

ENDEX: Evaluation of Dialogue Engagingness at Scale

Guangxuan Xu¹ Ruibo Liu²

Fabrice Harel-Canada¹ Nischal Reddy Chandra¹ Nanyun Peng¹

¹University of California, Los Angeles ²Dartmouth College

{gxu21, violetpeng} @cs.ucla.edu

Abstract

We propose ENDEX, the first human-reaction based model to evaluate dialogue engagingness. ENDEX is trained on 80k Reddit-based Engagement Dataset (RED) curated using a novel distant-supervision framework. Engagingness is a key measure that captures high-level quality of AI dialogue systems and closely reflects actual user experience. However, data shortage, plus the abstract and extensive definition of engagingness makes it challenging to develop an automatic metric. Our work departs from mainstream approaches that use synthetic negative examples to train binary classifiers, and instead, proposes a solution using distant-supervision from human-reaction feedback. To support the soundness of our ENDEX metric, we offer a theoretical foundation for engagement, an extensive ablation study, and empirical evidence of high correlation on five engagingness related datasets.¹

1 Introduction

Many modern generative language models are trained to maximize a likelihood objective, but this paradigm tends to assign high probability to generic responses (Li et al., 2016), such as “I don’t know.”. Prior research has established that people prefer to converse with interesting, creative, and informative agents (See et al., 2019), all concepts broadly related to the notion of *engagingness*. Furthermore, engagingness is recognized as a key evaluation metric for the quality of dialogue systems (Zhang et al., 2018; Ghazarian et al., 2020). For example, FAIR’s ParlAI (Miller et al., 2017) incorporated Engagingness as the default testing metric in the Blenderbot system (Roller et al., 2021); dialogue data challenges, like ConvAI2 (Dinan et al., 2019), Amazon Alexa Prize², and ensemble metrics like FED (Mehri and Eskenazi, 2020), all measure engagingness to benchmark dialogue quality.

However, the current evaluation of engagingness still primarily relies on expensive human annota-

¹Off-the-shelf ENDEX model and the RED dataset is available at <https://github.com/gxxu-ml/EnDex>.

²<https://www.amazon.science/alexaprize>

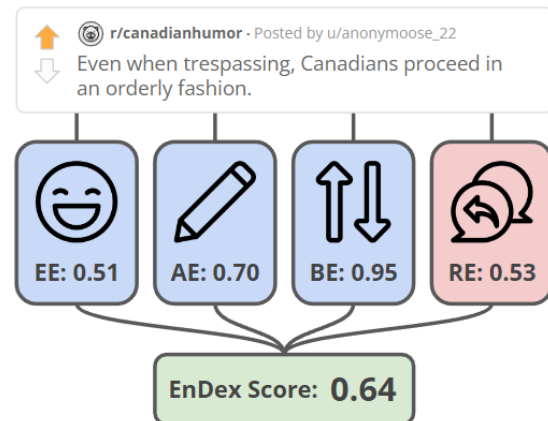


Figure 1: Example of an online post with scores for emotional engagement (EE), attentional engagement (AE), and behavioral engagement (BE) in blue to represent the 3 dimensions of *human engagement*; reply engagement (RE) in red; and the aggregated ENDEX score in green. We apply z-score to EnDex Score and pick a hyper-parameter threshold to cluster posts into positive and negative samples.

tion rather than off-the-shelf automatic tools, due to several theoretical and technical challenges: firstly, unlike more well-characterized properties such as fluency, the definition of engagingness is significantly more abstract and multi-dimensional (See et al., 2019), requiring well-tuned quality metrics for each sub-dimension to aggregate a final score. Secondly, what qualifies as engaging is open-ended and many different answers may embody the concept (Ghazarian et al., 2020). Therefore, reference-based metrics requiring unique ground truth, such as BLEURT (Sellam et al., 2020) and BERTScore (Zhang et al., 2020), cannot apply. Thirdly, there’s an acute shortage of large-scale, high-quality data annotated for engagingness.

