

On Automating Conversations

Ting-Hao (Kenneth) Huang

College of Information Sciences and Technology
Pennsylvania State University, University Park, PA, USA
txh710@psu.edu

Abstract

From 2016 to 2018, we developed and deployed Chorus, a system that blends real-time human computation with artificial intelligence (AI) and has real-world, open conversations with users. We took a top-down approach that started with a working crowd-powered system, Chorus, and then created a framework, Evorus, that enables Chorus to automate itself over time. Over our two-year deployment, more than 420 users talked with Chorus, having over 2,200 conversation sessions. This line of work demonstrated how a crowd-powered conversational assistant can be automated over time, and more importantly, how such a system can be deployed to talk with real users to help them with their everyday tasks. This position paper discusses two sets of challenges that we explored during the development and deployment of Chorus and Evorus: the challenges that come from being an “agent” and those that arise from the subset of conversations that are more difficult to automate.

Deploying and Automating Chorus

Interaction in rich natural language enables people to exchange ideas efficiently and achieve their goals more quickly. Modern personal intelligent assistants, such as Apple’s Siri and Amazon’s Echo, utilize conversation as their primary communication channel and illustrate a future where conversing with computers will be as easy as talking to a friend. However, despite decades of research, modern conversational assistants are still limited in domain, expressiveness, and robustness.

In response, we created a system that blends real-time human computation with AI and has real-world, open conversations with users. Instead of bootstrapping automation from the bottom up with only automatic components, we took a top-down approach that started with a working crowd-powered system. We developed and deployed a crowd-powered conversational assistant, *Chorus* (Huang et al. 2016), and then created a framework, *Evorus*, that enables Chorus to automate itself over time (Huang, Chang, and Bigham 2018). *Evorus* allowed external task-oriented chatbots and chatterbots to be added into Chorus to take over

parts of conversations, reuse crowd-submitted responses to answer future similar questions, and gradually learn to select high-quality responses to reduce its reliance on crowd oversight. Figure 1 shows the worker interface and a system overview of *Evorus*. To make the deployment more robust, we invented a new recruiting method, the Ignition model (Huang and Bigham 2017), to hire workers quickly. In our two-year deployment, more than 420 users talked with Chorus, having over 2,200 conversation sessions.

During the deployment of *Evorus*, conversations averaged 9.90 messages sent from users, 13.58 messages sent from crowd workers, and 1.93 messages sent from automatic chatterbots (Huang, Chang, and Bigham 2018). Thus, **automated responses were chosen 12.44% of the time**. As a comparison, without any automation, a conversation averaged 8.73 user messages and 12.98 crowd messages. To examine conversation quality, we sampled conversations with accepted automatic responses and a matching set without automated contributions. For each, 8 MTurk workers rated [Satisfaction, Clarity, Responsiveness, Comfort], which was based on the PARADISE’s objectives for evaluating dialogue performance (Walker et al. 1997) and the Quality of Communication Experience metric (Liu, Chua, and Stahl 2010), on a 5-point Likert scale (with 5 being the best.) The original conversations (N=46) had an average rating of [3.47, 4.04, 3.88, 3.56], while those with automatic responses (N=54) had scores of [3.57, 3.74, 3.52, 3.66]. The similar results suggest that the automatic components did not make conversations worse.

Having human workers in the system loop enabled Chorus to talk with real users to help them with their everyday tasks and demonstrated to us the challenges of conversational assistants that have not yet been fully explored by state-of-the-art automated chatbots. In this position paper, we discuss two sets of challenges that are less discussed in the AI community (Huang 2018): the common drawbacks of being an “agent” and human conversations beyond the scope of task-oriented dialogues.

Being an “Agent”

Chorus, *Evorus*, and the concept of a “crowd agent” (Lasecki and Bigham 2013b) strongly imply the interaction metaphor

