InterviewBot: Real-Time End-to-End Dialogue System to Interview Students for College Admission

Anonymous ACL submission

Abstract

We present the InterviewBot that dynamically integrates conversation history and customized topics into a coherent embedding space to conduct 10 mins hybrid-domain (open and closed) conversations with foreign students applying to U.S. colleges for assessing their academic and cultural readiness. To build a neural-based end-to-end dialogue model, 7,361 audio recordings of human-to-human interviews are automatically transcribed, where 440 are manually corrected for finetuning and evaluation. To overcome the input/output size limit of a transformer-based encoder-decoder model, two new methods are proposed, context attention and topic storing, allowing the model to make relevant and consistent interactions. Our final model is tested both statistically by comparing its responses to the interview data and dynamically by inviting professional interviewers and various students to interact with it in real-time, finding it highly satisfactory in fluency and context awareness.

1 Introduction

With the latest advancement of Conversational AI, end-to-end dialogue systems have been extensively studied (Zhang et al., 2020; Adiwardana et al., 2020; Roller et al., 2021). One critical requirement is context awareness; robust dialogue systems must consider relevant parts in conversation history to generate pertinent responses (Serban et al., 2017; Mehri et al., 2019; Bao et al., 2020; Zhou et al., 2021; Xu et al., 2022). However, these systems still suffer from issues such as hallucination, inconsistency, or lacking commonsense (Bao et al., 2021), hindering them to take a place in real applications.

Numerous admission interviews are given every year to students located in 100+ countries applying to colleges in the U.S., where the interviews are often conducted online. Those interviews are usually unscripted with an emphasis on asking the applicants thought-provoking questions based on their interests and experiences. The main objective is to provide decision makers (e.g,. admissions officers, faculty members) with an unfiltered look at those students in a daily academic environment.

Building an interview chatbot, called InterviewBot, will save time and effort of the interviewers, and provide foreign students a cost-efficient way of practicing interviews when native speakers are not available. Nonetheless, there are a few hurdles for developing an end-to-end InterviewBot. First, it is hard to collect a sufficient amount of data covering dialogues crossing open & closed domains (Sec. 3). Second, most transformer-based encoder-decoder models adapted by current state-of-the-art systems are not designed to handle long context; thus, they often repeat or forget previously discussed topics (Sec. 4). Third, it is demanding to find appropriate people to interactively test such a dialogue system with the professional objective (Sec. 5).

This paper presents an end-to-end dialogue system that interacts with international applicants to U.S. colleges. The system questions on critical perspectives, follows up on interviewee's responses for in-depth discussions, and makes natural transitions from one topic to another until the interview ends, which lasts about 30 turns (5 mins for text-based, 10 mins for spoken dialogues). To the best of our knowledge, it is the first real-time system using a neural model, completely unscripted, conducting such long conversations for admission interviews.

2 Related Work

Dialogue systems can be categorized into closed- and open-domain systems (Ilievski et al., 2018). Closed-domain systems require efficient access to domain knowledge (Lian et al., 2019) and serve specific professions such as education (Cunningham-Nelson et al., 2019), healthcare (Fan et al., 2021; Amiri and Karahanna, 2022), or customer service (Baier et al., 2018; Nichifor et al., 2021). Open-domain systems converse across multiple domains with natural transitions (Adiwardana et al., 2020) and conduct interactions in a broader horizon (Ahmadvand et al., 2018; Wang et al., 2017; Finch et al., 2020). For admission interviews, however, the conversation is often a mixture of closed (job-related questions) and open domains (general aspects of the applicant) dialogues, which makes it more challenging to build an end-to-end system.

Several dialogue systems have been developed to communicate with humans for information exchange or elicitation (Finch et al., 2020; Li et al., 2017; Kim et al., 2019) across multiple domains (Safi et al., 2020; Khoa, 2021; Okonkwo and Ade-Ibijola, 2021). Conversational agents for interviews have been experimented with for law enforcement (Minhas et al., 2022), healthcare (Ni et al., 2017), job application (Xiao et al., 2019), and psychology (Siddig and Hines, 2019), among which most are proof of concept. A few interviewbots have been developed on commercial platforms such as Google Dialogflow and IBM Watson Assistant, with the limitation of pre-scripted interviews; thus, they cannot proactively follow up to the user contents.

