

SMRT Chatbots: Improving Non-Task-Oriented Dialog with Simulated Multiple Reference Training

Huda Khayrallah
Johns Hopkins University
huda@jhu.edu

João Sedoc
New York University
jsedoc@stern.nyu.edu

Abstract

Non-task-oriented dialog models suffer from poor quality and non-diverse responses. To overcome limited conversational data, we apply Simulated Multiple Reference Training (SMRT; Khayrallah et al., 2020), and use a paraphraser to simulate multiple responses per training prompt. We find SMRT improves over a strong Transformer baseline as measured by human and automatic quality scores and lexical diversity. We also find SMRT is comparable to pretraining in human evaluation quality, and outperforms pretraining on automatic quality and lexical diversity, without requiring related-domain dialog data.

1 Introduction

Non-task-oriented dialog is a low-resource NLP task. While large and noisy related corpora exist (e.g. movie subtitles, social media, and irlogs; Serban et al., 2018), the publicly-released curated corpora are small. Serban et al. note that smaller corpora have lower lexical diversity and topic coverage, leading to models with poor quality non-diverse responses. Pretraining on larger data may improve performance, but requires a large dialog corpus in the right language and related domain.

We leverage Simulated Multiple Reference Training (SMRT; Khayrallah et al., 2020) to overcome sparse dialog data. SMRT uses a word-level knowledge distillation-inspired objective and a paraphraser to simulate multiple references per training example. Khayrallah et al. introduce SMRT for machine translation (MT) and simulate training on *all* translations for a source sentence, assuming: (1) all paraphrases of a target are translations of the source; and (2) all translations of the source are paraphrases of the target. (1) is true for dialog, but (2) is not—valid chatbot responses vary in meaning. SMRT captures *syntactic* diversity though it cannot represent all *semantic* variations.

prompt:	Study, study, study. I want to learn a lot.
response:	You are going to take courses?

paraphrases:

Table 1: A DailyDialog training pair and paraphrases. The tree of paraphrases includes some possible paraphrases of the original prompt, a **sampled path** and some of the other *tokens also considered in the training objective*.

We apply SMRT to chatbots and find that it: (1) improves human and automatic quality scores; (2) improves lexical diversity; (3) performs as well as pretraining in human evaluation with better performance on automatic measures of diversity and quality.

2 Method

We model the non-task-oriented dialog system (chatbot) task as conditional language modeling. These models are typically trained using Negative Log Likelihood (NLL) with respect to a single reference. An alternative approach is Knowledge Distillation (Hinton et al., 2015; Kim and Rush, 2016) which assumes access to a teacher distribution ($q(y|x)$) and minimizes the cross entropy with the teacher’s probability distribution.

Simulated Multiple Reference Training

SMRT is structured similarly to word-level Knowledge Distillation, but uses a paraphraser as the teacher distribution ($q(y'|y)$). The paraphraser conditions on the reference y (rather than the

source x) and generates a paraphrase y' . Additionally, SMRT *samples* a new paraphrase of the reference every epoch. The SMRT training objective for the i^{th} target word in the reference y , given the prompt x , with a target vocabulary \mathcal{V} is:

$$\mathcal{L}_{\text{SMRT}} = - \sum_{v \in \mathcal{V}} \left[p_{\text{PARAPHRASER}}(y'_i = v \mid y, y'_{j < i}) \times \log(p_{\text{CHATBOT}}(y'_i = v \mid x, y'_{j < i})) \right]$$

The paraphraser and chatbot each condition on the previously sampled paraphrase tokens ($y'_{j < i}$).

3 Experimental Setup

3.1 Dialog models

We train Transformer (Vaswani et al., 2017) chatbots in FAIRSEQ using parameters from the FLORES¹ benchmark for low-resource MT (Guzmán et al., 2019) for both a standard NLL baseline and SMRT.² Following Khayrallah et al. (2020), we sample from the 100 highest probability tokens from the paraphraser distribution at each time-step (Fan et al., 2018).

We train and evaluate on DailyDialog (Li et al., 2017), a high quality corpus with multiple references for evaluation. We train on the $\sim 80,000$ turns of English-learners practicing ‘daily dialogues’ in various contexts, e.g., chatting about vacation or food.

See Appendix A for full details for replication.

3.2 Paraphraser

We use the state-of-the-art PRISM multilingual paraphraser Thompson and Post (2020a,b).³ It is trained as a multilingual MT model on ~ 100 million sentence pairs in 39 languages. Paraphrasing is treated as zero-shot translation (e.g., English to English).

3.3 Evaluation Protocols

Human Evaluation We use Amazon Mechanical Turk to collect human judgments. For every HIT we display a prompt and two responses; the worker indicates their preferred response (or tie). Following Baheti et al. (2018), we employ the pairwise bootstrap test (Efron and Tibshirani, 1994) and report statistical significance at the 95% confidence level.