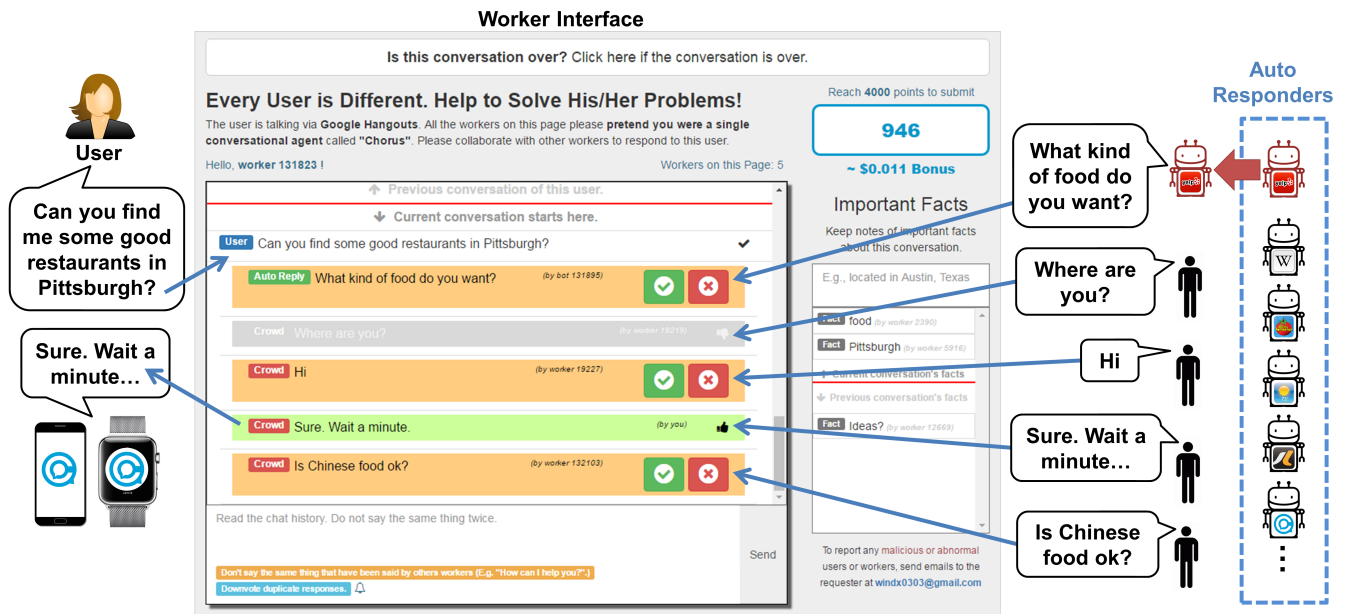


Figure 1: The Evorus worker interface allows workers to propose responses and up/down vote candidate responses. The up/down votes give Evorus labels to use to train its machine-learning system to automatically gauge the quality of responses. Evorus automatically expires response candidates upon acceptance of another in order to prevent workers from voting through candidate responses that are good but no longer relevant. Workers can tell each message is sent by the end-user (blue label), a worker (red label), or a chatbot (green label) by the colored labels.

of an “agent.” We are aware of the historical debate about the pros and cons of the “agent” metaphor for human-computer interaction. Ben Shneiderman famously opposes this model. He argued that people are different from computers, and “human-human interaction is not necessarily an appropriate model for human operation of computers.” (Shneiderman 2010) In 1997, Shneiderman and Pattie Maes had a landmark debate about the agent metaphor (Shneiderman and Maes 1997) and the pros and cons of “direct manipulation” and “software agents.” Shneiderman argued for the power of good user interfaces and visualizations and again stated that human-human interaction should not be the model for human operation of computers. Maes argued that having full control of unlimited actionable items and overly complex tasks is impractical, and thus future users will need “extra eyes and extra ears” to which to delegate tasks. In 2017, they revisited this debate. In this latest round, Maes argued for the concept of “human-computer integration,” while Shneiderman countered that computers are tools rather than equal partners of humans, and the goal of interactive technology should be “ensuring human control, while increasing the level of automation.” (Farooq et al. 2017)

We would like to start our discussion with an example of a real-world complex task: tax preparation. To prepare their taxes, most people either hire human experts (*i.e.*, accountants or tax professionals) or use specialty tax software (*e.g.*, TurboTax). These reflect the two sides of the “agent” metaphor debate: Human experts are hired agents who use conversations and natural languages (*e.g.*, email) as primary means of communication; conversely, tax programs use a

nicely designed workflow and user interface to transform the complex tax preparation process into a task that most users can accomplish. Each solution has its advantages and disadvantages. Human experts are more intelligent at understanding user needs and preferences, more flexible in communication, better at acknowledging context to provide customized suggestions, and able to save time and effort for users. However, the gaps between a human professional’s capability and what state-of-the-art technologies can do is significant. Without proper automated solutions, hiring a tax professional is expensive, and an expert’s time is not scalable. On the other hand, tax software gives users more direct control and instant feedback on the interface and is typically more affordable. However, software is less flexible, harder to customize for complex tax situations, and most importantly, requires much more time and effort from users, who must manually input data.

Our work aimed to relax the limitations set by hiring humans and attempted to reveal the true potential of the agent metaphor. The gap between what people want to accomplish (for example, interacting naturally) and what technology is capable of – known as the “social-technical gap” – is one of the central challenges for studying computer-supported cooperative work (CSCW) and HCI (Ackerman 2000). Our work showed how to use automation to reduce the gaps between human agents and automated solutions, but to fully explore the real trade-offs of agents as compared to direct manipulation (tools), we need more competent systems. We agreed with Maes’s insight that the software agent is powerful. However, we are also aware that to become as useful

as we imagine agents could be, we must be much closer to human-level capabilities.

