

Measuring the ‘I don’t know’ Problem through the Lens of Gricean Quantity

Huda Khayrallah
Johns Hopkins University
huda@jhu.edu

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

Abstract

We consider the intrinsic evaluation of neural generative dialog models through the lens of Grice’s Maxims of Conversation (1975). Based on the maxim of Quantity (be informative), we propose Relative Utterance Quantity (RUQ) to diagnose the ‘I don’t know’ problem, in which a dialog system produces generic responses. The linguistically motivated RUQ diagnostic compares the model score of a generic response to that of the reference response. We find that for reasonable baseline models, ‘I don’t know’ is preferred over the reference the majority of the time, but this can be reduced to less than 5% with hyperparameter tuning. RUQ allows for the direct analysis of the ‘I don’t know’ problem, which has been *addressed* but not *analyzed* by prior work.

1 Introduction

Neural generative dialog models have a tendency to produce generic, safe responses, such as ‘I don’t know’ (Serban et al., 2016; Li et al., 2016a). The repetition of such phrases is annoying to users, and contributes nothing to the conversation.

Evaluating chatbots is an active area of research, partly due to their open-ended nature (Hashimoto et al., 2019; Sedoc et al., 2019; Li et al., 2019;

Mehri and Eskenazi, 2020b; Deriu et al., 2020). To the best of our knowledge, no prior work focuses on *analyzing* systems for generic, safe responses, such as ‘I don’t know.’ While prior work (Li et al., 2016a,b; Csáky et al., 2019; Welleck et al., 2020) *addresses* the ‘I don’t know’ problem, the lack of *analysis* leaves it unclear if a method improves models by mitigating *this* problem, or another.

One linguistic framework for analyzing conversations is Grice’s Cooperative Principle (1975), which consists of Maxims of Conversation that function as guidelines for effective communication. Grice considered conversations between humans, but there has also been some exploration in NLP (Bernsen et al., 1996; Harabagiu et al., 1996; Qwaider et al., 2017; Jwalapuram, 2017).

We discuss each of the categories of maxims and the ways a chatbot might violate them in Table 1.

We propose a novel automatic diagnostic inspired by the Gricean QUANTITY maxim. Relative Utterance Quantity checks if the model favors a generic response (such as ‘I don’t know.’) over the reference it was trained on for each prompt. We apply our diagnostic to a method designed to address this problem (Csáky et al., 2019), and find that method does mitigate it, though not by as much as a hyperparameter search.

Maxim	Definition	Violated by...	Prompt: What color is grass?
QUANTITY	Be informative.	not answering a question (fully), or giving too much information.	I don’t know.
QUALITY	Be truthful.	lying, or saying something without evidence.	Grass is purple.
RELATION	Be relevant.	off-topic responses.	I like pizza.
MANNER	Be clear, brief, and orderly.	disfluent responses	is green grass usually.

Table 1: Gricean maxims, with examples of how they can be violated for the prompt ‘What color is grass?’

2 Relative Utterance Quantity (RUQ)

If a system responds ‘I don’t know.’ when it could have given a better or more informative answer, this is by definition a violation of QUANTITY. Based on this interpretation we propose a method for diagnosing the problem. We compare the model score of producing ‘I don’t know.’ to the model score of producing the reference response. This can be done on the training data, or the test data. Particularly on the training data, we should expect the model to ‘know’ the data it was trained on and therefore score it higher than ‘I don’t know.’

We propose two diagnostic measures to compute the Relative Utterance Quantity of a model: (1) We plot the average model score for each token across sentences. We compare the original reference, beam search output, and two ‘I don’t know’ (IDK) variants: ‘I don’t know.’ and ‘I don’t know what to do.’ allowing for the visualization of the relative gap in scores at different points in the sentence. (2) We compute the (length normalized) model score for ‘I don’t know.’ and the reference of each training prompt, and count how many times the reference is preferred. We denote the later as RUQ score. Both generalize to other generic responses, as might be appropriate for other corpora or other languages.

If there are multiple references we would recommend comparing the lowest likelihood reference for RUQ score, since all valid references should be better than I don’t know.

