

Getting Virtually Personal: Chatbots Who Actively Listen to You and Infer Your Personality

Michelle X. Zhou
Juji, Inc.
mzhou@juji-inc.com

Wenxi Chen
Juji, Inc.
wchen@juji.com

Ziang Xiao
University of Illinois at Urbana
Champaign
zxiao5@illinois.edu

Huahai Yang
Juji, Inc.
hyang@juji-inc.com

Tracy Chi
Juji, Inc.
tchi@juji-inc.com

Ransom Williams
Juji, Inc.
arwilliams@juji-inc.com

ABSTRACT

We build conversational agents to serve as AI interviewers who engage a user in a one-on-one, text-based conversation. Our live demos showcase two special skills of an AI interviewer: (a) ability to actively listen to a user during an interview—responding to complex and diverse user input empathetically; and (b) the ability to automatically infer the user’s Big 5 personality traits from the interview. We provide an overview of the technologies that enable these two abilities and real-world applications of such AI interviewers.

CCS CONCEPTS

• Computing Methodologies → Intelligent Agents • Human-centered computing → Interactive systems and tools

KEYWORDS

Chatbot, AI Interviewer, Personality Inference

ACM Reference format:

Michelle X. Zhou, Wenxi Chen, Ziang Xiao, Huahai Yang, Tracy Chi and Ransom Williams. 2019. Getting Virtually Personal: Chatbots Who Actively Listen to You and Infer Your Personality. In *24th International Conference on Intelligent User Interfaces (IUI '19 Companion)*, March 17-20, Marina Del Rey, CA, USA. ACM, New York, NY, USA, 2 Pages. <https://doi.org/10.1145/3308557.3308667>

1 INTRODUCTION

We have built conversational agents to serve as AI interviewers who can conduct personalized interviews through 1:1 text-based

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IUI '19, March 17–20, 2019, Marina del Rey, CA, USA
© 2019 Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-6272-6/19/03...\$15.00
<https://doi.org/10.1145/3308557.3308667>

conversations. Figure 1 shows an example conversation of our AI interviewer with a user.

To make AI interviewers act like effective human interviewers [1], we want to power them with two sets of skills. First, we want an AI interviewer to actively listen to an interviewee, and respond empathetically to elicit quality responses from the interviewee. For example, the AI interviewer shown in Figure 1 is able to grasp the gist of the user’s self-introduction and summarize it in its response, which makes the user feel heard. Second, we want an AI interviewer to read between the lines during an interview to assess the unique characteristics of an interviewee, just like what a human interviewer would do.

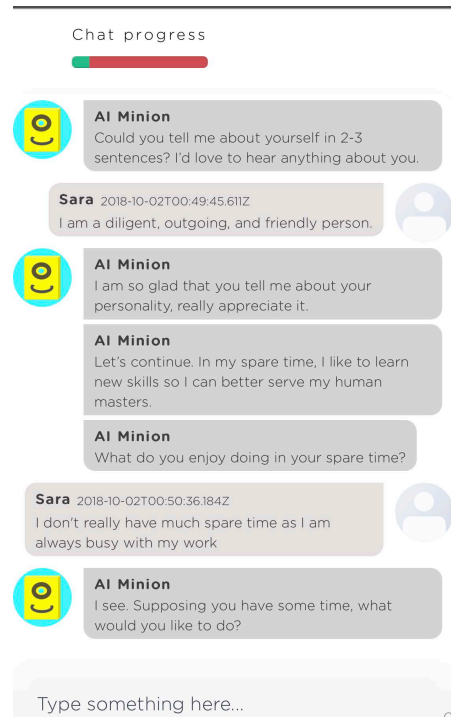


Figure 1. A screenshot of an interview conducted by an AI interviewer (AI Minion) and a user (Sara).

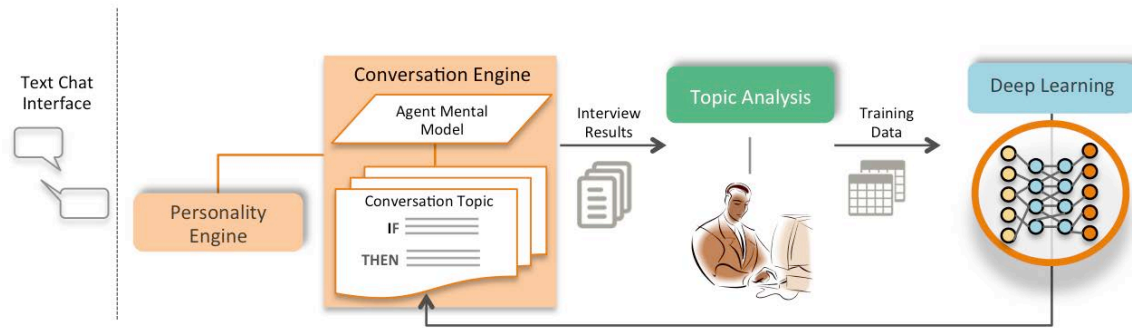


Figure 2. Overview of a hybrid approach to an AI interviewer and its continuous improvement.