Miscommunication Users need to communicate problems and context to an agent to delegate tasks, which provides opportunities for miscommunication. Although Chorus enabled open conversation with users, the flexibility and expressiveness of natural language mean ambiguity is inevitable in human languages. Therefore, users’ verbal specifications could be interpreted in different ways. Misinterpretation of language was not the only cause of miscommunication: Users also omitted important contextual information until prompted by the agent. The agent cannot know beforehand all the context, preferences, and constraints of the user, so miscommunication is inevitable. However, one of the primary advantages of our system is its ability to adapt to the user through open, multi-turn communication.

Hard or Expensive for Hands-on Controls Authorizing an agent to complete tasks means users do not have hands-on control of all task details. The agent often needs to make small decisions without constantly checking with the user, as such interruptions would devalue the automated component of the system. However, lack of control could be a problem for tasks that need to be precisely customized or when users simply prefer full control. In other words, if users hire an agent to personalize something for them (*e.g.*, a gift for mom), it would be difficult or expensive for the agent to mimic the same level of personalization that the user could provide. This is one of the fundamental limitations when hiring any agency.

Potentially Longer Task Completion Time Another limitation of outsourcing a task to an agent is that the end-to-end task completion time could be longer than that required if the users did the task themselves, especially for tasks where users were proficient. Users sometimes did not use Chorus due to the task being difficult. For tasks that were familiar or easy, it might take Chorus longer than the user because of the need to communicate with the agent, iterate back and forth to narrow down results, and finally obtain a result. In our experiments with InstructableCrowd, in some scenarios, the time the system took to produce IF-THEN rules was longer than that required by users on a mobile interface. While the users do not need to commit any time or effort while waiting, we acknowledge that communicating with Chorus takes time and the end-to-end completion time could be longer.

Higher User Expectation Finally, another fundamental limitation is that anthropomorphic agents create higher expectations of users, which can harm the user experience or create greater frustration when problems arise. For instance, a user might become more frustrated with an Amazon Echo device that failed to execute a command to “turn off the light” than with a manual switch that did not work. However, we also argue that higher user expectations will motivate users to explore new tasks that they did not think were possible and thus push the boundaries of the capability of

conversational assistants. While the system ended up disappointing some users (Chen 2016), it opened up many new research problems and use cases, influencing how today’s users interact with personal assistants. Over the long term, higher expectations are not a limitation but an opportunity.

The Conversations That Are Harder to Automate

While having human workers in the loop enables Chorus to hold realistic conversations with users, we observed limitations in terms of handling or automating social conversation and open-ended discussions. These limitations reflect the current bottleneck of conversational assistants and could indicate the future challenges of automating conversations.

Automating Domain-Independent Tasks and Social Conversations The Evorus framework focused only on automating domain-specific tasks. However, some real-world tasks are cross-domain or even domain-independent. It would be difficult for Evorus to automate these tasks because its bot-selection algorithm used topic similarity to choose chatbots. For instance, negotiation (*e.g.*, salary or price) is a common task in the real world that can happen in many different domains, including real estate, shopping, or trading. If Evorus included a “negotiation bot” that could negotiate with users, the bot-selection algorithm would have a hard time using domain similarity to decide when to use it. To date, our work has focused more on tasks than on social conversations. While the current version of Evorus included four chatbots to provide general responses, we did not incorporate approaches or knowledge about social conversations in our system. As in non-domain tasks, the bot-selection algorithm is also inefficient in selecting social chatbots because they do not represent a clear, single domain.

Open-Ended Discussions The deployed version of Chorus did not help workers organize information in a structured way, something that would help support open-ended discussions. The system only had a “memory board” where workers could write down a *list* of important facts. While this approach was sufficient to support most conversations, a list is insufficient for organizing all the arguments and information for long, open-ended long discussions such as “which school should I go to?” or “what’s the meaning of working?”

Ackerman studied how computer systems can be used to enhance organizational or group memory (Ackerman 1998; Ackerman and Malone 1990). Lasecki and Bigham experimented using workers’ ratings to rank notes recorded on Chorus’ memory board according to their relevance to the current conversation (Lasecki and Bigham 2013a). Gouravajhala *et al.* created a more advanced version of a memory board, Mnemo, where workers can summarize a note and estimate its longevity by answering questions like “Will this fact still be true by [the end of today/next week/next month/next year]?” (Gouravajhala *et al.* 2018). However, we did not focus on improving the memory board to enable more advanced structures for organizing information.

