1371/journal.pone.0202019

Niccolo Pescetelli, Manuel Cebrian, and Iyad Rahwan. 2019. From BeeMe to Bandersnatch: The Youniverse of science and fiction. (2019). https://medium.com/mit-media-lab/from-beeme-to-bandersnatch-the-youniverse-of-science-and-fiction-46867aab2174

Andrea Phillips. 2012. *A Creator's Guide to Transmedia Storytelling: How to Captivate and Engage Audiences across Multiple Platforms*. McGraw-Hill.

Peter Richerson and Joe Henrich. 2009. Tribal Social Instincts and the Cultural Evolution of Institutions to Solve Collective Action Problems. *SSRN Electronic Journal* 3, 1 (2009). DOI:http://dx.doi.org/10.2139/ssrn.1368756

P. Turchin, T. E. Currie, E. A. L. Turner, and S. Gavrilets. 2013. War, space, and the evolution of Old World complex societies. *Proceedings of the National Academy of Sciences* 110, 41 (10 2013), 16384–16389. DOI:http://dx.doi.org/10.1073/pnas.1308825110