3 Interview Dataset

Audio recordings of 7,361 interviews are automatically transcribed with speaker identification by the online tool RevAI,[1] where 440 are manually corrected on speaker ID assignment for finetuning and evaluation of our models (Table 1). Each recording contains an average of a ≈15-min long dialogue between an interviewer and an interviewee. The interviews were conducted by 67 professionals in 2018 - 2022. The largest age group of interviewees is 18-years-old with 59.3%, followed by 17-years-old with 29.4%. The male-to-female ratio is 1.2:1. The major country of origin is China with 81.4% followed by Belgium with 10.5%, alongside 37 other countries. Appendix A.1 provides detailed demographics of the interviewees.

4 Dialogue Generation

Figure 1 depicts an overview of our dialogue generation model. Since inputs to the encoder $\mathcal{E}$ and the decoder $\mathcal{D}$ are limited by the total number of tokens that the pretrained language model accepts, *sliding window* (Sec. 4.1) and *context attention* (Sec. 4.2) are proposed to handle long utterances and contexts in the previous utterances, respectively. In addition, *topic storing* is used to remember user-oriented topics brought up during the interview (Sec. 4.3). The input to $\mathcal{E}$ and output of $\mathcal{D}$ include the speaker ID `S1` or `S2` as the first token followed by an utterance from the interviewer or interviewee, respectively. Hyperparameters are finetuned by cross validations.

[1] https://www.rev.ai

	D	U	S1	S2
`TRN`	140	43.8	39.3	64.0
`DEV`	150	45.0	36.2	60.3
`TST`	150	44.3	37.8	61.3
`RAW`	6,921	40.4	41.5	67.6

Table 1: Distributions of our data. D: num of dialogues, U: avg-num of utterances per dialogue, S1/S2: avg-num of tokens per utterance by interviewer/interviewee. `TRN`/`DEV`/`TST`: training/development/evaluation (annotated) sets. `RAW`: unannotated set (auto-transcribed).

4.1 Sliding Window

Let $n = m + e$ be the max-number of tokens that $\mathcal{E}$ and $\mathcal{D}$ accept ($e < m < n$). Every utterance U whose length is greater than n is split into U^1 and U^2 as follows (w_i is the i'th token in U):

$$U^1 = \{w_1, \ldots, w_m, w_{m+1}, \ldots, w_n\}$$
$$U^2 = \{w_{m+1}, \ldots, w_n, w_{n+1}, \ldots, w_{n+m}\} \quad (141)$$

In our case, $n = 128$, $m = 100$, and $e = 28$ such that $n + m = 228$ is sufficiently long enough to handle most utterances based on our stats. $\mathcal{E}$ takes U^1 and U^2 then produces $E^1 = \{e^1_1, \ldots, e^1_n\}$ and $E^2 = \{e^2_{m+1}, \ldots, e^2_{n+m}\}$ where $e^*_i \in \mathbb{R}^{1\times d}$ is the embedding of w_i. Finally, the embedding matrix $E \in \mathbb{R}^{(n+m)\times d}$ of U is created by stacking all of the following embeddings:

$$\{e^1_1, \ldots, \frac{1}{2}\sum_{i=1}^{2}(e^i_{m+1}), \ldots, \frac{1}{2}\sum_{i=1}^{2}(e^i_n), \ldots, e^2_{n+m}\}$$

For utterances whose lengths are less than or equal to n, zero-padding is used to transform $\mathcal{E}$'s output from $\mathbb{R}^{n\times d}$ to $\mathbb{R}^{(n+m)\times d}$.

4.2 Context Attention

Let U_i be the i'th utterance to be generated as output. Let $C \in \mathbb{R}^{\ell\times d}$ be the context matrix stacking the embedding matrices of the previous utterances $\{E_{i-k}, .., E_{i-1}\}$, where k is the number of previous utterances to be considered and $\ell = k(n + m)$. The transpose of C is multiplied by the attention

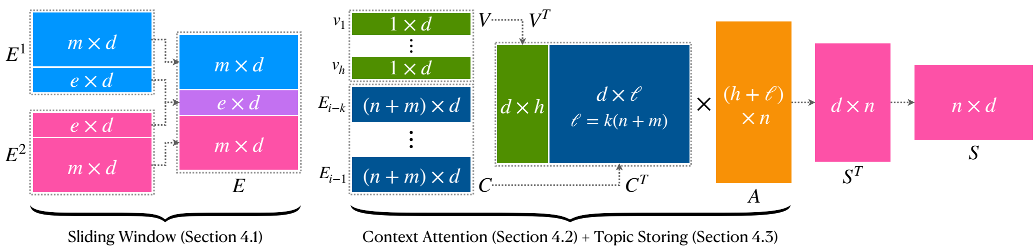

Figure 1: The overview of our dialogue generation model.

matrix $A \in \mathbb{R}^{\ell \times n}$ such that $C^T \cdot A \to S^T \in \mathbb{R}^{d \times n}$. Thus, $S \in \mathbb{R}^{n \times d}$ represents the context summary of $U_{i-k}, .., U_{i-1}$, which is fed into the decoder $\mathcal{D}$.