Furthermore, Chorus and Evorus used worker consensus to choose high-quality responses, assuming that most work-

ers likely agreed on the same high-quality responses. However, sometimes this assumption did not hold. In realistic open-ended discussions, topics are often subjective, where workers do not agree with each other easily or quickly. Chorus currently cannot hold these types of conversations without breaking the single-agent illusion. One potential solution would be to automatically detect such topics and alter the worker interface or voting mechanism correspondingly.

The “Top-Down” Research Approach

One lesson of our work is the effectiveness of the “top-down” research approach. One core advantage of this top-down approach is that users can start using the system on day one, providing realistic data to guide future automation. Furthermore, thanks to the oversight of crowd workers, such human-in-the-loop frameworks allow automated components to make more mistakes, opening more spaces for algorithms to try different strategies.

At a higher level, crowd-powered approaches are robust, intelligent, and do not require training data, but they are slow and costly. In contrast, automated approaches are fast, scalable, and more affordable, but they require large amounts of training data and have many more capability limitations. Our work attempted to answer this critical question: How do we combine crowd-powered and automated components wisely to leverage the advantages of each while minimizing the disadvantages? Our take was slightly different from the approaches of many prior attempts, which added pieces of human work into a larger automated system architecture. Instead, we introduced pieces of automation into human-powered systems. We believe that – even in this era of AI, machine learning, and deep learning – combining human intelligence with automated approaches will result in systems that are more robust, usable, and scalable, thereby creating a better world.

References

- Ackerman, M. S., and Malone, T. W. 1990. Answer garden: A tool for growing organizational memory. In *Proceedings of the ACM SIGOIS and IEEE CS TC-OA Conference on Office Information Systems*, COCS '90, 31–39. New York, NY, USA: ACM.
- Ackerman, M. S. 1998. Augmenting organizational memory: A field study of answer garden. *ACM Trans. Inf. Syst.* 16(3):203–224.
- Ackerman, M. S. 2000. The intellectual challenge of cscw: the gap between social requirements and technical feasibility. *Human-Computer Interaction* 15(2-3):179–203.
- Chen, B. X. 2016. Siri, alexa and other virtual assistants put to the test. *The New York Times*. Retrieved from: http://www.nytimes.com/2016/01/28/technology/personaltech/siri-alexa-and-other-virtual-assistants-put-to-the-test.html?_r=0.
- Farooq, U.; Grudin, J.; Shneiderman, B.; Maes, P.; and Ren, X. 2017. Human computer integration versus powerful tools. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '17, 1277–1282. New York, NY, USA: ACM.
- Gouravajhala, S. R.; Jiang, Y.; Kaur, P.; Chaar, J.; and Lasecki, W. S. 2018. Finding mnemo: Hybrid intelligence memory in a crowd-powered dialog system. In *Collective Intelligence Conference*.
- Huang, T.-H. K., and Bigham, J. P. 2017. A 10-month-long deployment study of on-demand recruiting for low-latency crowdsourcing. In *Proceedings of The fifth AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2017)*. AAAI.
- Huang, T.-H. K.; Lasecki, W. S.; Azaria, A.; and Bigham, J. P. 2016. “is there anything else i can help you with?”: Challenges in deploying an on-demand crowd-powered conversational agent. In *Proceedings of AAAI Conference on Human Computation and Crowdsourcing 2016 (HCOMP 2016)*. AAAI.
- Huang, T.-H. K.; Chang, J. C.; and Bigham, J. P. 2018. Evorus: A crowd-powered conversational assistant built to automate itself over time. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, 295:1–295:13. New York, NY, USA: ACM.
- Huang, T.-H. K. 2018. *A Crowd-Powered Conversational Assistant That Automates Itself Over Time*. Ph.D. Dissertation, Carnegie Mellon University.
- Lasecki, W. S., and Bigham, J. P. 2013a. Automated support for collective memory of conversational interactions. In *First AAAI Conference on Human Computation and Crowdsourcing*.
- Lasecki, W. S., and Bigham, J. P. 2013b. Interactive crowds: real-time crowdsourcing and crowd agents. In *Handbook of Human Computation*. Springer. 509–521.
- Liu, L. A.; Chua, C. H.; and Stahl, G. K. 2010. Quality of communication experience: definition, measurement, and implications for intercultural negotiations. *Journal of Applied Psychology* 95(3):469.
- Shneiderman, B., and Maes, P. 1997. Direct manipulation vs. interface agents. *interactions* 4(6):42–61.
- Shneiderman, B. 2010. *Designing the user interface: strategies for effective human-computer interaction*. Pearson Education India.
- Walker, M. A.; Litman, D. J.; Kamm, C. A.; and Abella, A. 1997. Paradise: a framework for evaluating spoken dialogue agents. In *Proceedings of the Association for Computational Linguistics*, ACL '98, 271–280. Stroudsburg, PA, USA: Association for Computational Linguistics.