We note that RUQ captures some types of QUANTITY violations, but not all violations of this maxim.

3 Data

Following [Khayrallah and Sedoc \(2020\)](#), we train and evaluate on DailyDialog ([Li et al., 2017](#)),¹ which consists of $\sim 80,000$ turns of English-learners practicing ‘daily dialogues’ in various contexts, e.g., chatting about vacation or food.

We also use Entropy-Based Data Filtering ([Csáky et al., 2019](#)), which filters out high entropy utterances² with the goal of removing generic ones. We use the recommended filtering threshold of

¹As released by ParlAI ([Miller et al., 2017](#)). The ParlAI release of DailyDialog is tokenized and lowercased. Following [Khayrallah and Sedoc \(2020\)](#) we detokenize and recase the DailyDialog data for training.

²Prompts that solicit many different responses and responses that can apply to many different prompts.

1.0 and ‘IDENTITY’ clustering. We filter based on their ‘source’, ‘target’, and ‘both’ settings. We consider ‘target’ as the baseline, as they find it works best. We denote models trained on DailyDialog as DD and models trained on [Csáky et al.](#)’s entropy filtered version as EF.

4 Evaluation Metrics

4.1 Standard Automatic Metrics

We use the single-reference and multi-reference³ automatic evaluation framework for DailyDialog released by [Gupta et al. \(2019\)](#),⁴ which is computed using NLG-EVAL ([Sharma et al., 2017](#)).⁵ We primarily consider multi-reference METEOR ([Lavie and Agarwal, 2007](#)); see Appendix A.7 for all metrics.⁶

4.2 Human Evaluation

For human evaluation of the different systems we use crowdworkers on Amazon Mechanical Turk to judge the fluency, coherence, and interestingness of utterances on a 1-5 Likert scale (see Appendix A.4 for full details) for 100 randomly sampled evaluation set prompts. Four annotators judge the responses from all systems for each prompt in a single turn context. We remove any annotators with a linear Cohen’s Kappa < 0.1 from the results.

5 Models

Following [Khayrallah and Sedoc \(2020\)](#), we train Transformer ([Vaswani et al., 2017](#)) chatbots in FAIRSEQ using parameters from the FLORES benchmark for low-resource MT ([Guzmán et al., 2019](#)):⁷ 5-layer encoder and decoder, 512 dimensional embeddings, and 2 encoder and decoder attention heads. The default regularization parameters are 0.2 label smoothing ([Szegedy et al., 2016](#)), 0.4 dropout, and 0.2 attention & ReLU dropout.

5.1 Hyperparameter Sweep

Some kinds of regularization (e.g., label smoothing and subword vocabularies) are not universally used

³For RUQ, we only use the original single-reference.

⁴github.com/prakharguptaz/multirefeval

⁵github.com/Maluuba/nlg-eval

⁶For reading ease, we report metrics scaled between 0 and 100 rather than 0 and 1.

⁷See § A.6 for full details for replication.