Ghazarian et al. (2020) jump-started efforts to automatically measure dialogue engagement, where they fine-tuned a BERT-based model (Devlin et al., 2019) on the ConvAI2 and DiallyDialog datasets (Li et al., 2017) to predict an engagingness score. However, finetuning on small size supervised data could

easily lead to overfitting and generalization problems. Another high performing metric on engagingness USL-H (Phy et al., 2020) assumes a positive set and generates synthetic negative samples to train model. However, credible positive samples are not always available, and synthetic negative samples may not be challenging enough to further advance classifier performance.

In light of the above challenges, we propose ENDEX, a novel metric trained with distantly supervised data to predict turn-level dialogue engagingness (Figure 1). ENDEX requires neither human annotations nor direct disentanglement of engagingness. Instead, we leverage observed *user reactions* to posts as distant signals to model engagingness, which marks a departure from mainstream approach to train on synthetic negative samples (Lan et al., 2020; Ghazarian et al., 2022; Tao et al., 2018; Sato et al., 2020). ENDEX trains on real conversations sourced from Reddit, that are automatically annotated as positive and negative examples with our framework. The novel dataset is named RED (Reddit Engagement Dataset) with over 80k labelled samples. ENDEX framework derives its theoretical underpinning from relevant HCI works, and has shown superior performance on five benchmark datasets.

2 EnDex Metric

Engagingness is not only a linguistic concept useful for dialogue systems, but also manifests itself in multi-modalities and is extensively leveraged to benchmark gaming and online learning experiences (Silpasuwanchai et al., 2016; Chen et al., 2005; McMahan, 2003; Schoenau-Fog, 2011). Our work is inspired by HCI study of *Human Engagement* (Ma, 2018), which decomposes engagingness into three major dimensions including *attentional engagement* (e.g., clicks and scrolls), *behavior engagement* (e.g., facial expressions), and *emotional engagement* (e.g., heart rate).

ENDEX metric follows the same intuition: we can infer engagingness of a text by analyzing human reactions to it, for which there is abundant data in social media. ENDEX metric learns from our distant-supervised RED dataset, which measures dialogue engagement along four dimensions as shown in Figure 1; three-dimensions correspond to the original Human Engagement definition, and one distinct Reply Engagement dimension for the dialogue specific task.

	Engaging	Non-engaging
# of samples	40,162	40,162
Emotional	.605 ± .273	.152 ± .120
Attentional	.759 ± .127	.203 ± .100
Behavioral	.659 ± .274	.318 ± .285
Reply	.718 ± .154	.354 ± .980
ENDEX	.709 ± .048	.259 ± .033

Table 1: RED dataset has two classes, engaging and non-engaging, clustered by applying z-score on ENDEX score. This table shows the mean and standard deviation of sub-dimension scores for both classes; the last row displays the distribution of the overall ENDEX score.

2.1 Reddit Engagement Dataset (RED)

We curate the Reddit Engagement Dataset (RED), a distant-supervision set, with 80k single-turn conversations. We source RED from Reddit, sampling from 43 popular subreddits, and processed a total of 5 million posts, filtering out data that was either non-conversational, toxic, or posts not possible to ascertain popularity; the resulting data distribution of RED is shown in Table 1. The following sections will explain the procedure to automatically annotate ENDEX scores and cluster samples into positive and negative sets.

We also curated a RED testset with 150 human annotated samples obtained from a different split from RED. The inter-annotator agreement is 0.34 Fleiss-Kappa, indicating fair agreement, which reflects the challenge of determining engagingness.

2.2 Distantly-Supervised Engagingness Scores

We use distant-supervision to provide samples in RED an ENDEX Score, which is the aggregate of 4 engaging dimensions. Section 2.2 discusses the intuition for each engagingness dimension; section 2.3 explains how to adjust raw score by thread popularity; section 2.4 lays out the formula to normalize and aggregate sub-dimensions into the overall engagingness score; section 2.5 explains sampling with z-score to convert the task into binary classification.

- **Emotional Engagement (EE):** Emotional connection is a key sign of human engagement (Savin-Baden et al., 2014); and we model EE using a multi-class emotional classifier (Demszky et al., 2020) on post replies. If post receives positive and emotional replies, it’s engaging; negative or neutral replies indicates non-engaging.

