

Training Conversational Agents with Generative Conversational Networks

Yen-Ting Lin, Alexandros Papangelis, Seokhwan Kim, Dilek Hakkani-Tür
{ytl, papangea, seokhwk, hakkanit}@amazon.com

Abstract

Rich, open-domain textual data available on the web resulted in great advancements for language processing. However, while that data may be suitable for language processing tasks, they are mostly non-conversational, lacking many phenomena that appear in human interactions and this is one of the reasons why we still have many unsolved challenges in conversational AI. In this work, we attempt to address this by using Generative Conversational Networks to automatically generate data and train social conversational agents. We evaluate our approach on TopicalChat with automatic metrics and human evaluators, showing that with 10% of seed data it performs close to the baseline that uses 100% of the data.

1 Introduction

Conversational Artificial Intelligence (AI) has progressed a lot in the recent past, partly due to advances in large pre-trained language models (LM) and partly due to commercial conversational agents (Alexa, Siri, Cortana, Google Assistant, and others). It is evident, however, that many challenges still remain, such as handling idioms, humour, expressing empathy, building and maintaining rapport, and so on. One big factor for this is the lack of large and rich conversational data that include these complex aspects of human communication. While the research community is making great efforts in collecting such data, these are still small compared to the amount of data needed to train deep networks. In this work, we take a first step into automatically generating conversational data using Generative Conversational Networks (GCN) (Papangelis et al., 2021), hoping that in the future we will be able to generate many of the aforementioned phenomena. GCN learns how to generate labelled, diverse, and targeted data that are optimised with Reinforcement Learning (RL) to maximise task performance (more details in the following section).

There are many recent studies on data augmentation, but most of them are geared towards language processing tasks rather than training conversational agents. Due to lack of space we only mention a few. PROTODA

(Kumar et al., 2021) uses prototypical networks to augment data for intent classification while GenSF (Mehri and Eskenazi, 2021) uses DialoGPT (Zhang et al., 2020) for zero-shot slot tagging; DINO (Schick and Schütze, 2021) uses pre-trained language models (PLM) to generate data for semantic textual similarity; (Campagna et al., 2020) focus on zero-shot dialogue state tracking and use an abstract dialogue model to generate data and improve performance. Soloist (Peng et al., 2020) uses a PLM fine-tuned on large dialogue corpora and is designed for transactional (goal-oriented) dialogues. (Mohanapatra et al., 2020) use PLM to train user simulators from crowd-generated conversations and their instructions, for transactional conversations. (Lin et al., 2021) train domain-independent user simulators for transactional dialogues.

2 Modeling Multi-Turn Conversations

2.1 Generative Conversational Networks

GCN (Papangelis et al., 2021) consist of two models in a meta-learning architecture: a data generator and a learner. The data generator creates a labelled dataset that is used to train the learner (conversational agent in our case) in a supervised fashion. The learner is then evaluated on an external validation set and its performance is used as a proxy of the quality of the dataset. This quality measure is used as a reward in a RL setup that trains the generator. This way, the generator learns to create data that lead the learner to perform well on the validation set. Both models can be pre-trained with seed data, if available, and paired with reward estimation, GCN can be used for continuous learning from user feedback. This approach has been proven to work well for intent detection and slot tagging for goal-oriented conversations and we here apply it to train social conversational agents.

2.2 Generator training

We use DialoGPT (Zhang et al., 2020) to frame the multi-turn conversation generation as autoregressive language modeling. We define a multi-turn conversation as a list of utterances: T_1, T_2, \dots, T_N where T_i is the utterance in turn i , and N is the number of turns in the conversation. Each utterance T is composed of tokens

Table 1: Automatic and human evaluation results. Human evaluators rate responses with the same conversation context on a scale of 1 to 5. BScore stands for BERTScore. Bold indicates statistical significance (t-test assuming unequal variance).

Model	Automatic Evaluation			Human Evaluation			
	BLEU	Rouge(1/2/L)	BScore	Engaging.	Fluency	Relevance	Overall
Ground Truth	-	-	-	3.85 ± 1.04	4.55 ± 0.85	3.77 ± 1.05	4.06 ± 1.04
BART-100	3.1	20.3/6.1/17.8	0.861	3.80 ± 1.06	4.58 ± 0.85	3.68 ± 1.10	4.02 ± 1.08
GCN Learner	1.3	15.8/2.7/13.6	0.851	3.79 ± 1.06	4.49 ± 0.95	3.58 ± 1.14	3.96 ± 1.12
BART-10	2.0	18.5/4.2/16.0	0.858	3.63 ± 1.04	4.50 ± 0.92	3.62 ± 1.12	3.92 ± 1.11
GCN-RL Learner	1.1	15.0/2.1/12.6	0.850	3.70 ± 1.10	4.47 ± 0.99	3.47 ± 1.17	3.88 ± 1.17

x_1, x_2, \dots, x_n , where n is the number of tokens in the utterance.

In our case, we use seed data D to aid training, that are split into training and validation. The generator is given the first three turns T_1, T_2 , and T_3 sampled from the seed data as a prompt and generates the remaining turns T'_4, T'_5, \dots, T'_N in a self-playing setting. This is done for a number of batches, resulting in a new dataset D' . A new learner is then spawned and trained on this dataset and evaluated on the seed validation set. Following (Ziegler et al., 2019) and (Papangelis et al., 2021), we use PPO (Schulman et al., 2017) to train the generator using the learner’s performance on the validation set as a reward, and iterate the data generation - learner training process until performance does not change significantly (within a small tolerance). We use a weighted sum of BLEU (Papineni et al., 2002), Rouge-1 (Lin, 2004), and BERTScore (Zhang* et al., 2020) as the reward for the generator RL training.¹ After the training phase, we create a final dataset D' and use that and the seed training data D to train a final learner that is used for evaluation.