¹github.com/facebookresearch/flores

²github.com/thompsonb/fairseq-smrt

³github.com/thompsonb/prism

Automatic Quality Evaluation We use MULTIREFEVAL for DailyDialog (Gupta et al., 2019). In §4 we report METEOR, ROUGE-L, and GREEDY MATCH for the original and multiple references. See Appendix B for all 14 metrics. For reading ease we report metrics scaled 0 to 100.

Automatic Diversity Evaluation To measure lexical diversity, we use the type/token ratio of unigrams, bigrams, and trigrams (Li et al., 2016).

4 Results

SMRT is preferred over the baseline system in human evaluation, as shown in Table 2. It outperforms the baseline in automatic quality too: see Table 3. Our *baseline* outperforms nearly all systems in Gupta et al. (2019) for these metrics,⁴ suggesting it is a strong baseline. SMRT has higher lexical diversity than the baseline, though not as high as the human reference response (Table 4).

baseline	SMRT	tie
35.8%	43.5%	20.6%

Table 2: Human preference judgments. The output of SMRT is preferred over the baseline system. This preference is statistically significant at the 95% confidence level.

	Multi-Ref			Single-Ref		
	M	R	GM	M	R	GM
baseline	12.8	34.0	76.9	6.9	20.9	71.2
SMRT	13.8	36.1	77.7	8.1	24.0	72.5

Table 3: SMRT outperforms the baseline on METEOR (M), ROUGE (R), and GREEDY MATCH (GM) for single and multi-reference scoring.

	1-grams	2-grams	3-grams
human reference	6.3%	38.9%	72.7%
baseline	2.9%	11.6%	20.4%
SMRT	3.8%	17.4%	32.2%

Table 4: Type/Token ratio for the baseline and SMRT. SMRT has higher lexical diversity than the baseline.

⁴Except CVAE on single reference METEOR.

5 Analysis

SMRT outperforms a strong baseline; here we analyze it in additional settings: pretraining and MMI.

5.1 Pretraining

Pretraining is another way of incorporating auxiliary data in the model. We pretrain on the OpenSubtitles corpus (OS; Lison and Tiedemann, 2016),⁵ which consists of ~ 200 million turns from movie subtitles. Similar to DailyDialog, it consists of conversational data on a variety of topics. After pretraining on OS, we fine-tune on DailyDialog.

Results In the human evaluation (Table 5), SMRT performs comparably to baseline pretraining. In automatic evaluation (Table 6), SMRT outperforms pretraining. We combine SMRT with pretraining⁶ and find that this again performs comparably to baseline pretraining in human evaluation, and pretraining with SMRT performs better in the automatic evaluation. Finally, we compare SMRT with and without pretraining, and find with pretraining is preferred in human evaluation, while they perform similarly on the automatic metrics.

Pretraining improves the NLL baseline’s diversity, but SMRT’s diversity is still better. Combining SMRT with pretraining improves diversity compared to pretraining alone: see Table 7.

Overall, SMRT performs on par with pretraining in terms of human evaluation of quality, with better diversity and automatic metrics of quality.⁷

Discussion It can be hard to find *dialog* corpora that are *large*, *domain relevant*, and *in-language*.

Unlike pretraining, SMRT incorporates non-dialog data. PRISM was trained to translate, and leveraged as a paraphrase model using zero-shot translation. It is not trained to generate dialog, yet we still leverage it to improve a chatbot.

The paraphraser is trained on less data (~ 100 million sentences pairs, with ~ 17 million English sentences) than is used for OpenSubtitles pretraining (~ 200 million turns—all in English), thus competitive performance is not a result of more data.

PRISM was trained on formal text: Wikipedia, news (Global Voices, and SETimes) parliamentary proceedings (EuroParl), and documents (United

⁵[opendatacommons.org](https://opendatacommons.org/licenses/by/4.0/)

⁶We pretrain with NLL then fine-tune with SMRT.

⁷We hypothesize a conversation-level evaluation would further highlight the strengths of SMRT, by allowing for human judgments of diversity but that is beyond our budget.

M1	M2	M1	M2	tie
PT + baseline	SMRT	31.6%	32.7%	35.8%
PT + baseline	PT + SMRT	34.9%	36.3%	28.8%
SMRT	PT + SMRT	32.3%	37.4%	30.2%

Table 5: Pretraining (PT) human preferences. SMRT and NLL pretraining perform comparably, adding SMRT to pretraining is comparable to NLL pretraining, and pretrained SMRT outperforms SMRT alone. None of the preferences between models in this table are statistically significant at the 95% confidence level.

	Multi-Ref			Single-Ref		
	M	R	GM	M	R	GM
baseline	12.8	34.0	76.9	6.9	20.9	71.2
SMRT	13.8	36.1	77.7	8.1	24.0	72.5
PT + baseline	13.6	35.8	77.5	7.1	21.7	71.5
PT + SMRT	13.9	36.6	77.6	7.9	23.7	72.3

Table 6: SMRT alone outperforms baseline pretraining (PT) on METEOR (M), ROUGE (R), and GREEDY MATCH (GM) for single and multi-reference scoring.

