

Investigating Robustness of Dialog Models to Popular Figurative Language Constructs

Harsh Jhamtani^{* 1} Varun Gangal^{* 1} Eduard Hovy¹ Taylor Berg-Kirkpatrick²

¹ School of Computer Science, Carnegie Mellon University

² Computer Science and Engineering, University of California San Diego

{jharsh, vgangal, hovy}@cs.cmu.edu, tberg@ucsd.edu

Abstract

Humans often employ figurative language use in communication, including during interactions with dialog systems. Thus, it is important for real-world dialog systems to be able to handle popular figurative language constructs like metaphor and simile. In this work, we analyze the performance of existing dialog models in situations where the input dialog context exhibits use of figurative language. We observe large gaps in handling of figurative language when evaluating the models on two open domain dialog datasets. When faced with dialog contexts consisting of figurative language, some models show very large drops in performance compared to contexts without figurative language. We encourage future research in dialog modeling to separately analyze and report results on figurative language in order to better test model capabilities relevant to real-world use. Finally, we propose lightweight solutions to help existing models become more robust to figurative language by simply using an external resource to translate figurative language to literal (non-figurative) forms while preserving the meaning to the best extent possible.

1 Introduction

Human frequently employ figurative language such as metaphors (Carbonell, 1982) and idioms (Jackendoff, 1995) for effective and/or stylistic communication. Thus, dialog models interacting with humans should be equipped to handle these forms of communication. However, understanding figurative language might be challenging for machines since figurative constructions often exhibit non-compositional semantics and may rely on shared cultural and common-sense knowledge (Carbonell and Minton, 1983). For example, a powerful GPT2 model fine-tuned on DailyDialog dataset is unable to handle the metaphor ‘built on the sand’ (Figure

^{*} HJ and VG contributed equally for this paper. Order decided by coin flip.

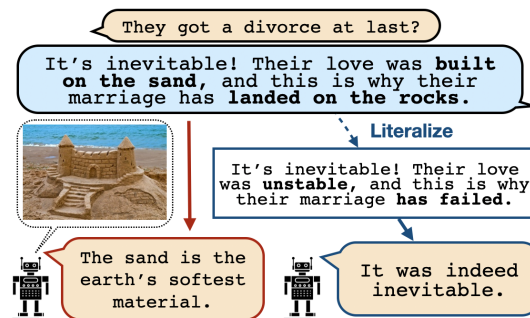


Figure 1: An example illustrating how model responses are affected by figurative constructs in dialog context. Here, the model conflates the metaphorical use of *build on the sand* with its literal meaning, leading to an inappropriate, atypical response.

1), and the response seems to rely on the unintended literal sense of ‘sand’.

In this work, we investigate the performance of existing dialog models when faced with inputs containing figurative language use. (1) First, we identify the subsets in existing datasets (such as DailyDialog (Li et al., 2017) and PersonaChat (Zhang et al., 2018)) which have figurative language use such as metaphors and similes. We observe that the performance of all the dialog models under consideration is lower on such subsets containing figurative language use compared to the dataset as a whole. (2) Second, we gather manually written literal/non-figurative equivalents of the dialog utterances in DailyDialog and PersonaChat which exhibit figurative language use. For example, literal version of ‘on the sand’ can be ‘unstable’ (Figure 1). We observe that performance of dialog models improves when using literal equivalents in place of figurative language. We release the resulting datasets, and encourage that new dialog models be tested separately on such datasets to understand and measure their ability to handle figurative language. (3) Finally, we propose a simple defense against occurrences of figurative language in dialog context. More specifically, we use existing classifiers to detect presence of certain types of figurative language

in dialog contexts, and use dictionary lookups to transform them to their literal counterparts before feeding them to the dialog models. The proposed technique is lightweight, does not require any re-training of the models, and is effective – though gaps still remain, leaving scope for interesting future explorations. ¹

2 Figurative Language In Open Domain Dialog

We experiment with **DailyDialog (DD)** dataset (Li et al., 2017), which is an open domain dialog corpus with 13.1K conversations on colloquial topics like Tourism, Health etc, of which 1K dialogs (6.74K utterances) form the test split. To carry out the desired analysis, we need to first identify the utterances in the dataset which have figurative language use. To achieve high precision labeling, we rely on manual annotations instead of using external figurative language detectors/classifiers. The task was performed manually by two graduate students (native English speakers) studying in a university with English as the primary language of instruction. Additionally, we request the annotators to also write down the literal equivalent versions of the utterances containing figurative language. We release the resulting subset of DailyDialog dataset as **DailyDialog-Figurative (DD-Fig)**, consisting of those dialog instances which contain figurative language, along with two manually written literal versions of each utterance.