- **Attentional Engagement (AE):** More user time spent indicates higher engagement (Attfield et al., 2011). We model AE of a post by examining whether it has edited replies, and the information specificity in the replies.
- **Behavioral Engagement (BE):** Human behavioral features closely correlate with their engagement state (Attfield et al., 2011), and we model BE by examining Reddit post scores, adjusted by popularity.
- **Reply Engagement (RE):** Following definition from (Ghazarian et al., 2020), if a certain post is very likely to be continued by following threads, it is considered engaging; reply_counts are also popularity adjusted.

2.3 Adjustment for Popularity

Raw score for Behavior Engagement(upvotes) and Reply Engagement(reply counts) are heavily influenced by the popularity of the thread in which the post appears. A non-engaging post may receive high user interaction because it simply receives a lot of exposure; on the flip side, a very engaging post may receive zero user interaction simply because it is rarely seen. To mitigate the imbalanced exposure problem, we calculate a popularity value for each thread, and adjust posts scores by the popularity value of the thread it resides.

Popularity Value(PV) The PV of a post is given by the amount of exposure its parent post attracts. Let the target post be θ and its parent σ , R_{reply} obtains the reply counts of a post, and U_{pvote} obtains the upvotes of a post. The PV is defined in equation (1), where coefficient 2 is adopted to give equal weight for reply and upvotes; popularity value adjusted RE score is given by PVRE in equation (6), where M_{pv} and M_{re} are the median of popularity value and reply counts in the entire dataset. Only popularity adjusted scores are used for calculating ENDEX score.

$$PV(\theta) = 2 * R_{reply}(\sigma) + U_{pvote}(\sigma) \quad (1)$$

$$PVRE(re) = re + \frac{M_{pv}}{M_{re}} * \frac{re}{PV(re)} * re \quad (2)$$

2.4 Monotone Submodular Normalization

The final ENDEX score is essentially a weighted sum of the 4 respective sub-dimension scores; an importance nuance is the usage of submodular normalization (shown in Eq. 8) for 3 dimensions to bring raw scores to the scale of 0-1. We observe

that unit increase in raw score lead to diminishing positive effect on engagingness. For example, a sentence with 100 replies should be more engaging than one with 1 replies, but not 99 times more; thus, we normalize engagingness score with a monotone submodular function $f(x) = \frac{x}{x+\alpha}$.

$$N(x) = \left(\sum_{n=1}^3 w_i * \frac{x_i}{x_i + \alpha_i} \right) + w_{EE} * x_{EE}, \quad (3)$$

N is the normalized score for sample x , x_i is x 's raw score on dimension i , where $i \in \{RE, BE, AE\}$; α_i is the median of i -th dimension; w_i is the weight for i dimension; w_{EE} is the weight for EE dimension. The weight can be tuned for your own usage of RED; ³.

2.5 Clustering with z-score

Essentially, engagingness prediction is a classification task, and we want to prepare dataset for binary classification. We use z -score on the ENDEX Score to easily sample and cluster the data according to standard deviation from mean. A confidence threshold κ (ours is 1) needs to be picked, which means that we regard samples that fall between κ standard deviation from mean as uncertain, and are thus discarded. And we cluster positive and negative samples using the following equation (9).

$$Polarity(x) = \begin{cases} 1 & \text{if } z_score(x) > \kappa \\ 0 & \text{if } z_score(x) < -\kappa \end{cases} \quad (4)$$

The EnDex metric is then trained as a binary classification task by finetuning a RoBERTa-large model (Liu et al., 2019) on turn-level RED data.