	1-grams	2-grams	3-grams
human reference	6.3%	38.9%	72.7%
baseline	2.9%	11.6%	20.4%
SMRT	3.8%	17.4%	32.2%
PT + baseline	3.5%	14.2%	24.6%
PT + SMRT	4.1%	17.9%	31.8%

Table 7: Type/Token ratios for pretraining (PT). Pretraining the baseline model increases lexical diversity. SMRT and PT SMRT outperform the baseline and PT + baseline in lexical diversity by similar amounts.

Nations), not casual or conversational data. So SMRT works in spite of a domain mismatch with the paraphraser, which is especially useful when there is no domain-relevant dialog pretraining data.

While dialog research currently focuses on English, PRISM has coverage of 39 languages meaning SMRT can be applied in other languages.

DailyDialog is well matched to OpenSubtitles, and yet SMRT performs as well as pretraining on OS. This suggests SMRT is effective at leveraging non-dialog data, which is crucial when no in-domain, in-language dialog data is available.

prompt:	(a) Listen, Karen, I need your help. I don't know anyone here yet.	(b) I try my best to eat only fruits, vegetables, and chicken.
baseline:	What's wrong with me?	I don't know what to eat.
SMRT:	I am glad to help you, what is wrong?	What kind of fruit do you prefer?
PT + baseline:	I'm sorry to hear that.	What kind of vegetables do you like?
PT + SMRT:	I am sorry, Karen, I can't help you.	What kind of food do you eat?

Table 8: Two example evaluation prompts, with various system outputs. SMRT outputs are better than the baseline.

baseline + MMI	SMRT + MMI	tie
34.7%	38.4%	26.9%

Table 9: Human preferences judgments. When comparing models with MMI decoding, SMRT is preferred. This preference is statistically significant at the 95% confidence level.

	Multi-Ref			Single-Ref		
	M	R	GM	M	R	GM
baseline	12.8	34.0	76.9	6.9	20.9	71.2
SMRT	13.8	36.1	77.7	8.1	24.0	72.5
baseline + MMI	12.7	33.5	76.7	6.6	20.1	70.8
SMRT + MMI	13.7	35.8	77.6	7.9	23.5	72.3

Table 10: MMI degrades both baseline and SMRT performance on METEOR (M), ROUGE (R), and GREEDY MATCH (GM) for single and multi-ref scoring. SMRT + MMI still outperforms baseline + MMI.

5.2 MMI

Maximum Mutual Information (MMI) decoding, $(1-\lambda) \log p(y|x) + \lambda \log p(x|y)$, is commonly used in dialog to increase response diversity (Li et al., 2016), however we did not find it helpful in our experiments. Following MMI-bidi, we rerank a 100-best list with a reverse model.⁸ When comparing both models with MMI, we find human prefer SMRT to the baseline, see Table 9. MMI degrades automatic measures of quality (Table 10) and diversity (Table 11) of both the baseline and SMRT models compared to standard decoding. The quality degradation is similar for both, but the degradation in diversity is more pronounced for SMRT.

5.3 Examples

For a training pair and paraphrased responses, see Table 1. SMRT decreases that number of dull and

⁸We sweep λ of 0.1, 0.2, 0.3, 0.4, 0.5. 0.1 performs best on the automatic quality metrics, so we use that for analysis.

	1-grams	2-grams	3-grams
human reference	6.3%	38.9%	72.7%
baseline	2.9%	11.6%	20.4%
SMRT	3.8%	17.4%	32.2%
baseline + MMI	2.9%	10.1%	17.5%
SMRT + MMI	3.6%	15.7%	28.6%

Table 11: Type/Token ratio comparison with MMI. MMI degrades lexical diversity for both methods.

off-topic answers, see Table 8. In prompt (a), the baseline is off-topic. Pretraining expresses sympathy, but is unhelpful. SMRT and pretrained SMRT give relevant responses. In (b), the baseline has the right general topic but is a poor response. Both SMRT variants and the pretrained baseline respond well. For more examples, see Appendix C.

6 Related work

Paraphrasing Neural paraphrasing is actively improving (Wieting et al., 2017, 2019; Li et al., 2018; Wieting and Gimpel, 2018; Hu et al., 2019a,b,c; Thompson and Post, 2020a,b); we expect future improved paraphrasers will improve SMRT.

Simulated Multiple Reference Training Khayrallah et al. use a paraphraser trained on PARABANK2 (an English paraphrase dataset created using back-translation; Hu et al., 2019c) for SMRT. Thompson and Post (2020a) introduced PRISM; they show PRISM outperforms PARABANK2 and that PARABANK2 is biased against producing the input as the paraphrase, while PRISM is not. Thus, while Khayrallah et al. use a SMRT objective with 50% probability and standard NLL otherwise, we only use SMRT.