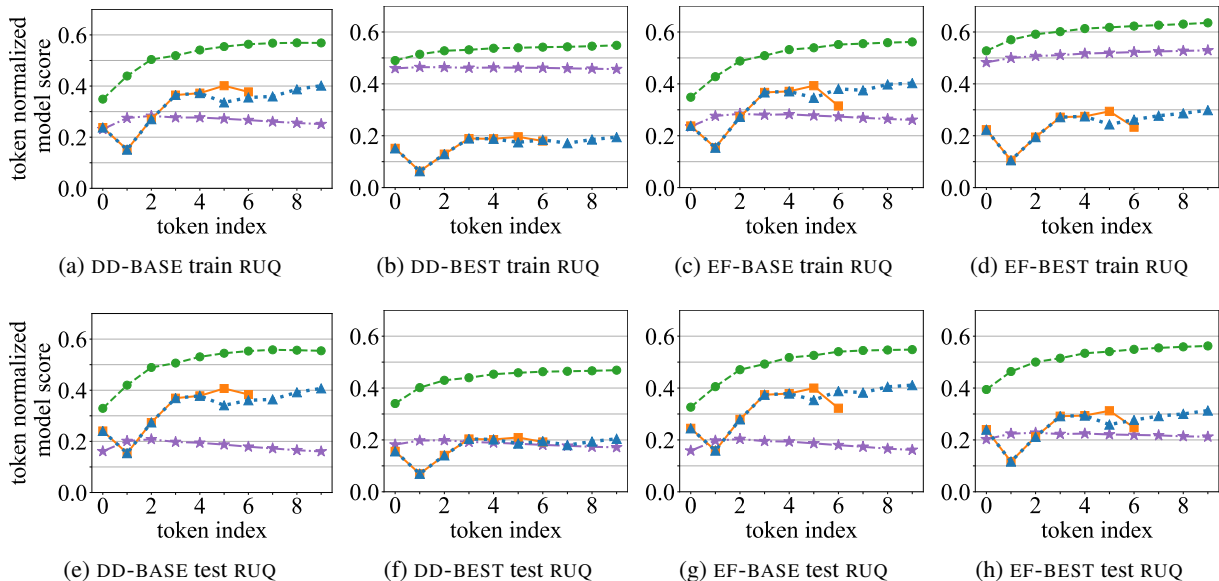


Figure 1: RUQ plots on the train (top) and test (bottom) data. We plot the token normalized model score for the reference (★), the beam-search output (●), ‘I don’t know.’ (■), and ‘I don’t know what to do.’ (▲). Points are per (subword) token, and averaged over all prompts.

training data	BASE	BEST
DAILYDIALOG	12.7	17.8
ENTROPY-FILTERED	13.2	17.2

Table 2: Multi-reference METEOR for the four systems we analyze in this work. BEST models are the result of the hyper parameter sweeps.

in dialog.⁸ Since we are concerned with the model over-fitting on IDK, we perform a hyperparameter sweep of regularization parameters, including SentencePiece (Kudo and Richardson, 2018) vocabulary size, learning rate, dropout, attention & relu dropout, and label smoothing.⁹

We denote models trained with the FLORES hyperparameters as BASE, and the best model from the hyperparameter searches for each data type (as selected by multiple-reference METEOR) as BEST.

We report the multi-reference METEOR scores for the BASE and BEST systems in Table 2.¹⁰ For the DailyDialog data we find that hyperparameter tuning can improve multiple-reference METEOR from 12.7 (DD-BASE) to 17.8 (DD-BEST).

We perform the same hyperparameter sweep after performing entropy filtering (Csáky et al.,

⁸For example popular toolkits for dialog (e.g., Hugging Face (Wolf et al., 2020) and ParlAI (Miller et al., 2017)) do not implement label smoothing.

⁹See Appendix A.1 for more hyperparameter details.

¹⁰We report hyperparameters of these models and their performance on the full set of automatic metrics in § A.7.

2019) on the data, but we find that the best model is still DD-BEST. Without hyperparameter tuning, entropy filtering improves performance by ~ 0.5 on multi-reference METEOR, but the improvement by hyperparameter sweeping is much larger (5.1 points).¹¹

We did a very thorough sweep (including values we expected to perform poorly), which led to some general takeaways:¹² Using a subword vocabulary (of 4-8k) is helpful. (2) Label smoothing interacts with subword vocabulary size, but is also helpful.

6 Relative Utterance Quantity

6.1 RUQ Plots

We show plots for the four models in Figure 1. We plot the token normalized model score for reference and ‘I don’t know.’ For additional comparison, we also plot the model scores for the

¹¹We note that Csáky et al. (2019)—who proposed entropy filtering and an observed a 1 BLEU point improvement from using it (we observed a 0.3 improvement in single reference BLEU)—did not use any subwords units; they used a total vocab size of 16k. Our 10 best systems all had Sentencepiece vocab sizes of 2k, 4k, or 8k, so perhaps this difference may explain the discrepancy between their results and our replication. We note that for the 3 metrics which we believe our evaluations are comparable—single reference Embedding Average Cosine Similarity, and single reference Vector Extrema Cosine Similarity—our baseline outperforms their results. The BLEU scores are not directly comparable because they report sentence BLEU, while we report corpus BLEU following Gupta et al. (2019).

