Enterprise Crowd Computing for Human Aided Chatbots

Bozzon, Alessandro

DOI

[10.1145/3195555.3195566](https://doi.org/10.1145/3195555.3195566)

Publication date

2018

Document Version

Accepted author manuscript

Published in

Proceedings of the 1st International Workshop on Software Engineering for Cognitive Services

Citation (APA)

Bozzon, A. (2018). Enterprise Crowd Computing for Human Aided Chatbots. In *Proceedings of the 1st International Workshop on Software Engineering for Cognitive Services* (pp. 29-30). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3195555.3195566>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Enterprise Crowd Computing for Human Aided Chatbots

Alessandro Bozzon
Delft University of Technology
Delft, The Netherlands
a.bozzon@tudelft.nl

ABSTRACT

A chatbot is an example of cognitive computing system that emulates human conversations to provide informational, transactional, and conversational services. Despite their widespread adoption, chatbots still suffer from a number of performance issue due to limitations with their programming and training. In this paper we discuss *Human Aided Chatbots*, i.e. chatbots that rely on humans in the loop to operate. Human Aided Chatbots exploit human intelligence, brought for instance by crowd workers or full-time employees, to fill the gaps caused by limitations of fully automated solutions. A recent example of Human Aided Chatbots is Facebook M. To achieve broader adoption, Human Aided Chatbots must overcome a number of issues, including scalability, low-latency, and privacy. In this short paper, we discuss how Crowd Computing performed in the enterprise could help overcoming such issues. We present some recent findings in the field of Enterprise Crowd Computing, and introduce ECrowd, a platform for enterprise crowd computing designed for gathering training data for cognitive systems.

CCS CONCEPTS

• **Information systems** → **Crowdsourcing**; *Chat*; • **Computing methodologies** → *Intelligent agents*;

KEYWORDS

Crowd Computing, Chatbot, Human Aided Chatbots

ACM Reference Format:

Alessandro Bozzon. 2019. Enterprise Crowd Computing for Human Aided Chatbots. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The dream of a human-like highly intelligent computer assistant is not new. In science fiction books and movies, from Hal in “2001: A space odyssey” (1968) to Jarvis in “Iron Man” (2013), personal assistants helped heroes in their life and to manage their work duties. In research, the first conversational agents appeared in the 60’s. They were explicitly programmed, using rule-based approaches.

They had clear limitations in terms of performance and real-world utility. Recent advances in parallel processing hardware, natural language understanding machine learning, and artificial intelligence enabled the creation of a new generation of personal

assistants like Google Assistant, Siri and Alexa. *Chatbots* are text-based conversational agents, living in messenger applications (e.g. Facebook Messenger, Telegram, Whatsapp and WeChat), and emulating a conversation with a human to provide *informational*, *transactional*, or *conversational* services [5]. Figure 1 shows the reference chatbot architecture introduced in [6]. Given a user request, the language understanding (LU) component infer the user’s intent and the associated information; the action execution & information retrieval (AEIR) component performs the requested actions, or retrieves the information of interest from its data source; the response generation (RG) component builds a response to the user; finally, the dialog management (DM) component preserves and update the context of a conversation to request missing information, to process clarifications by users, and to ask follow-up questions.

Figure 1: Chatbot Architecture

State-of-the-art chatbots are still far from being perfect, struggling to serve well user requests and to carry on a meaningful conversation. Each of the chatbot’s components suffer from a number of limitations when employed in real world scenarios. For instance, the LU might fail interpreting user requests, the DM might miss clarification requests for missing information, the AEIR might not find the requested information or not execute the correct requested action, and the RG might not provide a satisfactory response.

All the issues above are accountable to limitations with rule-based algorithms, or with machine learning models trained with incomplete or biased datasets.

Human intelligence could be used to address the limitations of fully-automated solutions. Intuitively, a human being proficient with the language used in the conversation, and having access to the Web, could easily overcome such issues; a human can adapt and perform well even when input or instructions themselves change. *Crowd Computing* [3, 4] is a computational paradigm that advocates the use of human processing power to solve problems that computers cannot yet solve.

2 HUMAN AIDED CHATBOTS

Human Aided Chatbots are chatbots that utilise Crowd Computing in at least one of its components. *Facebook M*¹ was a personal assistant based on the Facebook Messenger platform, and a successful example of Human Aided Chatbot. Facebook M became operational in 2015 to help developing Facebook's artificial intelligence projects, and has been discontinued in January 2018. Facebook M allowed the execution custom transactional tasks with help of a crowd of dedicated Facebook employees.

In [5] we review the state-of-the-art in the field, and discuss a number of open research questions that need to be addressed to improve performance of chatbots, and the quality of their interaction. While human aided bots like Facebook M can serve thousands of users, widespread adoption of human aided chatbots require novel crowd computing solutions that would ensure that with the growth of the number of chatbot users, the costs associated with human computation would grow only gradually (*scalability*), and the latency due to human interaction would be contained (*real-timeliness*). Another challenge is related to *privacy*: as users interacting with chatbots often need to share their personal information, new methods need to be designed and developed to address privacy concerns.