4.3 Topic Storing

Even with the context attention, the model still has no memory of contexts prior to U_{i-k}, leading it to repeat the same topics that it has already initiated. To overcome this issue, topic storage is introduced to remember key topics derived by the interviewer. Every interview in our data came with 8-16 questions by the interviewer, who used those questions during the interview and thought they led to assess crucial aspects of the interviewee. Our final model considers these questions the "key topics" and dynamically stores them as the dialogue progresses.

Let $Q = \{q_1, .., q_h\}$ be the topical question set. During training, $\mathcal{D}$ learns to generate `Q` instead of `S1` as the first token of the interviewer's utterance that contains any $q_i \in Q$. In addition, it generates `B`/`E` if the interviewer begins/ends the current dialogue with that utterance (Table 3). Any utterance starting with `Q` is encoded by $\mathcal{E}$ that creates the utterance embedding $v_i \in \mathbb{R}^{1 \times d}$. These embeddings get stacked as the interview goes to create the topic matrix $V \in \mathbb{R}^{h \times d}$. If $|Q| < h$, then zero-padding is used to create V (in our case, $h = 16$). Finally, V is stacked with the context matrix C (Sec. 4.2), and $(V \oplus C)^T \in \mathbb{R}^{d \times (h+\ell)}$ is multiplied by the attention matrix $A \in \mathbb{R}^{(h+\ell) \times n}$ to create the transpose of the context summary matrix $S \in \mathbb{R}^{n \times d}$.

5 Experiments

For our experiments, the encoder and the decoder in BlenderBot 1.0 (Roller et al., 2021) are used.[2] Three models are developed as follows:

- BB: *Blender Baseline* that takes only up to 128 tokens in U_{i-1} as context.
- SW: BB + *Sliding Window* (Section 4.1), taking all tokens in U_{i-1} as context.
- CT: SW + *Context Attention* (Section 4.2) + *Topic Storing* (Section 4.3), taking all tokens in $\{U_{i-k}, .., U_{i-1}\}$ as context.

[2]There have been updated versions of BlenderBot introduced (Xu et al., 2022; Shuster et al., 2022). However, we chose the first version for our experiments because we found it to be as effective yet much more efficient than the newer versions.

All models are first trained on `raw` and finetuned on `TRN` in Table 1). To assess real-life performance, 10 interviews are conducted per model, where each interview consists of exactly 30 turns. A qualitative analysis is performed on the top-3 most frequently occurring errors as follows:

- Repetitions: how often it repeats topics already covered in the previous utterances.
- Early Ending (EE): implies to end the interview without covering a sufficient amount of topics.
- Off Topic (OT): how often it makes utterances that are not relevant to the current topic.

Table 2 shows the error analysis results. The repetition rates are significantly reduced as the model gets more advanced. Compared to the baseline, the CT model conducts 3.5 times longer conversations before it attempts to end the interview, while generating twice fewer off-topic utterances, which is very promising. Examples of these error types are provided in Appendix A.2.

Model	Repetitions (R)	Early Ending (EE)	Off Topic (OT)
BB	30.0	13.3	20.0
SW	16.7	23.3	26.7
CT	**6.7**	**46.7**	**10.0**

Table 2: The error analysis of all generation models. R: avg-% of repeated topics, EE: avg-% of the interview conducted before the model attempts to end (higher is better), OT: avg-% of off-topic utterances.