¹²See chateval.org/RUQ for automatic metrics on the full hyper parameter sweep.

training data	BASE	BEST
DAILYDIALOG	28.5%	95.3%
ENTROPY-FILTERED	37.9%	89.2%

Table 3: Training data RUQ scores. Entropy filtering improves how often the reference is preferred to ‘I don’t know.’, but by less than the hyperparameter sweeps (which are denoted BEST).

	Fluency	Coherence	Interestingness
Human	4.9	4.6	4.0
DD-BASE	4.8	3.5	2.6
DD-BEST	4.8	3.8	2.7
EF-BASE	4.4	3.3	2.8
EF-BEST	4.4	3.1	3.3

Table 4: Average human judgement ratings on 1-5 pointwise scale for DailyDialog (DD) and the entropy filtered (EF) data. The result of the hyperparameter sweep is denoted BEST.

beam-search output and ‘I don’t know what to do.’ Overall, we observe that for the BASE models the IDKs are higher probability than the reference, even on the training data. This is problematic, because the model is ranking a response that is not providing enough QUANTITY of information higher than the reference despite the fact that it should ‘*know*’ the training data. The relative difference in probabilities is much better in DD-BEST than DD-BASE, particularly on the training set. Simply entropy filtering the data alone does not fix the problem.

6.2 RUQ scores

We summarize QUANTITY in a single statistic by counting how many times the reference has a higher probability than ‘I don’t know.’ on the training data.

Entropy filtering improves how often the reference is preferred to ‘I don’t know.’, but not by as much as the hyperparameter sweep does, see Table 3 for the RUQ scores on the training data.¹³ For both DD-BASE and EF-BASE, IDK is preferred over the reference response the model was trained on over half of the time (71.5% for DD, 62.1% for EF).

¹³RUQ scores on the on the test data are reported in § A.7. The overall trend is same, but the absolute values lower.

6.3 Human Evaluation

Table 4 shows human judgments of fluency, coherence, and interestingness.¹⁴ The models trained on DailyDialog have higher fluency and coherence, while the models trained on the filtered data have higher interestingness. For both kinds of data, the hyperparameter tuning (as selected by METEOR) improved interestingness. Fluency did not change. Coherence was reduced for the filtered models and improved for the base model. Improved RUQ may be reflected in either interestingness or coherence, but other factors can influence those judgments. Therefore, measuring RUQ directly is important to measuring progress on the IDK problem.

7 Discussion

The relative RUQ rankings of the four systems we consider in this work are the same as the relative rankings by multi-reference METEOR, and DD-BEST (the single best model according to multi-reference METEOR) is also the one with the highest RUQ score. Among all models in the hyperparameter sweep, RUQ is correlated with METEOR with Spearman’s ρ of 0.9 but this drops to 0.6 when considering only the top 20 systems, demonstrating that RUQ and METEOR do not capture the same phenomenon. We note that RUQ on the training data does not require a particular (multi-reference) test set like most automatic evaluation metrics. RUQ simply diagnoses how well the model learned the training data compared to a generic response.

The model’s relative preference of IDK over the (presumably) better reference response is not only a QUANTITY violation, but is also indicative of a fundamental problem with the models themselves, and should be fixed before decoding time (either by correcting the data, or by correcting the model).

Csáky et al. (2019) argue that the IDK problem is due to the one-to-many/many-to-one nature of dialog training data—if a single response applies to many different responses, it will become the canonical response. Therefore their entropy filtering method removes one-to-many/many-to-one pairs, by removing high entropy responses. While this data filtering reduces the problem, we found that the baseline model trained on the

¹⁴§ A.5 discusses head to head judgments. Models trained on the DailyDialog data are preferred over the filtered models, but there is no clear preference between base and best models.

entropy filtered data (EF-BASE) still preferred IDK over the reference the majority of the time, suggesting opportunities for future research on the IDK problem.