Paraphrastic Dialog Augmentation There is little work on data augmentation for chatbots, but there is a variety of work on *task-oriented dialog* augmentation. Kurata et al. (2016) use self-training with noisy decoding to create additional target side data. Using a seq-to-seq model, Hou et al. (2018) generate diverse lexical and syntactic alternatives within a semantic frame. Gao et al. (2020) jointly train a paraphrase model and a response generation model using dialog data. These works generate paraphrases using dialog training data; in contrast, we leverage additional corpora. Niu and Bansal (2018, 2019) include paraphrasing as one of several augmentation policies, using external paraphrase data. For other NLP tasks, Hu et al. (2019a) perform paraphrastic data augmentation for natural language inference, question answering and machine translation.

Diversity A variety of decoding approaches address diversity in chatbot output, including: MMI (Li et al., 2016), various random sampling (e.g. Fan et al. (2018)), modified beam search (Cho, 2016; Vijayakumar et al., 2016; Tam et al., 2019; Kulikov et al., 2019) and over-generating and clustering post-decoding (Ippolito et al., 2019). In this work we improve training, which can be combined with any decoding strategy. Zhang et al. (2018) and Xu et al. (2018) use adversarial training to encourage diversity. Ippolito et al. (2019) note such methods are ‘task-specific and difficult to implement.’ SMRT is general with simple public code. Jiang and de Rijke (2018) connect the low-diversity problem to overconfidence in the model distribution. Since it trains toward a distribution rather than a 1-hot vector, SMRT may have more reasonable confidence levels.

7 Conclusion

SMRT improves upon a strong Transformer baseline in quality and diversity. It also has human evaluation quality comparable to pretraining, with better automatic quality and lexical diversity. This method, which works even in settings where pretraining is impractical due to a lack of *in-domain same language dialog* data, has a high potential for impact in creating chatbots for more languages.

Acknowledgments

We thank Patrick Xia, Claire Daniele, Nathaniel Weir, Carlos Aguirre, and additional anonymous proofreaders for their helpful comments and feedback on the paper. We additionally thank the reviewers for their insightful comments.

This work was partially supported by the Amazon AWS Cloud Credits for Research program. This work was supported in part by DARPA KAIROS (FA8750-19-2-0034). The views and conclusions contained in this work are those of the authors and should not be interpreted as representing official policies or endorsements by DARPA or the U.S. Government.

References

- Ashutosh Baheti, Alan Ritter, Jiwei Li, and Bill Dolan. 2018. [Generating more interesting responses in neural conversation models with distributional constraints](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3970–3980, Brussels, Belgium. Association for Computational Linguistics.
- Kyunghyun Cho. 2016. [Noisy parallel approximate decoding for conditional recurrent language model](#).
- Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. CRC press.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Gabriel Forgues, Joelle Pineau, Jean-Marie Larchevêque, and Réal Tremblay. 2014. [Bootstrapping dialog systems with word embeddings](#). In *Modern Machine Learning and Natural Language Processing at NeurIPS*.
- Silin Gao, Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. [Paraphrase augmented task-oriented dialog generation](#).
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey Bigham. 2019. [Investigating evaluation of open-domain dialogue systems with human generated multiple references](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 379–391, Stockholm, Sweden. Association for Computational Linguistics.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. [The](#)

- FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111, Hong Kong, China. Association for Computational Linguistics.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#).
- Yutai Hou, Yijia Liu, Wanxiang Che, and Ting Liu. 2018. [Sequence-to-sequence data augmentation for dialogue language understanding](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1234–1245, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- J. Edward Hu, Huda Khayrallah, Ryan Culkin, Patrick Xia, Tongfei Chen, Matt Post, and Benjamin Van Durme. 2019a. [Improved lexically constrained decoding for translation and monolingual rewriting](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 839–850, Minneapolis, Minnesota. Association for Computational Linguistics.
- J. Edward Hu, Rachel Rudinger, Matt Post, and Benjamin Van Durme. 2019b. [ParaBank: Monolingual bitext generation and sentential paraphrasing via lexically-constrained neural machine translation](#). In *Proceedings of AAAI*.
- J. Edward Hu, Abhinav Singh, Nils Holzenberger, Matt Post, and Benjamin Van Durme. 2019c. [Large-scale, diverse, paraphrastic bitexts via sampling and clustering](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 44–54, Hong Kong, China. Association for Computational Linguistics.
- Daphne Ippolito, Reno Kriz, João Sedoc, Maria Kustikova, and Chris Callison-Burch. 2019. [Comparison of diverse decoding methods from conditional language models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3752–3762, Florence, Italy. Association for Computational Linguistics.
- Shaojie Jiang and Maarten de Rijke. 2018. [Why are sequence-to-sequence models so dull? understanding the low-diversity problem of chatbots](#). In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 81–86, Brussels, Belgium. Association for Computational Linguistics.
- Huda Khayrallah, Brian Thompson, Matt Post, and Philipp Koehn. 2020. Simulated multiple reference training improves low-resource machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Online. Association for Computational Linguistics.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Skip-thought vectors](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3294–3302. Curran Associates, Inc.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Iliia Kulikov, Alexander Miller, Kyunghyun Cho, and Jason Weston. 2019. [Importance of search and evaluation strategies in neural dialogue modeling](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 76–87, Tokyo, Japan. Association for Computational Linguistics.
- Gakuto Kurata, Bing Xiang, and Bowen Zhou. 2016. [Labeled data generation with encoder-decoder lstm for semantic slot filling](#). In *Interspeech 2016*, pages 725–729.
- Alon Lavie and Abhaya Agarwal. 2007. [METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 228–231, Prague, Czech Republic. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. [Paraphrase generation with deep reinforcement learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language*