5.1 Static Evaluation

Following previous work, static evaluation is performed on the CT model, where the input is every

Flag	Utterance
B	Hello, and thank you so much for coming today. uh, Let us start by having you just to state your full legal name and also the school that you are currently attending?
S2	Sure, my name is David, and I go to Shanghai high school in China.
Q	Great. um, yeah, let's talk about leadership experiences. Can you give me some specific examples?
S2	Sure, when I was a junior, I was elected to be the Vice President in the Student Union and I organized many activities to help students to enrich their life on campus.
S1	Okay, great. Thank you so much for sharing that. so, um, tell me a little bit about your leadership style. What do you think you bring to the student union?
S2	Sure, I am very open-minded, and with some of my colleagues, we together established a lot of new programs on top of our existing ones. Many of those programs became very popular.
S1	Ok, great. and so, um, what, uh, what would you say has been the most challenging aspect of being a leader?
...	
Q	Great. uh, um, hmm, let' talk about your courses. What courses did you take in high school and which courses do you like most?
S2	Sure, I took Math, English, PE, Chemistry, Art, and so on. I like Math most.
Q	Okay. so, um, tell me a little bit about your, your interest in mathematics.
...	
E	Okay, great. well, thank you very much for sharing that with me. um, and thank you for taking the time out of your day to interview with us.

Table 3: An interview dialogue conducted by our best model (CT in Section 5). S1/S2: interviewer/interviewee (chatbot/human), B/E: beginning/ending utterance (chatbot), Q: topical question (chatbot).

batch of k-utterances and prior topics per interview, and its output is compared to the corresponding human response in TST (Table 1). The average BLEU score is 0.08 and cosine similarity is 0.19, which are low. However, such static evaluation assesses each output independently and obstructs dialogue fluency by artificially inserting human utterances to the model, and thus, does not reveal its capability in conducting long contextualized interviews.

5.2 Real-time Evaluation

The CT model is deployed to an online text-based platform in a public cloud. For real-time evaluation, 5 professional interviewers and 10 students are invited to have conversations with our InterviewBot and give ratings from 1 to 5 to indicate their overall satisfactions. The average dialogue duration is 256 seconds. Almost half of the evaluators are satisfied (Scores 4 and 5) and another 40% indicate positive attitude on the coverage of topics and discussions (Score 3), implying that it performs reasonably well for this realistic setting (Table 4). Overall, with the average score of 3.5, the InterviewBot has shown great potential in applying to practical applications.

Rating Score	5	4	3	2	1	$\sum$
Evaluator Count	3	4	6	1	1	15

Table 4: The rating distribution of the InterviewBot conversations for real-time evaluation. 5: very satisfied, 4: satisfied, 3: neutral, 2: unsatisfied, 1: very unsatisfied.

6 Conclusion

Our InterviewBot is a model-based dialogue system equipped with contextual awareness and topic sensitivity that conducts college admission interviews. Questions covering diverse topics and discussions in extended follow-ups are carried along the conversations, which have been assessed by professional interviewers and student volunteers. The average satisfaction score of 3.5 projects prevailing deployment of the InterviewBot for thousands of college applicants, especially for international students.

With promising future applications, however, the current version of the InterviewBot has two major limitations. First, the early ending in Table 2 still happens, where an ending utterance gets generated after an insufficient amount of turns, in which case, the interview may not cover critical aspects of the applicant. Second, the bot makes good follow-ups to various topics; however, it needs to derive deeper discussions with more details.

In the future work, the main focus is to enrich the follow-up discussions on topics or sub-topics during interactions by training the InterviewBot with more structured topic-flow materials. This task would indirectly alleviate the early ending issue by deepening the discussions on certain topics.

7 Ethics Statement and Broader Impact

This research project is conducted under full compliance of research ethics in the ACL Code of Ethics. The data revealed in this paper are either generalized or generated from models that have no identifying or private information released. The interviewers and interviewees involved in the conversational data used in this paper have agreed to the usage of the data by the data collecting party and its collaborators for research purposes.

The future application of the Interviewbot will impact thousands of students applying to U.S. colleges. While institutions emphasizing on the research environment, the Interviewbot could help students understand the basics of research conduct of ethics, and even admission committees to make decisions on students' readiness in compliance of code of ethics in research, through interview conversations.

However, constant observations of its performance may be enforced to prevent negative impact due to malfunction, biased behaving or unbalanced demographic distributions in learning.

References

Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. arXiv preprint arXiv:2001.09977.

Ali Ahmadvand, Ingyu Choi, Harshita Sahijwani, Justus Schmidt, Mingyang Sun, Sergey Volokhin, Zihao Wang, and Eugene Agichtein. 2018. Emory irisbot: An open-domain conversational bot for personalized information access. Alexa Prize Proceedings.

Parham Amiri and Elena Karahanna. 2022. Chatbot use cases in the covid-19 public health response. Journal of the American Medical Informatics Association, 29(5):1000–1010.