8 Related Work

Gricean Maxims in NLP Gricean maxims have previously been discussed in NLP. [Bernsen et al. \(1996\)](#) examine the relationship between a new set of maxims for human-bot dialogs and relate them to Gricean maxims. They point out that these do not entirely overlap; however, the maxim of Quantity is preserved since unambiguous contributing responses are required in conversations in general. ([Harabagiu et al., 1996](#)) attempt to explicitly create an evaluation methodology using sets of primitive rules and WordNet. Our approach is different as RUQ is a diagnostic metric.

[Jwalapuram \(2017\)](#) propose a Gricean dialog evaluation where humans rate performance on a Likert scale for each category. [Qwaider et al. \(2017\)](#) consider the QUANTITY, RELATION, and MANNER maxims for ranking community question answers. They use other NLP tools to evaluate if the response has key elements or named entities (QUANTITY/RELATION), has high semantic similarity (RELATION), and includes/excludes positive/negative polarity terms (MANNER).

Chatbot evaluation Automatic evaluations for dialog typically measure lexical or semantic similarity between a produced response and a reference, under the assumption that the reference is a good response and responses similar to it will be good as well. Since there are often multiple valid responses to a prompt, this can be extended to multiple references too. In contrast, in our work we compare a model’s score of a reference to a model’s score of a generic response for directed analysis.

HUSE ([Hashimoto et al., 2019](#)) uses the model score combined with human judgments to evaluate diversity and quality, classifying a response as human- or machine-generated. Our work does not require human judgments, and compares the model score of a generic response to the reference response.

[Mehri and Eskenazi \(2020a\)](#) also use scoring from a model. Whereas that work is using an external model, we propose an intrinsic diagnostic for a particular phenomenon. Each serves a

different purpose, and an advantage of our method is our analysis does not require an external model, which might not be available in all languages and for all types of text.

Mitigating the IDK Problem A variety of approaches have been proposed to *mitigate* the IDK problem. These include active post-processing methods such as MMI ([Li et al., 2016a](#)), as well as training data filtration ([Csáky et al., 2019](#)), reinforcement learning ([Li et al., 2016b](#)) and unlikelihood training ([Welleck et al., 2020](#)). In our work, we propose an intrinsic model diagnostic to *analyze* the problem.

MMI Maximum Mutual Information was proposed as a ‘Diversity-Promoting Objective Function’ for dialog ([Li et al., 2016a](#)). MMI-bidi encourages the prompt to be predictable from the response, by using a reverse direction model. We argue this was not diversity broadly speaking, but actually tackling a RELEVANCY problem, since it is scoring how predictable the prompt is from the response.

[Li et al.](#) demonstrate MMI improves performance, though recent work found that it does not always ([Khayrallah and Sedoc, 2020](#)).

Copying in Machine Translation [Ott et al. \(2018\)](#) found that copying was overrepresented in the output of RNN NMT. Using an analysis that inspired RUQ plots they compare the score of the beamsearch output to that of the copied source. They also consider the probability at each position in the output, and find the model is unlikely to start copying; however, after starting to copy continuing to copy has high probability. We find IDK has a relatively high score from the start, though for some models the gap widens towards the end of the sentence.

9 Conclusion

We reframe the IDK problem as a violation of the Gricean maxim of QUANTITY, and introduce a new measure—Relative Utterance Quantity (RUQ)—which allows researchers to diagnose if their model is violating this particular conversational principle, and analyze methods that aim to address it.

We aim to encourage further discussion and research drawing on linguistic principles about discourse and pragmatics for analysis of dialog models.

Acknowledgments

We thank Patrick Xia, Nathaniel Weir, Rachel Rudinger, and Claire Daniele for their helpful comments and feedback on the paper. We additionally thank the reviewers for their insightful comments.

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 of DARPA or the U.S. Government.