- Processing*, pages 3865–3878, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Alexander H. Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, and Jason Weston. 2017. [Parlai: A dialog research software platform](#).
- Tong Niu and Mohit Bansal. 2018. [Adversarial over-sensitivity and over-stability strategies for dialogue models](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 486–496, Brussels, Belgium. Association for Computational Linguistics.
- Tong Niu and Mohit Bansal. 2019. [Automatically learning data augmentation policies for dialogue tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1317–1323, Hong Kong, China. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Vasile Rus and Mihai Lintean. 2012. [A comparison of greedy and optimal assessment of natural language student input using word-to-word similarity metrics](#). In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 157–162, Montréal, Canada. Association for Computational Linguistics.
- Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2018. A survey of available corpora for building data-driven dialogue systems: The journal version. *Dialogue & Discourse*, 9(1):1–49.
- Shikhar Sharma, Layla El Asri, Hannes Schulz, and Jeremie Zumer. 2017. [Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation](#). *CoRR*, abs/1706.09799.
- Yik-Cheung Tam, Jiachen Ding, Cheng Niu, , and Jie Zhou. 2019. Cluster-based beam search for pointer-generator chatbot grounded by knowledge. In *Dialog System Technology Challenges 7 at AAI*.
- Brian Thompson and Matt Post. 2020a. Automatic machine translation evaluation in many languages via zero-shot paraphrasing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Online. Association for Computational Linguistics.
- Brian Thompson and Matt Post. 2020b. Paraphrase generation as zero-shot multilingual translation: Disentangling semantic similarity from lexical and syntactic diversity. In *Proceedings of the Fifth Conference on Machine Translation (Volume 1: Research Papers)*, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R. Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. [Diverse beam search: Decoding diverse solutions from neural sequence models](#).
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- John Wieting, Kevin Gimpel, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. [Simple and effective paraphrastic similarity from parallel translations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4602–4608, Florence, Italy. Association for Computational Linguistics.
- John Wieting, Jonathan Mallinson, and Kevin Gimpel. 2017. [Learning paraphrastic sentence embeddings from back-translated bitext](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 274–285, Copenhagen, Denmark. Association for Computational Linguistics.

Jingjing Xu, Xuancheng Ren, Junyang Lin, and Xu Sun. 2018. Diversity-promoting GAN: A cross-entropy based generative adversarial network for diversified text generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3940–3949, Brussels, Belgium. Association for Computational Linguistics.

Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018. Generating informative and diverse conversational responses via adversarial information maximization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1810–1820. Curran Associates, Inc.

A Experiment Setup

A.1 Dialog Models

We train Transformer conditional language models in FAIRSEQ using parameters from the FLORES⁹ benchmark for low-resource machine translation (Guzmán et al., 2019) for both the baseline and SMRT. We use the publicly released SMRT fork of FAIRSEQ (Ott et al., 2019; Khayrallah et al., 2020),¹⁰ along with the PRISM M39V1 paraphraser (Thompson and Post, 2020a).¹¹

We use a 5-layer encoder and decoder, 512 dimensional embeddings, and 2 encoder and decoder attention heads. We regularize with 0.2 label smoothing, and 0.4 dropout. We optimize using Adam with a learning rate of 10^{-3} . We train 100 epochs, and select the best checkpoint based on validation set perplexity. We generate with a beam size of 10, and no length penalty.

Figure 1 shows the train command for SMRT, Figure 2 shows the train command for the NLL baseline.

We train and evaluate on the DailyDialog corpus (Li et al., 2017), as released by ParlAI (Miller et al., 2017).¹² We pretrain on the OpenSubtitles corpus (OS; Lison and Tiedemann, 2016).¹³

Since SMRT compares the distribution over tokens from the paraphraser and chatbot their vocabularies must match, so we apply the PRISM SentencePiece model (Kudo and Richardson, 2018) to the DailyDialog and OpenSubtitles corpora. The ParlAI release of DailyDialog is tokenized and lowercased. Since the data the paraphraser is trained on is not, we detokenize and recase the DailyDialog data. We then provide the PRISM dictionary when running FAIRSEQ-PREPROCESS (see Figure 3).