Daniel Baier, Alexandra Rese, Maximilian Röglinger, D Baier, A Rese, and M Röglinger. 2018. Conversational user interfaces for online shops? a categorization of use cases. In Proceedings of the International Conference on Information Systems, ICIS.

Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. PLATO: Pre-trained dialogue generation model with discrete latent variable. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 85–96, Online. Association for Computational Linguistics.

Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhihua Wu, Zhen Guo, Hua Lu, Xinxian Huang, et al. 2021. Plato-xl: Exploring the large-scale pre-training of dialogue generation. arXiv preprint arXiv:2109.09519.

Sam Cunningham-Nelson, Wageeh Boles, Luke Trouton, and Emily Margerison. 2019. A review of chatbots in education: practical steps forward. In 30th Annual Conference for the Australasian Association for Engineering Education (AAEE 2019): Educators Becoming Agents of Change: Innovate, Integrate, Motivate, pages 299–306. Engineers Australia.

Xiangmin Fan, Daren Chao, Zhan Zhang, Dakuo Wang, Xiaohua Li, Feng Tian, et al. 2021. Utilization of self-diagnosis health chatbots in real-world settings: case study. Journal of medical Internet research, 23(1):e19928.

Sarah E Finch, James D Finch, Ali Ahmadvand, Xiangjue Dong, Ruixiang Qi, Harshita Sahijwani, Sergey Volokhin, Zihan Wang, Zihao Wang, Jinho D Choi, et al. 2020. Emora: An inquisitive social chatbot who cares for you. arXiv preprint arXiv:2009.04617.

Vladimir Ilievski, Claudiu Musat, Andreea Hossmann, and Michael Baeriswyl. 2018. Goal-oriented chatbot dialog management bootstrapping with transfer learning. arXiv preprint arXiv:1802.00500.

Bui Thanh Khoa. 2021. The impact of chatbots on the relationship between integrated marketing communication and online purchasing behavior in the frontier market. Jurnal The Messenger, 13(1):19–32.

Soomin Kim, Joonhwan Lee, and Gahgene Gweon. 2019. Comparing data from chatbot and web surveys: Effects of platform and conversational style on survey response quality. In Proceedings of the 2019 CHI Conference, pages 1–12.

Jingyi Li, Michelle X Zhou, Huahai Yang, and Gloria Mark. 2017. Confiding in and listening to virtual agents: The effect of personality. In Proceedings of the 22nd International Conference on Intelligent User Interfaces, pages 275–286.

Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to select knowledge for response generation in dialog systems. arXiv preprint arXiv:1902.04911.

Shikib Mehri, Evgeniia Razumovskaia, Tiancheng Zhao, and Maxine Eskenazi. 2019. Pretraining methods for dialog context representation learning. arXiv preprint arXiv:1906.00414.

Rashid Minhas, Camilla Elphick, and Julia Shaw. 2022. Protecting victim and witness statement: examining the effectiveness of a chatbot that uses artificial intelligence and a cognitive interview. AI & SOCIETY, 37(1):265–281.

Lin Ni, Chenhao Lu, Niu Liu, and Jiamou Liu. 2017. Mandy: Towards a smart primary care chatbot application. In International symposium on knowledge and systems sciences, pages 38–52. Springer.

Eliza Nichifor, Adrian Trifan, and Elena Mihaela Nechifor. 2021. Artificial intelligence in electronic commerce: Basic chatbots and the consumer journey. Amfiteatru Economic, 23(56):87–101.

Chinedu Wilfred Okonkwo and Abejide Ade-Ibijola. 2021. Chatbots applications in education: A systematic review. Computers and Education: Artificial Intelligence, 2:100033.

Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. Recipes for building an open-domain chatbot. In Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, pages 300–325, Online. Association for Computational Linguistics.

Zeineb Safi, Alaa Abd-Alrazaq, Mohamed Khalifa, Mowafa Househ, et al. 2020. Technical aspects of developing chatbots for medical applications: scoping review. Journal of medical Internet research, 22(12):e19127.

Iulian Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 31.

Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, William Ngan, Spencer Poff, Naman Goyal, Arthur Szlam, Y-Lan Boureau, Melanie Kambadur, and Jason Weston. 2022. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage.

Abubakr Siddig and Andrew Hines. 2019. A psychologist chatbot developing experience. In AICS, pages 200–211.