3 ECROWD: ENTERPRISE CROWD COMPUTING

Enterprises have been adopting crowd computing to sustain their business needs and processes, transferring the practices of crowd-sourcing from the online environment to the internal crowd of the enterprise: the *employees*. *Enterprise Crowd Computing* differs from traditional crowd computing in terms of both the crowd it involves (i.e., employees) and the problem it targets (i.e., business problems) [8]. Crowd Computing provides an effective way to exploit the internal knowledge profiles of employees and to leverage on their non-utilised working capacity to solve business critical tasks [8]. Moreover, employees operate under contracts signed in order to commit to corporate norms and values, including, for instance, intellectual property rights and *privacy* [8]. Such commitment makes employees always available for contribution, possibly in an opportunistic manner [1].

Given its characteristics, we advocate for the use of Enterprise Crowd Computing in the context of Human Aided Chatbots.

In the context of a collaboration with the IBM Benelux Centre for Advanced Studies, we developed ECrowd [1, 7], a platform for Enterprise Crowd Computing designed to support the creation of custom mobile crowd computing applications for data creation, enrichment, and analysis. ECrowd has been successfully employed in a number of cognitive computing scenario, including the training of chatbots. Figure 2 depicts some examples of cognitive tasks deployed in ECrowd, which include: *Information Extraction*, to support the development of domain-specific language understanding (LU) chatbot module; *Moral Machine* (a task based on a research on the morality of future Artificial Intelligence [2]); and *Cell Count*, for the development of machine learning application in the medical domain. ECrowd has been developed with real-time task assignment, privacy-preservation functionalities, and notification capabilities.

However, the solicitation of prompt contributions from employees remain a challenge. In a recent work [1], we have shown that employees are generally willing to perform crowd computing tasks during their break times, but are willing to devote only short attention spans to it (2–10 minutes). Employees are not bothered by smartphone notifications, although an excessive amount of notification can facility the feared “crowd-out” effect. Moreover, the nature of the task is also of importance, where learning and purpose (i.e. support to the company's goals) appear to be the most compelling reason for participation.

Figure 2: Cognitive Computing Tasks in ECrowd

4 CONCLUSIONS

Our experience with the ECrowd platform suggests that Enterprise Crowd Computing can support Human Aided Chatbots in terms of capacity and privacy, but more work is needed. For instance, it is essential to develop employees engagement strategies able to cope with the real-time requirements of conversational agents.

REFERENCES

- [1] Sarah Bashirieh, Sepideh Mesbah, Judith Redi, Alessandro Bozzon, Zoltán Szilávik, and Robert-Jan Sips. 2017. Nudge Your Workforce: A Study on the Effectiveness of Task Notification Strategies in Enterprise Mobile Crowdsourcing. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*. ACM, New York, NY, USA, 4–12. <https://doi.org/10.1145/3079628.3079692>
- [2] Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2016. The social dilemma of autonomous vehicles. *Science* 352, 6293 (2016), 1573–1576.
- [3] Alessandro Bozzon, Marco Brambilla, Stefano Ceri, Andrea Mauri, and Riccardo Volontario. 2014. Pattern-Based Specification of Crowdsourcing Applications. In *Web Engineering*, Sven Casteleyn, Gustavo Rossi, and Marco Winckler (Eds.). Springer International Publishing, Cham, 218–235.
- [4] Alessandro Bozzon, Piero Fraternali, Luca Galli, and Roula Karam. 2014. Modeling Crowdsourcing Scenarios in Socially-Enabled Human Computation Applications. *Journal on Data Semantics* 3, 3 (01 Sep 2014), 169–188. <https://doi.org/10.1007/s13740-013-0032-2>
- [5] Pavel Kucherbaev, Alessandro Bozzon, and Geert-Jan Houben. 2018. Human Aided Bots. *IEEE Internet Computing* In printing (2018).
- [6] Michael McTear, Zoraida Callejas, and David Griol. 2016. *The Conversational Interface: Talking to Smart Devices* (1st ed.). Springer Publishing Company, Incorporated.
- [7] Benjamin Timmermans, Zoltán Szilávik, Manfred Overmeier, and Alessandro Bozzon. 2017. ECrowd: Enterprise Crowdsourcing for Training Cognitive Systems using the Workforce. In *Proceedings of the 29th Benelux Conference on Artificial Intelligence. BNAIC 2017*.
- [8] Maja Vukovic and Claudio Bartolini. 2010. Towards a research agenda for enterprise crowdsourcing. In *Proceedings of the 3rd International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (ISoLA)*. Springer, 425–434.

¹<http://bit.ly/2qKLaBX>