References

- Niels Ole Bernsen, Hans Dybkjær, and Laila Dybkjær. 1996. [Cooperativity in human-machine and human-human spoken dialogue](#). *Discourse Processes*, 21(2):213–236.
- Richárd Csáky, Patrik Purgai, and Gábor Recski. 2019. [Improving neural conversational models with entropy-based data filtering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5650–5669, Florence, Italy. Association for Computational Linguistics.
- Jan Deriu, Don Tuggener, Pius von Däniken, Jon Ander Campos, Alvaro Rodrigo, Thiziri Belkacem, Aitor Soroa, Eneko Agirre, and Mark Cieliebak. 2020. [Spot the bot: A robust and efficient framework for the evaluation of conversational dialogue systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3971–3984, Online. 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*.
- H. P. Grice. 1975. *Logic and Conversation*, pages 41 – 58. Brill, Leiden, The Netherlands.
- 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.
- Michael Alexander Kirkwood Halliday. 1989. *Spoken and Written Language*. Language education. Oxford University Press.
- Sanda Harabagiu, Dan Moldovan, and Takashi Yukawa. 1996. [Testing gricean constraints on a wordnet-based coherence evaluation system](#). In *Working Notes of the AAAI-96 Spring Symposium on Computational Approaches to Interpreting and Generating Conversational Implicature*, pages 31–38.
- Tatsunori Hashimoto, Hugh Zhang, and Percy Liang. 2019. [Unifying human and statistical evaluation for natural language generation](#). 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 1689–1701, Minneapolis, Minnesota. Association for Computational Linguistics.
- Prathyusha Jwalapuram. 2017. [Evaluating dialogs based on Grice’s maxims](#). In *Proceedings of the Student Research Workshop Associated with RANLP 2017*, pages 17–24, Varna. INCOMA Ltd.
- Huda Khayrallah and João Sedoc. 2020. [SMRT chatbots: Improving non-task-oriented dialog with Simulated Multiple Reference Training](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4489–4505, Online. 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.
- Batia Laufer and Paul Nation. 1995. [Vocabulary Size and Use: Lexical Richness in L2 Written Production](#). *Applied Linguistics*, 16(3):307–322.
- 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. 2016a. [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.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016b. [Deep reinforcement learning for dialogue generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Austin, Texas. Association for Computational Linguistics.
- Margaret Li, Jason Weston, and Stephen Roller. 2019. [Acute-eval: Improved dialogue evaluation with optimized questions and multi-turn comparisons](#).
- 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.
- 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.
- Shikib Mehri and Maxine Eskenazi. 2020a. [Unsupervised evaluation of interactive dialog with DialoGPT](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 225–235, 1st virtual meeting. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 2020b. [USR: An unsupervised and reference free evaluation metric for dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. [ParlAI: A dialog research software platform](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 79–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018. [Analyzing uncertainty in neural machine translation](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3956–3965, Stockholmsmässan, Stockholm Sweden. PMLR.
- 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.
- Mohammed R. H. Qwaider, Abed Alhakim Freihat, and Fausto Giunchiglia. 2017. [TrentoTeam at SemEval-2017 task 3: An application of Grice maxims in ranking community question answers](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 271–274, Vancouver, Canada. 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.
- João Sedoc, Daphne Ippolito, Arun Kirubakaran, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2019. [ChatEval: A tool for chatbot evaluation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 60–65, Minneapolis, Minnesota. Association for Computational Linguistics.
- Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI’16*, page 3776–3783. AAAI Press.
- 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.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. [Rethinking the inception architecture for computer vision](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826.
- 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.

Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

A Appendix

A.1 Hyperparameter Search

We sweep SentencePiece (Kudo and Richardson, 2018) vocabulary size (1k,4k,8k,16k), learning rate (1e-2, 1e-3, 1e-4), dropout (0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6), attention & ReLU dropout (0.0, 0.1, 0.2, 0.3, 0.4), and label smoothing (0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8).

A.2 Standard Automatic Metrics

In § A.7 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),¹⁵ which is computed using NLG-EVAL¹⁶ (Sharma et al., 2017). This includes 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 report metrics scaled between 0 and 100 rather than 0 and 1. See chateval.org/RUQ for automatic metrics on the full hyper parameter sweep.