Zihao Wang, Ali Ahmadvand, Jason Ingyu Choi, Payam Karisani, and Eugene Agichtein. 2017. Emersonbot: Information-focused conversational ai emory university at the alexa prize 2017 challenge. 1st Proceeding of Alexa Prize.

Ziang Xiao, Michelle X Zhou, and Wat-Tat Fu. 2019. Who should be my teammates: Using a conversational agent to understand individuals and help teaming. In Proceedings of the 24th International Conference on Intelligent User Interfaces, pages 437–447.

Jing Xu, Arthur Szlam, and Jason Weston. 2022. Beyond goldfish memory: Long-term open-domain conversation. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 5180–5197, Dublin, Ireland. Association for Computational Linguistics.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DIALOGPT : Large-scale generative pre-training for conversational response generation. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pages 270–278, Online. Association for Computational Linguistics.

Hao Zhou, Pei Ke, Zheng Zhang, Yuxian Gu, Yinhe Zheng, Chujie Zheng, Yida Wang, Chen Henry Wu, Hao Sun, Xiaocong Yang, et al. 2021. Eva: An open-domain chinese dialogue system with large-scale generative pre-training. arXiv preprint arXiv:2108.01547.

A Appendix

A.1 Interviewee Demographics

We summarize the demographics of the interviewees in this section. Table 2 shows the distribution of the ages of applicants. Most interviewees are between 17 to 19, which is an accurate reflection of the ages of high school students applying to colleges.

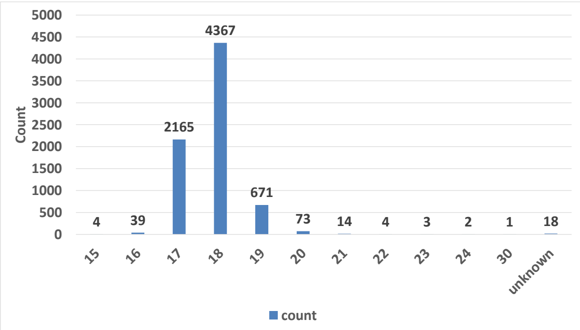

Figure 2: The interviewee age demographics.

The distribution of countries of origin of applicants is shown in Figure 3. There are 38 countries in total. The majority of applicants come from China. Other major countries are Belgium, Bangladesh, Canada, India and Belarus.

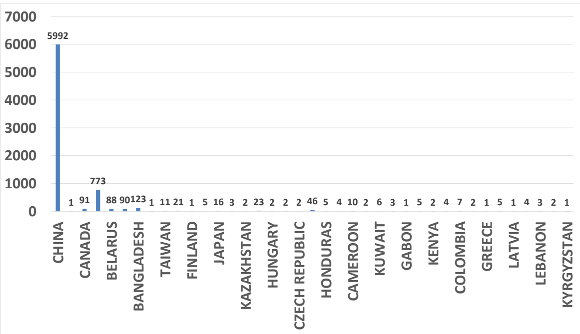

Figure 3: The interviewee country demographics.

The gender distribution of applicants is shown in Figure 4. The numbers of male and female applicants are close, with exclusion of applicants not providing gender information.

A.2 Examples of Generation Limitations

The following are examples for each type of generation errors (underlined). Early ending prevents the Interviewbot from covering critical aspects of applicants and further discussion. The Interviewbot cannot name mentions properly, either skipping over or generating wrong entities in the conversations. Random Generation generates random words

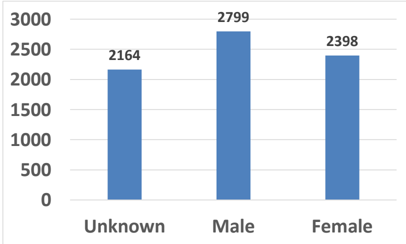

Figure 4: The interviewee gender demographics.

and punctuation, and concatenates them incorrectly. Not following up on previous topics hinders the Interviewbot from becoming proactively engaged with interviewees.

- Early Ending
  - InterviewBot: Could you please state your name?
  - Interviewee: Sure. My name is David.
  - ... 4 more turns ...
  - InterviewBot: Thank you and best wishes to your applications.
- Repetition
  - InterviewBot: Now why, why do you think receiving an education is important?
  - ... more turns ...
  - InterviewBot: Why do you think people value education so much?
- Off Topic
  - Interviewee: Sure, I really like Mathemathcs.
  - InterviewBot: What do you do for fun?