For MMI we use SMRT for the reverse model as well. For pretraining + SMRT we use standard NLL for pretraining on OpenSubtitles, and fine-tune on DailyDialog with SMRT.

A.2 Evaluation Protocols

A.2.1 Human Evaluation

We randomly sample 500 prompt-response pairs from the test set, and filter out any that are not distinct, leaving 482 pairs.

A.2.2 Automatic Quality Evaluation

In Appendix B we report the full automatic evaluation results of the 14 metrics across both the single reference and multi-reference evaluation from the the multi-reference automatic evaluation framework for DailyDialog released by Gupta et al. (2019), which is computed using NLG-EVAL¹⁴ (Sharma et al., 2017). This include word-overlap metrics: BLEU (Papineni et al., 2002), METEOR (Lavie and Agarwal, 2007), and ROUGE-L (Lin, 2004) as well as embedding based metrics: SkipThought (Kiros et al., 2015), embedding average (Forgues et al., 2014), vector extrema and Greedy Matching (Rus and Lintean, 2012). For reading ease, we reports metrics scaled between 0 and 100 rather than 0 and 1.

A.2.3 Automatic Diversity Evaluation

We compute the type/token ratio on tokenized text, using the same spaCy¹⁵ tokenization used in the quality evaluation scripts.¹⁶

⁹<https://github.com/facebookresearch/flores/tree/5696dd4ef07e29977d5690d2539513a4ef2fe7f0>

¹⁰<https://github.com/thompsonb/fairseq-smrt/tree/fdaad1faa01a630beba0c969bd26b65941787752>

¹¹<https://github.com/thompsonb/prism/tree/d2c94b1160f76b3a817eba7f9aba3436deb44731>

¹²<https://github.com/facebookresearch/ParlAI/tree/1e905fec8ef4876a07305f19c3bbae633e8b33af>

¹³<https://www.opensubtitles.org>

¹⁴<https://github.com/Maluuba/nlg-eval/tree/846166566bf0fdccbaa9e5b41da97147470b525b>

¹⁵<https://spacy.io/>

¹⁶<https://github.com/Maluuba/nlg-eval/tree/846166566bf0fdccbaa9e5b41da97147470b525b>

```
python fairseq-smrt/train.py \  
$DATADIR \  
--source-lang src \  
--target-lang tgt \  
--seed 10 \  
--save-dir $SAVEDIR --paraphraser-lang-prefix "<en>" \  
--patience 50 --criterion smrt_cross_entropy \  
--paraphraser-model prism/m39v1/checkpoint.pt \  
--paraphraser-data-dir prism/m39v1/ \  
--paraphraser-sample-topN 100 \  
--prob-use-smrt 1.0 \  
--label-smoothing 0.2 \  
--share-all-embeddings \  
--arch transformer --encoder-layers 5 --decoder-layers 5 \  
--encoder-embed-dim 512 --decoder-embed-dim 512 \  
--encoder-ffn-embed-dim 2048 --decoder-ffn-embed-dim 2048 \  
--encoder-attention-heads 2 --decoder-attention-heads 2 \  
--encoder-normalize-before --decoder-normalize-before \  
--dropout 0.4 --attention-dropout 0.2 --relu-dropout 0.2 \  
--weight-decay 0.0001 \  
--optimizer adam --adam-betas '(0.9, 0.98)' --clip-norm 0 \  
--lr-scheduler inverse_sqrt --warmup-updates 4000 --warmup-init-lr 1e-7 \  
--lr 1e-3 --min-lr 1e-9 --no-epoch-checkpoints \  
--max-tokens 4000 \  
--max-epoch 100 --save-interval 10 --update-freq 4 \  
--log-format json --log-interval 100
```

Figure 1: SMRT training command.

```

python fairseq-smrt/train.py \
  $DATADIR \
  --source-lang src \
  --target-lang tgt \
  --seed 10 \
  --save-dir $SAVEDIR \
  --patience 50 --criterion label_smoothed_cross_entropy \
  --label-smoothing 0.2 \
  --share-all-embeddings \
  --arch transformer --encoder-layers 5 --decoder-layers 5 \
  --encoder-embed-dim 512 --decoder-embed-dim 512 \
  --encoder-ffn-embed-dim 2048 --decoder-ffn-embed-dim 2048 \
  --encoder-attention-heads 2 --decoder-attention-heads 2 \
  --encoder-normalize-before --decoder-normalize-before \
  --dropout 0.4 --attention-dropout 0.2 --relu-dropout 0.2 \
  --weight-decay 0.0001 \
  --optimizer adam --adam-betas '(0.9, 0.98)' --clip-norm 0 \
  --lr-scheduler inverse_sqrt --warmup-updates 4000 --warmup-init-lr 1e-7 \
  --lr 1e-3 --min-lr 1e-9 --no-epoch-checkpoints \
  --max-tokens 4000 \
  --max-epoch 100 --save-interval 10 --update-freq 4 \
  --log-format json --log-interval 100

```

Figure 2: Baseline NLL training command.

```

python fairseq-smrt/preprocess.py \
  --source-lang src --target-lang tgt \
  --trainpref $path_to_sentencepieced_data/train.sp \
  --validpref $path_to_sentencepieced_data/valid.sp \
  --testpref $path_to_sentencepieced_data/test.sp \
  --srcdict prism/m39v1/dict.tgt.txt \
  --tgtdict prism/m39v1/dict.tgt.txt \
  --destdir $databin

```

Figure 3: fairseq-preprocess command.