3 Evaluation

We use TopicalChat (Gopalakrishnan et al., 2019) in our experiments, sampling 10% of it as seed for GCN. We use DialoGPT (Zhang et al., 2020) and BART (Lewis et al.) as initial weights for the generator and the learner, respectively. As mentioned in the previous section, our GCN learner was trained on both seed data and synthetic data, created by the GCN generator. We compare the performance of the GCN learner against the ground truth and against three baselines: BART trained with the same seed data (BART-10), BART trained with the complete training set of TopicalChat (BART-100), and a GCN learner trained on seed and synthetic data but without updating the generator via RL (GCN-RL).

Automatic evaluation. We calculate the reference-based automatic evaluation metrics on the TopicalChat test set (frequent) (Table 1, left). The learner agent trained on the seed and synthetic data by the generator performs lower on most metrics. However, due to the intrinsic one-to-many property of conversation,

¹The weights are used to normalize the metrics to a scale of 0 to 1 and determined empirically: 0.1, 0.01, 0.95, for BLEU, Rouge-1 and BERTScore, respectively.

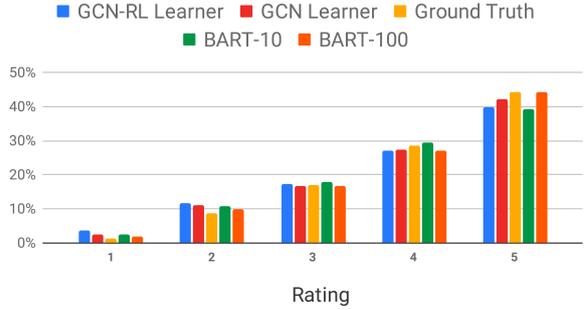


Figure 1: Distribution of ratings from human evaluators

reference-based metrics may not correlate well with human judgments; our generated conversation may be appropriate for the dialogue context but different from the reference responses. For this reason, we also conduct human evaluation.

Human evaluation. We conduct a study where human evaluators rate the output of the GCN learner, the baselines, and the ground truth. Specifically, we rate how engaging, fluent, and relevant each response from each model is, on a scale from 1 to 5. We generate 1000 samples for each condition using the same context and make sure we have 3 ratings per sample per condition. Table 1 (right) shows the results of the evaluation, where we see that the GCN learner outperforms the BART-10 and GCN-RL models and is close to BART-100’s performance. All models are outperformed by the ground truth, which may be due to the size of our models or the number of training iterations. Figure 1 shows the distribution of ratings for all conditions.

4 Discussion

We present GCN for multi-turn conversations, where our agent can learn to generate the data it needs for social conversations. The lower results in the reference-based metrics (Table 1, left) indicate that GCN generates new data; paired with the good scores in the human evaluation (Table 1, right), we see that the generated data is also useful for the task at hand, a result in line with the results on intent detection and slot tagging in (Papangelis et al., 2021). As future work, we are conducting thorough analyses and evaluations of the generated data, exploring more generator the learner architectures, and working towards knowledge-grounded conversations.

References

- Giovanni Campagna, Agata Foryciarz, Mehrad Moradshahi, and Monica Lam. 2020. Zero-shot transfer learning with synthesized data for multi-domain dialogue state tracking. In *Proceedings of the 58th Association for Computational Linguistics*.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, Dilek Hakkani-Tür, and Amazon Alexa AI. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895.
- Manoj Kumar, Varun Kumar, Hadrien Glaude, Cyprien de Lichy, Aman Alok, and Rahul Gupta. 2021. Proton: Efficient transfer learning for few-shot intent classification. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 966–972. IEEE.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension.
- Chin-Yew Lin. 2004. *ROUGE: A package for automatic evaluation of summaries*. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Hsien-chin Lin, Nurul Lubis, Songbo Hu, Carel van Niekerk, Christian Geishausser, Michael Heck, Shutong Feng, and Milica Gašić. 2021. Domain-independent user simulation with transformers for task-oriented dialogue systems. *arXiv preprint arXiv:2106.08838*.
- Shikib Mehri and Maxine Eskenazi. 2021. Gensf: Simultaneous adaptation of generative pre-trained models and slot filling. *SIGDial*.
- Biswesh Mohapatra, Gaurav Pandey, Danish Contractor, and Sachindra Joshi. 2020. Simulated chats for task-oriented dialog: Learning to generate conversations from instructions. *arXiv e-prints*, pages arXiv–2010.
- Alexandros Papangelis, Karthik Gopalakrishnan, Aishwarya Padmakumar, Seokhwan Kim, Gokhan Tur, and Dilek Hakkani-Tur. 2021. Generative conversational networks. In *SIGDial*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. *Bleu: a method for automatic evaluation of machine translation*. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2020. Soloist: Building task bots at scale with transfer learning and machine teaching. *arXiv preprint arXiv:2005.05298*.
- Timo Schick and Hinrich Schütze. 2021. Generating datasets with pretrained language models. *arXiv preprint arXiv:2104.07540*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with bert*. In *International Conference on Learning Representations*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B Dolan. 2020. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Association for Computational Linguistics: System Demonstrations*, pages 270–278.
- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.