Though figurative constructs are only mildly frequent at an utterance level (2.2% in DD), their frequency of occurring at least once in a dialog is ≈ 6 times higher (13.1%). This means that model sensitivity to figurative constructs is *more significant* an issue than *mere utterance-level frequency suggests*, since figurative constructs occurring anywhere in a dialog potentially affect all model responses for that dialog. Additionally, handling figurative constructs is still critical to robust long-tailed performance (Bamman, 2017), which matters for worst-case behaviour and user satisfaction (Goel et al., 2010; Ilievski, 2019). Lastly, most the prevalent figurative constructs we observe are metaphors (45.8%), idioms (47.3%), rhetorical questions (5.3%), hyperbole (6.9%), and personification (3.8%). Appendix §B explores further properties of figurative prevalence.

¹Data and code can be found at <https://github.com/vgtomahawk/Dialog-Fig-Speech-Robust>.

Metric	Model	DD	DD-Fig	% Drop	Rank-change
Bleu-4	GPT2	0.005	0.000	100.00	1
Bleu-4	CVAE	0.012	0.000	100.00	-2
Bleu-4	DualEnc	0.003	0.000	100.00	0
Bleu-4	HRED	0.017	0.009	50.00	-1
Bleu-4	Seq2Seq	0.012	0.015	-23.53	2
Meteor	GPT2	0.130	0.105	18.84	-2
Meteor	CVAE	0.135	0.114	15.54	-1
Meteor	DualEnc	0.111	0.105	5.31	1
Meteor	HRED	0.134	0.114	15.24	-1
Meteor	Seq2Seq	0.121	0.115	5.61	3
Rouge-L	GPT2	0.166	0.146	11.90	0
Rouge-L	CVAE	0.319	0.283	11.47	-1
Rouge-L	DualEnc	0.233	0.214	8.48	0
Rouge-L	HRED	0.329	0.297	9.91	-1
Rouge-L	Seq2Seq	0.315	0.302	4.10	2
Human	GPT2	3.278	2.712	17.27%	0
Human	CVAE	2.302	1.771	23.06%	-1
Human	DualEnc	1.699	1.397	17.79%	0
Human	HRED	2.353	2.115	10.10%	0
Human	Seq2Seq	2.146	1.870	12.85%	1

Table 1: We compare the model response quality for DD-Fig subset compared to the full DD. Across most of the models, we observe a considerable drop in response quality. Additionally, we note that relative ranks of various models as per automated metrics on Daily Dialog dataset show major changes when evaluating just on the figurative subset – i.e. the best model overall might not be the best on the figurative subset, highlighting the need to separately report results on the data subset which has figurative language use.

3 Experiments

Do the dialog models perform worse on the data subset having figurative language use ? We compare model performances on DD vs DD-Fig. For DailyDialog data, we consider model outputs provided by Gupta et al. (2019) from the following methods: CVAE (Zhao et al., 2017), HRED (Serban et al., 2016), Seq2Seq (Vinyals and Le, 2015), Dual-encoder (DualEnc) (Lowe et al., 2015), and GPT2-medium (Radford et al., 2019) fine-tuned on DD. To report automated metrics, we use the multi-reference annotation set collected by Gupta et al. (2019). However, automated metrics may not be well correlated with output response quality (Sai et al., 2020; Gangal et al., 2021). Therefore, we also carry out human evaluations, wherein human annotators (on Amazon Mechanical Turk) are asked to judge the appropriateness of a dialog response on a 1-5 Likert scale. We observe that most of the models perform much worse on DD-Fig (Table 1), with the drop being close to 99% in some cases.

Relative Performance of Models: We additionally investigate if the relative ranks of the mod-

Data/Model:	DD-Fig/GPT2	PC-Fig/GPT
Figurative in last utterance		
Bleu-3/4	+98% / +477%	NA
Meteor/Rouge-L	+5.3% / -3.6%	NA
Human-rating	+12.1%	+17.2%
Figurative anywhere in dialog history		
Bleu-3/4	+13% / +79%	NA
Meteor/Rouge-L	+3.9% / -3.0%	NA
Human-rating	+13.7%	+19.5%

Table 2: We compute the % change in performance when figurative language replaced with manually written literal counterparts in DD-Fig and PC-Fig data subsets. We observe that replacing figurative text in dialog context with literal version leads to improved dialog response quality.

els change when evaluating only on the subset with figurative language use compared to the entire test split. We note that there are some substantial changes in the relative ranks of the models (Table 1). For instance, as per Meteor and human-ratings, Seq2Seq performs better than CVAE on the figurative subset, while doing worse on the complete test dataset. Such changes in relative ranks further highlight the need to separately report results for the proposed data subsets. Interestingly, Seq2Seq improves its relative rankings in general on the DD-Fig subset. We hypothesize this is because it generates very generic responses with very little information overlap with the context.

Does replacing figurative language with semantically equivalent literal translations lead to better performance? Above analysis only reports correlation in performance with contexts containing figurative language. However, there could be certain confounding factors involved. Thus, to make a more direct comparison, we will next compare results when using figurative contexts versus their literal counterparts. To perform this experiment, we utilize the human written literal versions in DD-Fig, and experiment with the GPT2 model (which is the best performing model as per human rating on the overall dataset). We report results under two setups: (1) when figurative language is present in the last utterance of the dialog history, and (2) when figurative language is present anywhere in the dialog history. Human ratings are collected using the same procedure as described for Table 1.