3 Experiments

3.1 Experiment Set-up

We test the performance of the ENDEX metric on 5 golden evaluation sets that have turn-level labels. Among them, BETTER (Ghazarian et al., 2019), PREDEENG-600 (Ghazarian et al., 2020) are annotated specifically for engagingness with high annotator agreement. BETTER samples are taken from human conversation, while half of PREDEENG-600 are chatbot generations. FED (Mehri and Eskenazi, 2020) annotates dialogue for 9 different dimensions, and we use their engagingness scores as

³For ENDEX, α for three dimensions RE, BE, AE are 1, 2, 18, respectively; we also applied weights of 3, 3, 2, 1 for RE, AE, EE, and BE

Method	BETTER		PREDENG-600		FED-ENG		RED-TEST		GRADE	
	P	S	P	S	P	S	P	S	P	S
Random (<i>ref.</i>)	0.025	0.025	-0.012	-0.013	0.080	0.081	-0.053	-0.053	0.053	0.045
Question Specificity	0.167	0.167	0.073	0.074	0.320	<u>0.320</u>	0.194	0.194	0.009	0.008
USL-H	0.356	0.343	0.688	0.699	0.267	0.277	0.121	0.125	0.234	0.243
PRED_EN	0.234	0.310	-0.137	-0.134	0.250	0.340	0.044	0.178	-0.090	-0.060
PRED_EN (FT+DD)	0.338	0.368	0.390	0.450	0.253	0.195	0.237	0.258	0.194	0.173
Ours: ENDEX	0.414*	0.397*	0.397	0.348	0.235*	0.225*	0.381**	0.375**	0.266	0.248
Ours: ENDEX+NS	<u>0.478*</u>	<u>0.455*</u>	0.597**	0.577**	0.229*	0.214*	<u>0.389**</u>	<u>0.378**</u>	<u>0.308</u>	<u>0.282*</u>
Ours: ENDEX-BEST	0.521	0.511	<u>0.620</u>	<u>0.629</u>	<u>0.286</u>	0.253	0.414	0.405	0.406	0.352

Table 2: The correlation between engagement scores and ground truth human judgment. Best scores are **emboldened** and second-best are underlined. We train ENDEX and ENDEX+NS 10 times and report the mean with * and ** indicating a $stdev < 0.05$ and < 0.03 , respectively. ENDEX-BEST is the best score observed over the 10 runs. Compared to existing metrics, the ENDEX-framework achieves SOTA correlation with human judgement on engagingness, leading by far on our newly proposed RED-TEST dataset with more complex and longer texts than chitchats.

target labels. GRADE (Huang et al., 2020) contains quality annotations for dialogue coherence, and we include it to test whether our model can also have good zero-shot performance on related tasks. Lastly, our own RED-TEST is sourced from Reddit, contains discussions on various topics. A table of evaluation set statistics is provided in the Appendix 3.

3.2 Ablation Study on Engaging Dimensions

To test the robustness of the 4 engagingness dimensions 2.2 of ENDEX, we conducted ablation study to train model using only signals from each of the 4 dimensions. We hypothesize that dimensions with high positive contribution towards final results should have very successful clustering of engaging and non-engaging samples by itself; so, if we train model on data clustered by such dimension, we can still get good performance models.

We train five different models on different subsets of RED. All datasets included the same 40k negative (i.e. non-engaging) samples drawn according to our overall engagement score. However, the other 40k positive (i.e. engaging) samples were selected according to a particular dimension score (e.g. EE, AE, BE, and RE), except for ENDEX, which is our aggregate score model. Figure 2 shows that all four dimensions correlate with engagingness to some degree, but RE, AE, and EE are especially effective. We also observe a synergistic effect of training on a composite score rather than any one dimension individually. The experiment highlights and corroborates the multi-dimensionality of engagingness previously reported in the literature (See

et al., 2019). Overall, having an aggregate score is crucial for successful distant-supervised annotation of negative and positive examples.

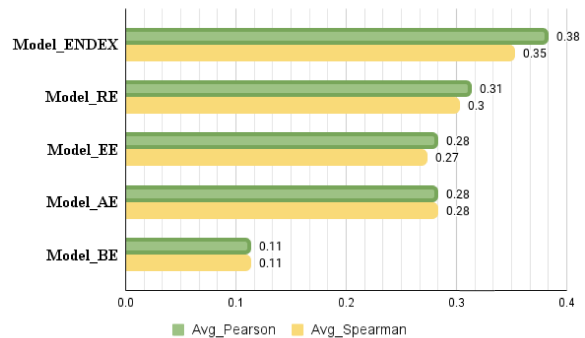


Figure 2: Ablation study of our four engagement dimensions. The ENDEX model was trained on our aggregate engagingness score, while RE, EE, AE, and BE indicate models trains only on scores reflecting that particular dimension.