B Extended Automatic Results

Table 12 and Table 13 show the evaluation against the multiple references for the word based and embedding based metrics. Table 14 and Table 15 show the evaluation against the original single reference for the word based and embedding based metrics.

	Average Max Sentence BLEU				Corpus BLEU				METEOR	ROUGE
	BLEU1	BLEU2	BLEU3	BLEU4	BLEU1	BLEU2	BLEU3	BLEU4		
baseline	27.9	14.3	9.8	7.3	48.3	25.1	15.3	10.0	12.8	34.0
SMRT	29.2	16.4	11.6	8.9	49.9	28.1	18.1	12.4	13.8	36.1
baseline + MMI	27.8	13.8	9.3	7.0	48.2	24.3	14.6	9.5	12.7	33.5
SMRT + MMI	29.2	16.2	11.5	8.7	50.1	27.9	17.9	12.2	13.7	35.8
PT + baseline	29.5	15.9	11.0	8.3	49.9	27.1	16.9	11.3	13.6	35.8
PT + SMRT	29.7	16.6	11.8	9.0	50.7	28.4	18.1	12.3	13.9	36.6

Table 12: Word-overlap based metrics on multiple references.

	Cosine Similarity			GreedyMatching
	SkipThought	Embed. Avg.	VectorExtrema	
baseline	71.7	90.6	62.2	76.9
SMRT	73.6	90.5	63.4	77.7
baseline + MMI	71.6	90.7	62.3	76.7
SMRT + MMI	73.5	90.5	63.3	77.6
PT + baseline	72.5	90.9	63.2	77.5
PT + SMRT	73.8	90.5	63.5	77.6

Table 13: Embedding based metrics on multiple references

	Average Max Sentence BLEU				Corpus BLEU				METEOR	ROUGE
	BLEU1	BLEU2	BLEU3	BLEU4	BLEU1	BLEU2	BLEU3	BLEU4		
baseline	15.8	7.5	5.4	4.2	14.0	6.6	4.1	2.8	6.9	20.9
SMRT	18.0	10.0	7.4	5.8	15.1	8.2	5.5	3.9	8.1	24.0
baseline + MMI	15.4	7.0	5.0	3.9	13.7	6.2	3.8	2.6	6.6	20.1
SMRT + MMI	17.9	9.7	7.2	5.7	15.1	8.0	5.3	3.8	7.9	23.5
PT + baseline	16.4	8.0	5.7	4.5	14.6	7.0	4.4	3.0	7.1	21.7
PT + SMRT	17.9	9.8	7.3	5.7	15.2	8.1	5.3	3.8	7.9	23.7

Table 14: Word-overlap based metrics on the single reference test set

	Cosine Similarity			GreedyMatching
	SkipThought	Embed. Avg.	VectorExtrema	
baseline	64.8	86.1	50.0	71.2
SMRT	67.2	86.5	52.1	72.5
baseline + MMI	64.6	86.0	49.9	70.8
SMRT + MMI	67.0	86.4	51.9	72.3
PT + baseline	65.4	86.4	50.6	71.5
PT + SMRT	67.1	86.5	52.0	72.3

Table 15: Embedding based metrics on the single reference test set

C Examples

C.1 Paraphrase Examples

Table 16 and Table 17 each show a training pair and 20 independent random paraphrases of the response. Sampling is limited to the top 100 tokens per time-step. During training a new sample is taken in each of the 100 epochs. While there are a few small errors, overall the paraphrases are of high quality and remain valid responses. Since sampling is redone each epoch, an error would only be seen once in training.

C.2 Dialog Examples

Tables 18 through 24 show example evaluation prompts and example outputs.

prompt:	It's a wonderful Spanish style.
response:	Oh, I love the roof tiles on Spanish style houses.

paraphrases:

- Ahmed says he loves the cladding of Spanish-style houses.
- Oh, I like the roof tiles on houses in Spanish style.
- Oh, I love the roof tiles of Spanish style houses.
- Oh, I love the roof tiles on Spanish-style houses.
- Oh, I love roof tiles on Spanish-architecture homes.
- Oh, I enjoy roof tiles, in Spanish-style homes.
- Oh, I love the roof tiles on Spanish-American style houses.
- Oh, I love the roof tills on Spanish style houses.
- Oh, I love the roof tiles in Spanish-style house.
- Oh, I love the roof tiles to house in Spanish style
- OH, I love the tiled roofs that are found in Spanish style homes.
- Oh, I love the roof tiles on Spanish-style houses.
- Oh, I like roof tiles on Spanish-style houses.
- Oh, I love the roof tiles on Spanish-style houses.
- Oh I'm very fond of roof tiles in Spanish-style homes.
- Oh, I adore the roof tiles on Spanish-style homes.
- Oh, I love the roof tiles on Spanish style houses.
- Oh, I love roof tiles in Spanish-style houses.
- Oh, I love roof tilting on Spanish-style houses.
- Oh, I love the roof tiles on Spanish style houses.