A.3 Lexical Diversity

The Gricean maxims focus on ensuring cooperation between speakers, but there is more to a conversation than cooperation—especially in an open ended conversation that might be had with a chatbot. This is where additional desiderata may come in to play, such as interestingness. One (indirect) automatic way of measuring interestingness is lexical diversity (Halliday, 1989; Laufer and Nation, 1995), by computing the n-gram type/token ratio (Li et al., 2016a). We use the same spaCy¹⁷ tokenization used in the automatic evaluation scripts (§ A.2).¹⁸

A.4 Pointwise Human Evaluation

We presented Amazon Mechanical Turk workers the task with 4 prompts and all system responses along with a human reference. The annotators had a maximum time allotted of 20 minutes. Our criteria

¹⁵github.com/prakharguptaz/multirefeval

¹⁶github.com/Maluuba/nlg-eval

¹⁷spacy.io

¹⁸github.com/Maluuba/nlg-eval

for inclusion were over 500 approved HITs, an approval rate over 98%, and location set to US. Each HIT was paid \$0.15 with an overlap of 4 annotators per HIT.¹⁹ A screenshot of the HIT is in Figure 2.

A.5 Head to Head Human Evaluation

In addition to the point-wise evaluation, we also test head-to-head pairwise performance on the evaluation set of 480 unique prompt/response pairs, as shown in § A.5. Models trained on the DailyDialog data outperform the filtered models, but there is no clear preference between base and best models.

A.6 Dialog Models

We train Transformer conditional language models in FAIRSEQ using parameters from the FLORES²⁰ benchmark for low-resource machine translation (Guzmán et al., 2019).

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 3 shows the train command.

We train and evaluate on the DailyDialog corpus (Li et al., 2017), as released by ParlAI (Miller et al., 2017).²¹

A.7 Full Automatic Results

Table 6 shows the hyperparameters for each system. Table 7 and Table 8 show the evaluation against the multiple references for the word based and embedding based metrics. Table 9 and Table 10 show the evaluation against the original single reference for the word based and embedding based metrics. Table 11 shows the lexical diversity, and Table 12 shows the RUQ sores.

¹⁹We aimed to compensate the crowdworkers fairly (\$ 8 per hour) and did this by annotating a set of data ourselves to estimate the timing of the task

²⁰<https://github.com/facebookresearch/flores/tree/5696dd4ef07e29977d5690d2539513a4ef2fe7f0>

²¹<https://github.com/facebookresearch/ParlAI/tree/1e905fec8ef4876a07305f19c3bbae633e8b33af>

Please Note

- You have to be an **English Native Speaker**.
- You have to complete the ratings for all responses. **All fields are required**.

Informed Consent

This is a linguistic experiment performed at [REDACTED]. If you have any question about this study, feel free to contact [REDACTED]. Participation in this research is voluntary. You have the right to withdraw from the experiment at any time. The collected data will be used for research purposes only. Personal data will be kept confidential and will not be shared with third parties.

Instructions

In this task you will read a turn of a conversation. For each conversation you will see possible responses to the last turn generated by a computer program. The programs attempts to generate responses that are relevant, while also making an interesting contribution to the conversation.

For each prompt, you will read several system possible responses, and judge each response on its appropriateness on a 1 to 5 scale from not at all appropriate to extremely appropriate. Appropriateness is defined as follows:

1. **Fluency**: Flows naturally and sounds like what an experienced speaker/writer of the language might say.
2. **Coherence**: Facts and topics are consistent, i.e. non-contradictory with previous parts of the conversation.
3. **Interestingness**: Contains rich, non-repetitious, and interesting information that adds to the conversation.

Example: Below we show a prompt and several possible responses, along with suggested scores for each response.

Prompt: What do you do for work?

Responses Suggested scores

I doctor. Fluency: 1, Coherence: 5, Interestingness: 5. This does not sound like a native speaker's response, but there are no contradictions and it answers the question.
 I like pancakes. Fluency: 5, Coherence: 1, Interestingness: 3. This is grammatical but not coherent as it does not answer the question and only some information was added to the conversation.
 I work. Fluency: 4, Coherence: 4, Interestingness: 1. This is grammatical and mostly coherent, but no information is added to the conversation..
 I work as a teacher. Fluency: 5, Coherence: 5, Interestingness: 5. An excellent on-topic response that does not contradict any previous statements.

Note: ignore things like "i ' m" this should be read as "I'm" and not docked points for grammar. There are ATTENTION CHECKS please make sure you select the number.