3.3 Comparison with Related Works

We compare our ENDEX metric, and heuristics-augmented ENDEX+NS metric with five baselines. Three baselines are rule-based, including **Random**, information **Specificity** (See et al., 2019) that counts number of non-stopword tokens, and **Inquisitiveness** (Ghandeharioun et al., 2019) that examines question asking ability. We included them because in some dataset, rule-based system could work surprisingly well (Yeh et al., 2021).

We selected **USL-H** (Phy et al., 2020) as a baseline because it is the top performing metric on the PREDENG-600 and FED engagingness evaluation sets Yeh et al. (2021). USL-H is designed to

measure high-level dialog quality, including understandability, sensibleness, and likability; it trains 3 BERT-based (Devlin et al., 2019) classifiers for each component, and uses a composite score named USL-H for overall assessment. PRED_EN (Ghazarian et al., 2020) uses BERT embedding plus MLP layer and train on ConvAI dataset (Dinan et al., 2019) to make engagement score predictions. PRED_EN (FT+DD) further finetunes the original PRED_EN metric on the DailyDialogue dataset, to get better results.

Our model has two versions: ENDEX is solely trained on *human-reaction* based data. +NS means non-engaging samples set is mixed with some rule-based negative samples, created by random insertion, random deletion, copying, and generic replies;

The experiments in Table 2 demonstrate that our model achieves strong performance on 4 engagingness related datasets, and good correlation with one coherence dataset(GRADE). ENDEX surpasses PRED_EN and USL-H by a large margin on two real human conversations, BETTER and RED-TEST. USL-H still leads in PREDENG-600, and ENDEX+NS’s best model is a close second. Yeh et al. (2021) shows achieving high score on FED-ENG is challenging, with no one surpassing 0.3 spearman in 12 tested metrics. A strong rule-based question detection algorithm surprisingly claims the highest result, and ENDEX a close second.

We find that training solely with *human reaction* distant supervision signals suffices for building competitive models on par or even surpassing mainstream metrics, and it shows better generalization capability in new domains, which seems to echo recent success on modeling human preferences via upvotes in Reddits (Gao et al., 2020).

4 Conclusion

This paper proposes the first *human reaction* based model, ENDEX, to evaluate dialogue engagingness, and curates an 80k Reddit Engagement Dataset (RED) using a novel distant-supervision framework. The success of ENDEX demonstrates the validity of training automatic metrics with *human reaction* signals, offering a strong complement to a synthetic negative sampling approach. We also release an off-the-shelf ENDEX model, and a large scale dataset to facilitate future research.

Limitation

One limitation is that we only curated data for turn-level dialogue. Multi-turn dialogues could also be useful, but it was computationally infeasible to interactively query Reddit for entire threads of conversation. Future work can explore this direction to produce dialogue-level and system-level engagingness metrics.

We also haven’t fully explore our model’s performance on non-dialogue domains, such as on story or creative generations. The training data distribution from the Reddit corpus is diverse enough that it could potentially achieve good performance in non-dialogue settings. A valuable direction of future work is to adapt our method for more general engagingness, or another evaluation metric for open-domain generation.

Ethics

A caveat of using framing our approach around *human attention* is that not all texts attracting high attention are good and ethical. Since being engaging often carries a positive connotation, we made a deliberate design decision to mitigate forms of *negative engagement* in our metric. For example, we assign lower scores to samples flagged by Reddit as controversial, and our behavioral engagement dimension subtracts downvotes from upvotes to punish negative, biased (Liu et al., 2021), and aggressive comments. Moreover, we implemented our emotional engagement algorithm to reward posts with positive emotional replies and punish posts that prompt negative emotions. Future may try to account for the darker aspects of engagingness into our framework and improve the ENDEX metric to differentiate between positive and negative engagement.