Table 16: An example training pair with 20 independent random paraphrase samples of the response. Sampling is limited to the top 100 tokens per time-step. During training a new sample is taken each epoch.

prompt:	Oh, man. I hope you get to feeling better.
response:	I would feel better if you came over and took care of me.

paraphrases:

- I'd be well if you came up and took care of me.
- I would feel better if you came down and took care of me.
- I would be better off if you came up and took care of me.
- And I wish you could make care of me and find me.
- I would do better if you came and took care of me.
- I'd feel better if You came up and took care of me.
- I would feel better if you came up and took care of me.
- I would feel better if you come over and take care of me.
- I will feel best if you come over and take care of me.
- I would feel better if you came and took care of me.
- This is something I would feel better if you came up and took care of me.
- I would feel better if you came and cared for me.
- I would feel better if you came up and took care of me.
- I would feel better if you came up and took care of me.
- I would feel better if you came up and took care of me.
- I would feel better if you came out and took care of myself.
- I'd be better off if you came over and took care of me.
- I hope you'd grow up and take care of me.
- I would feel better if you got over and looked after me.
- I would feel better if you came down and picked me up.

Table 17: An example training pair with 20 independent random paraphrase samples of the response. Sampling is limited to the top 100 tokens per time-step. During training a new sample is taken each epoch.

prompt:	that sounds nice. Do you want to share some kebabs as well?
baseline:	sure. How about some cookies?
SMRT:	sure, what about some lambs?
PT + baseline:	yes, I'd like to share some kebabs.
PT + SMRT:	yes, I would like to share some kebabs.

Table 18: An evaluation prompt and example output. The baseline mentions the wrong kind of food. The remaining systems give good responses, though perhaps SMRT should be suggesting 'lamb' (kebab) rather than 'lambs'.

prompt:	I'll use ginger, garlic scallion, hot pepper and vinegar.
baseline:	I'm afraid I'm not sure.
SMRT:	I'm sorry, I'm going to buy it!
PT + baseline:	Anything else?
PT + SMRT:	That sounds like a good idea.

Table 19: An evaluation prompt and example output. The baseline is off-topic, but SMRT is better here. Baseline and SMRT pretraining each give good responses.

prompt:	Believe it or not, tea is the most popular beverage in the world after water.
baseline:	Really? What's the nature of this city?
SMRT:	Really? I am so surprised.
PT + baseline:	That's true. Tea is the most popular in the world.
PT + SMRT:	Yes, I agree with you, but I don't know what to do.

Table 20: An evaluation prompt and example output. The baseline is off-topic, but SMRT and baseline pretraining give good (but very different) responses. Pretrained SMRT is less ideal, though could be a fine response in some contexts.

prompt:	Oh, good. How old is your niece?
baseline:	I'm afraid I'm not sure.
SMRT:	Eighteen.
PT + baseline:	She's 42.
PT + SMRT:	She is ten years old.

Table 21: An evaluation prompt and example output. The baseline gives an 'I don't know' variant, but the remaining systems have good output.

prompt:	Can you tell me how to get to the Jade Buddha Temple?
baseline:	Sure. You can buy yourself a ID card.
SMRT:	You can buy a ticket from the buyer's cheque.
PT + baseline:	Sure. Go straight ahead and turn left.
PT + SMRT:	Go straight ahead and turn left, turn left, turn left and turn left.

Table 22: An evaluation prompt and example output. The baseline does not answer the question, but the remaining systems have good output.

prompt:	I know what you mean. I ought to wear right clothes at the right time.
baseline:	I don't know. I'm not sure what to do.
SMRT:	I guess you are right.
PT + baseline:	I'm sorry to hear that. I don't know what to wear.
PT + SMRT:	I don't know what to wear at the right time.

Table 23: An evaluation prompt and example output. The baseline gives an 'I don't know' variant. While the pretrained baseline has some uncertainty, it is at least unsure about clothing. The SMRT systems have good output.

prompt:	Do you speak only English in the class, or does your teacher explain everything to you in Spanish?
baseline:	No, I don't know how to speak Spanish.
SMRT:	Oh, we never speak Spanish.
PT + baseline:	I don't know.
PT + SMRT:	No, I don't speak English in the class.

Table 24: An evaluation prompt and example output. The baseline pretraining gives an 'I don't know' variant. The baseline is okay, the SMRT systems have better output.