{% for prompt in prompts %} {% set promptloop = loop %}

1. **Prompt:** {% for turn in prompt.prompt %}

{{ turn }}

{% endfor %}

{% for model in prompt.models %} {% set modelloop = loop %}

System {{ modelloop.index }}: {{ model.response }}

	Least (1)				Most (5)
Fluency	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coherence	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Interestingness	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

{% endfor %}

Figure 2: Instructions for AMT task.

M1	M2	M1	M2	tie
EF-BEST	DD-BEST	38.5%	42.4%	19.1%
EF-BASE	EF-BEST	34.4%	34.0%	31.7%
EF-BASE	DD-BASE	37.0%	41.6%	21.4%
DD-BASE	DD-BEST	36.3%	36.5%	27.2%

Table 5: Head to head comparison between various systems. Models trained on the DailyDialog data outperform the filtered models, but there is no clear preference between base and best models.

```
python 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 3: Training command.

Data	Params	bpe	lr	dropout	otherdropout	labelsmooth
DD	BASE	4	0.001	0.4	0.2	0.2
DD	BEST	4	0.001	0.0	0.1	0.4
EF	BASE	4	0.001	0.4	0.2	0.2
EF	BEST	2	0.001	0.0	0.1	0.2

Table 6: Hyperparameters for each of the four models we consider.

Data	Params	Average Max Sentence BLEU				Corpus BLEU				METEOR	ROUGE
		BLEU1	BLEU2	BLEU3	BLEU4	BLEU1	BLEU2	BLEU3	BLEU4		
DD	BASE	27.8	14.7	10.3	7.9	48.1	25.6	16.2	11.2	12.7	34.3
DD	BEST	33.9	21.9	17.7	15.3	53.9	36.1	28.9	25.1	17.8	39.7
EF	BASE	27.8	14.0	9.4	7.0	46.9	24.1	14.6	9.8	13.2	33.4
EF	BEST	31.7	19.1	14.9	12.7	51.0	32.8	25.5	21.8	16.9	37.2

Table 7: Word-overlap based metrics on multiple references.

Data	Params	Cosine Similarity			GreedyMatching
		SkipThought	Embed. Avg.	VectorExtrema	
DD	BASE	72.4	90.8	62.9	77.2
DD	BEST	73.8	92.2	65.4	79.3
EF	BASE	71.9	91.2	62.2	77.0
EF	BEST	72.8	91.6	62.7	77.9

Table 8: Embedding based metrics on multiple references.

Data	Params	Average Max Sentence BLEU				Corpus BLEU				METEOR	ROUGE
		BLEU1	BLEU2	BLEU3	BLEU4	BLEU1	BLEU2	BLEU3	BLEU4		
DD	BASE	15.3	7.6	5.6	4.5	12.9	6.3	4.1	3.0	6.7	20.6
DD	BEST	24.3	16.7	14.3	12.8	23.2	16.7	14.2	12.9	11.9	29.2
EF	BASE	15.9	7.4	5.2	4.1	15.8	7.5	4.7	3.3	7.2	20.4
EF	BEST	22.1	14.0	11.8	10.5	22.9	15.8	13.2	11.8	11.1	26.6

Table 9: Word-overlap based metrics on the single reference test set.

Data	Params	Cosine Similarity			GreedyMatching
		SkipThought	Embed. Avg.	VectorExtrema	
DD	BASE	65.3	86.3	50.6	71.3
DD	BEST	68.2	88.5	54.7	74.6
EF	BASE	64.9	86.9	50.2	71.3
EF	BEST	67.0	87.7	52.3	73.1

Table 10: Embedding based metrics on the single reference test set.

Data	Params	1-grams	2-grams	3-grams
DD	BASE	2.4	10.3	18.8
DD	BEST	3.5	18.0	35.5
EF	BASE	2.3	10.7	20.1
EF	BEST	3.8	18.3	34.6

Table 11: Type/Token ratios.

Data	Params	RUQ-train	RUQ-test
DD	BASE	28.5	12.2
DD	BEST	95.3	35.7
EF	BASE	37.9	15.5
EF	BEST	89.2	30.7

Table 12: RUQ scores on the train and test data.