Human annotations for RED-TEST were obtained via Amazon Mechanical Turks. We filtered out toxic samples to reduce the likelihood of offensive content and paid \$0.30 USD per instance for an expected hourly wage of \$20 USD.

Acknowledgement

We want to give special thanks to Sarik Ghazarian for great advice and help with the baseline model; we also appreciate the brainstorming session with Hongyuan Yang and the suggestion to use z-score for sampling from Zeyi Wang.

References

- Simon Attfield, Gabriella Kazai, and Benjamin Piwowarski. 2011. Towards a science of user engagement (position paper).
- Mark Chen, Beth E. Kolko, Elisabeth Cuddihy, and Eliana Medina. 2005. Modeling and measuring engagement in computer games. In *DiGRA Conference*.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. *GoEmotions: A dataset of fine-grained emotions*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, abs/1810.04805.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander H. Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur D. Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhumoye, Alan W. Black, Alexander I. Rudnicky, Jason Williams, Joelle Pineau, Mikhail S. Burtsev, and Jason Weston. 2019. *The second conversational intelligence challenge (convai2)*. *ArXiv preprint*, abs/1902.00098.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. Dialogue response ranking training with large-scale human feedback data. *ArXiv*, abs/2009.06978.
- Asma Ghandeharioun, Judy Hanwen Shen, Natasha Jaques, Craig Ferguson, Noah Jones, Àgata Lapedriza, and Rosalind W. Picard. 2019. *Approximating interactive human evaluation with self-play for open-domain dialog systems*. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13658–13669.
- Sarik Ghazarian, Johnny Wei, Aram Galstyan, and Nanyun Peng. 2019. *Better automatic evaluation of open-domain dialogue systems with contextualized embeddings*. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 82–89, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sarik Ghazarian, Ralph M. Weischedel, Aram Galstyan, and Nanyun Peng. 2020. *Predictive engagement: An efficient metric for automatic evaluation of open-domain dialogue systems*. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7789–7796. AAAI Press.
- Sarik Ghazarian, Nuan Wen, A. G. Galstyan, and Nanyun Peng. 2022. *Deam: Dialogue coherence evaluation using amr-based semantic manipulations*. *ArXiv preprint*, abs/2203.09711.
- Laura Hanu and Unitary team. 2020. Detoxify. Github. <https://github.com/unitaryai/detoxify>.
- Lishan Huang, Zheng Ye, Jinghui Qin, Liang Lin, and Xiaodan Liang. 2020. *GRADE: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9230–9240, Online. Association for Computational Linguistics.
- Tian Lan, Xian-Ling Mao, Wei Wei, Xiaoyan Gao, and Heyan Huang. 2020. *Pone: A novel automatic evaluation metric for open-domain generative dialogue systems*. *ArXiv preprint*, abs/2004.02399.
- 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.
- Ruibo Liu, Chenyan Jia, Jason Wei, Guangxuan Xu, Lili Wang, and Soroush Vosoughi. 2021. *Mitigating political bias in language models through reinforced calibration*. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17):14857–14866.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. *Roberta: A robustly optimized bert pretraining approach*. *ArXiv preprint*, abs/1907.11692.
- Xiaojuan Ma. 2018. *Towards human-engaged AI*. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5682–5686. ijcai.org.
- Alison McMahan. 2003. Immersion, engagement, and presence: A method for analyzing 3-d video games. *The Video Game Theory Reader*, pages 67–86.
- Shikib Mehri and Maxine Eskenazi. 2020. *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.

- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. [ParIAI: 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.
- Vitou Phy, Yang Zhao, and Akiko Aizawa. 2020. [Deconstruct to reconstruct a configurable evaluation metric for open-domain dialogue systems](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4164–4178, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.
- Shiki Sato, Reina Akama, Hiroki Ouchi, Jun Suzuki, and Kentaro Inui. 2020. [Evaluating dialogue generation systems via response selection](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 593–599, Online. Association for Computational Linguistics.
- Maggi Savin-Baden, Gemma Tombs, Roy Bhakta, and David J. H. Burden. 2014. Students’ experiences of emotional connection with pedagogical agents.
- Henrik Schoenau-Fog. 2011. The player engagement process - an exploration of continuation desire in digital games. In *DiGRA Conference*.
- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. [What makes a good conversation? how controllable attributes affect human judgments](#). 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 1702–1723, Minneapolis, Minnesota. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Chaklam Silpasuwanchai, Xiaojuan Ma, Hiroaki Shigemasa, and Xiangshi Ren. 2016. [Developing a comprehensive engagement framework of gamification for reflective learning](#). pages 459–472.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. [RUBER: an unsupervised method for automatic evaluation of open-domain dialog systems](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 722–729. AAAI Press.
- Yi-Ting Yeh, Maxine Eskénazi, and Shikib Mehri. 2021. [A comprehensive assessment of dialog evaluation metrics](#). *ArXiv preprint*, abs/2106.03706.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

A Appendix

A.1 RED Processing Steps

Our Reddit data is downloaded from Pushshift.io⁴, and we processed approximately 5M data to curate an 80k sample RED dataset. We deleted posts that do not have an immediate parent thread, because we need pair turn-level data. Our preprocessing removes non-conversational data, such as posts including & gt(reply to symbol). We also removed explicitly toxic data filtered by Detoxify (Hanu and Unitary team, 2020).

We also applied a key data processing trick to reduce noisy signals – the exposure variable. It helps measure the amount exposure each post receives to help normalize its upvote/reply score. We reward posts that are in low-exposure, unpopular threads, while penalizing posts in high-exposure, popular threads, because high upvotes and replies in popular threads may be more due to exposure than true engagingness.

After computing the normalized score given in Equation 8, we also apply another z-score to normalize the final ENDEX score according to standard deviation, so that we can easily sample our data from it. A score with higher standard deviation will imply a higher probability that the sample is

⁴<https://files.pushshift.io/reddit/comments/>

engaging according to our metric. We apply a cut-off to sample high probability engaging and non-engaging samples, and arrive at the RED dataset.

A.2 Model training details

Our RED-TEST set contains 300 human labeled data. The train validation split during training is 0.8 and 0.2.

We used 4 Nvidia A6000 GPUs for training, and 1 Nvidia A6000 GPU for inference. The average runtime for training one model is 2 minutes per epoch, and inference time is in seconds, negligible for the testset. The estimated energy cost per model is, assuming per second gpu energy cost of 245W: $245W * 4 * 60 = 58800$ per model.

We trained our model for 2 epochs, and only save the best checkpoints, with learning rate of $5e-5$ with no extensive hyperparameter search.

We used specificity and question examination inspired from (See et al., 2019); USL-H (Phy et al., 2020) and PREDEN is taken from a GitHub repo⁵ and modified to use a local bert-base-uncased since the original ‘bert-as-service’ code no longer functions.

The Formula for calculating each dimension is given in the following:

- **Reply Engagement;** The raw Reply Engagement score(re) is just the reply counts of a post. Popularity value adjusted RE score is given by PVRE in equation (6), where M_{pv} and M_{re} are the median of popularity value and reply counts in the entire dataset. Please refer to equation (1) for calculation of the popularity value.

$$PVRE(re) = re + \frac{M_{pv}}{M_{re}} * \frac{re}{PV(re)} * re \quad (5)$$

- **Behavioral Engagement;** The raw BE score of a certain post(be) is obtained by subtracting downvotes to upvotes and set to 0 if given controversy flag. Popularity value adjusted BE score is given by PVBE in equation (??), where M_{pv} and M_{be} are the median of popularity value and raw BE score in the entire dataset. Please refer to equation (1) for calculation of the popularity value.

$$PVBE(be) = be + \frac{M_{pv}}{M_{be}} * \frac{be}{PV(be)} * be \quad (6)$$

⁵<https://github.com/exe1023/DialEvalMetrics>

- **Attentional Engagement;** It is calculated using maximum information specificity, or the maximum number of non-stopword tokens in a post’s replies, and whether its children posts are edited; t is the maximum reply specificity, and e stands for number of edited replies.

$$AE(x) = t + 10 * e \quad (7)$$

- **Emotional Engagement;** The EE score is the aggregate probability for all positive emotion categories, produced by the go-emotion classifier (Demszky et al., 2020).