Despite the recent advances in conversational agents, it is non-trivial to power AI interviewers with such capabilities due to the complex nature of open-ended interview conversations. For example, it is challenging to interpret complex and diverse user input to an open-ended interview question such as “*What’s the top challenge you are facing*”, let alone responding to such input empathetically. To exacerbate the problem, we also want an AI interviewer to analyze a user’s conversation on the fly to assess his/her characteristics, such as intellectual curiosity.

To address all of the challenges, we have developed two sets of technologies: (1) a hybrid conversation engine that couples deep learning with a topic-based conversation model; and (2) an evidence-based personality engine that automatically mines a user’s conversation text to infer his/her Big 5 personality traits.

In this demo, we will showcase the key capabilities of our AI interviewers in multiple real-world use cases and explain the two corresponding technologies behind the scene.

2 OVERVIEW OF OUR APPROACH

Figure 2 provides an overview of our approach to building an AI interviewer. The *conversation engine* defines an agent’s mental model (e.g., the model of an interviewer) that drives a set of conversation topics. Each topic is governed by a set of rules, each of which would be triggered by a user’s input (e.g., a user’s response to an interview question) or context parameters (e.g., the allotted interview time is reached). A triggered rule drives the AI interviewer to act (e.g., responding to a user’s input). Since rules alone cannot handle complex and diverse user input, we augment them with deep-learning trained models based on human interviewee responses collected by our AI interviewers.

The *topic analysis* module uses enhanced Latent Dirichlet Allocation (LDA) [2] to analyze the training data and identify various semantic themes. A human annotator then examines and rectifies the topic analysis results. The labeled training data can be used to build various data-driven models, for example, training text classification models to predict latent semantics of interviewee responses. The trained models are then incorporated into topic-specific rules to interpret the semantics of a user input and generate empathetic responses during an interview.

To personalize each conversation, we have developed a personality engine that analyzes a user’s conversation text and

automatically infers the user’s Big 5 personality traits [3]. An AI interviewer can use the inferred personality traits to personalize a conversation, e.g., asking users different interview questions based on their personality traits.

As shown in Figure 2, our approach supports continuous enhancements of an AI interviewer. The rules within each topic bootstrap an AI interviewer to conduct interviews and collect interview results. The interview results are then analyzed and labeled as training data to train data-driven models, which are to enhance the capabilities of AI interviewers.

3 USE CASE DEMONSTRATIONS

Our AI interviewers can be used in a wide variety of real-world applications, such as interviewing job applicants and surveying users for product or market research. We will show an AI interviewer’s active listening and personality inference capabilities through live demos, including opinion elicitation and personalized recommendation (e.g., recommending food or travel locations) based on user personality traits inferred from a chat. Moreover, we show how our approach enables rapid adaptation of an AI interviewer to a new interview task, which is critical to the adoption of conversational agents. We will also explain the limitations of our current approach and discuss future research directions to further improve AI interviewers.

ACKNOWLEDGMENTS

This work is in part supported by the Air Force Office of Basic Research under FA9550-15-C-0032. We’d also like to thank all our current and previous colleagues and interns at Juji for making this demo possible.

REFERENCES

- [1] Bauer, C., Figl, K. and Motschnig-Pitrik, R. (2010), Introducing ‘Active Listening’ to Instant Messaging and E-mail: Benefits and Limitations. IADIS International Journal on WWW/Internet.
- [2] Song, Y., Pan, S., Liu, S., Wei, F., Zhou, M.X., and Qian, W. (2013) Constrained text co-clustering with supervised and unsupervised constraints. IEEE TKDE 25(6): 1227-1239.
- [3] Zhou, M.X., Mark, G., Li, Y., and Yang, H. (2018) Trusting Virtual Agents: Effect of Personality. ACM Transactions on Interactive Intelligent Systems, Special issue Highlights of IUI 2017. To appear.