A.3 Submodular Normalization and z-score Clustering

After we obtained the sub-dimension scores, we want to aggregate them into a single normalized ENDEX Score, and lastly cluster them into positive and negative sets to train a binary classifier. The formulas are list in the following:

$$ENDEX(x) = \left(\sum_{n=1}^3 w_i * \frac{x_i}{x_i + \alpha_i} \right) + w_{EE} * x_{EE}, \quad (8)$$

N is the normalized score for sample x , x_i is x ’s raw score on dimension i , where $i \in \{RE, BE, AE\}$; α_i is the median of i -th dimension; w_i is the weight for i dimension; w_{EE} is the weight for EE dimension. The weight can be tuned for your own usage of RED;

A confidence threshold κ (ours is 1) needs to be picked, which means that we regard samples that fall between κ standard deviation from mean as uncertain, and are thus discarded. And we cluster positive and negative samples using the following equation (9).

$$Polarity(x) = \begin{cases} 1 & \text{if } z_score(x) > \kappa \\ 0 & \text{if } z_score(x) < -\kappa \end{cases} \quad (9)$$

A.4 Annotation data and test data

We performed annotation on Amazon Mechanical Turk, and selected annotators based in the United State; in implemented restrictions to annotator with 98% pass rate. We give four examples and clear

WARNING: The data may contain offensive contents. If you feel uncomfortable at any time, feel free to stop.

Our research rewards objective annotations, and rejects every reckless labelling.

Does the Reply express engagement in a dialogue context?

Tip: The Reply is **Engaged** if it satisfies one or more of the following:

1. Emotional/informative Interaction.
2. Interesting Language (funny, creative, using rhetorical devices).
3. Actively seeking conversation, asking good questions.

Tip: The Reply is **Non-Engaged** if it satisfies one or more of the following:

1. Generic replies/not making sense.
2. Plain language, and not actively seeking continuation.
3. Self-talking.

Examples:

Context: "Boten Anna by Basehunte."

Reply: "You are a god and I love you."

Analysis: *The reply is praising in an emotional and explicit way; we can expect the speaker is flattered and may want to respond.*

Verdict: 0. *Engaging.*

Context: "You sir have opened my eyes to a whole new world or reddit. Thanks... Thanks alot..."

Reply: "You are most welcome."

Analysis: *Generic response isn't going to help continue the topic by any means.*

Verdict: 1. *Non-engaging.*

Figure 3: The screenshot of the task description of our Amazon MTurk questionnaire. We have prepared instructions, demonstrations, and proper warning of offensive content.

Set Number (Please ignore this. Recording purpose only): 3

Dialog ID: 1

Context: But how can he steer?

Reply: By rotating the steering wheel in the direction he wishes the car to go.

Yes, the Reply is engaging. No, the Reply is NOT engaged.

Dialog ID: 2

Context: Does anybody have a free how to guide to meditation, that they ve used, and that worked for them?

Reply: Here s mine Sit Cross Legs Think

Yes, the Reply is engaging. No, the Reply is NOT engaged.

Figure 4: The screenshot of the labeling area of our Amazon MTurk questionnaire. Each pair will be labelled by three annotators.

instruction for the task carried out. A screenshot of our annotation interface is provided below.

Table 3 gives summary of the evaluation datasets we used.

Dataset	# of Samples	Context Length	Response Length	Source	Agreement Rate
BETTER	297	6	8	Human	N/A
PREDENG-600	600	12	14	Human+Bot	0.51
FED-ENG	261	26	12	Human+Bot	N/A
RED-TEST	150	16	17	Human	0.34
GRADE	150	12	14	Human	N/A

Table 3: Dataset Statistics for the 5 golden evaluation sets, with number of samples, context-length, response length, and if applicable, inter-annotator agreement rate.