Table 2 shows the main results. For some metrics such as Bleu-4, models perform more than 5 times better when fed with literal translations instead of figurative language. Between the two setups under consideration, we observe slightly higher impact

(as per human evaluation ratings) when one or more figurative language constructs are in use anywhere in the dialog history.

Experiments with Personachat (PC): PersonaChat (Zhang et al., 2018) is a persona grounded dialog corpus, with 1K dialogs (7.8K utterances) forming the test split. We consider a GPT model fine-tuned on the train split of the PC dataset. We follow the same tagging and annotation procedure for the test split, and refer to the resulting dataset as PersonaChat-Figurative dataset (**PC-Fig**). Results in Table 2 demonstrate reduction in performance of the dialog model in human evaluations (Automated overlap metrics in the case of PC are considered unreliable since PC contains only one reference per dialog context).

We notice that compared to DailyDialog, PersonaChat utterances tend to be shorter, more informal, and highly spoken-language like utterances, with fast topic transitions. On replacing figurative language with its literal counterpart into such contexts, the replaced literal text, which is typically English of a more formal and written variety, ends up being much more out of sync with the context in the lexico-syntactic/stylistic sense than it is for DailyDialog. This has a slight downward effect on metrics, offsetting some of the gains from replacing away the figurative language.

4 Mitigation

We propose a lightweight mitigation approach wherein we use existing resources to detect and then construct literal translations for two popular figurative constructs: metaphors and idioms. Thus, the proposed mitigation approach does not require any retraining of dialog models.

Metaphor Detection Through Classifier: We train a metaphor detection classifier based on the VUA corpus (Steen et al., 2010; Gao et al., 2018)². To better generalize to external data via recent contextual models, we skip using model by Gao et al. (2018), and instead learn a classifier C_{bert}^{met} by fine-tuning the *bert-base-uncased* (Devlin et al., 2018) checkpoint from Wolf et al. (2019). On VUA, C_{bert}^{met} gets a test F1 of 0.724, which is close to Gao et al. (2018)’s value of 0.726. Next, we run each test utterance in dilaog dataset through C_{bert}^{met} to get its probability p^{met} of being metaphorical. To retain only more reliable predictions, especially consid-

²<https://github.com/gao-g/metaphor-in-context>

ering domain shift w.r.t VUA, we only choose utterances with $p^{met} > 0.9$. The set of metaphorical utterances thus identified is D_{auto}^{met} .

Idiom Detection Through Lexicon: Idioms are frequently used expressions, which have a fixed, typically non-compositional meaning understood by most native speakers from cultural precedent. We curate a lexicon of 2048 commonly used idioms (e.g. *filling his shoes*) from an online source³ – see Appendix A.2 for more details). All utterances with at least one lexicon entry as a substring are identified to create the set of automatically detected idiomatic utterances, D_{auto}^{idiom} . We unify the sets detected above to form $D_{auto}^{fig} = D_{auto}^{met} \cup D_{auto}^{idiom}$. $|D_{auto}^{fig}|$ constitutes 1520 of 6740 utterances for **DD** (22.5%) and 911 of 7801 utterances for **PC** (11.7%) respectively.

Dictionary Replacement: Wiktionary (Zesch et al., 2008) - the collaboratively created online dictionary, provides a curation of entries corresponding for phrases with “idiomatic”⁴ usages. These entries encompass conventionalized metaphors⁵, idioms, euphemisms, commonly used similes etc. Each entry lists the surface form of the figurative construct paired with a gloss. Glosses are for the most part literal interpretations of the figurative construct. However, they often bear other details like dialect (“US”), etymology (“archaic”) etc, which we remove through simple regex-based rules. This allows direct use of the now-cleaned gloss as a literal interpretation in-context. Furthermore, we expand entries whose surface forms contain uninflected verb forms or unrealized pronouns indicated by *someone*, *one’s* etc, spawning one new entry per pronoun-inflection combination⁶. This yields us a dictionary with 17, 743 tuples of the form $\{fig_i, Lit(fig_i)\}$ ⁷. Finally, for each detected utterance $u \in |D_{auto}^{fig}|$, each matched occurrence of fig_i is replaced by $Lit(fig_i)$, $\forall 1 \leq i \leq n$.

Results: From Table 3, we see that mitigation-based literalization leads to higher quality model responses as per most automatic metrics as well as human evaluation. Though the proposed approach

³<https://www.englishclub.com/ref/Idioms/>

⁴Overloaded use of the term to refer to several figurative phenomena at once, and not just idioms proper.

⁵Commonly used metaphors with a fixed, nearly universally accepted meaning.

⁶We use *pyinflect* python library

⁷See Appendix §C for further analysis of the dictionary

Data/Model:	DD-Fig/GPT2	PC-Fig/GPT
Figurative in last utterance		
Bleu-3/4	+46.3%/+54.1%	NA
Meteor/Rouge-L	+2.0%/+5.7%	NA
Human-rating	+13.7%	+0.4%
Figurative anywhere in dialog history		
Bleu-3/4	+8.1%/+0.6%	NA
Meteor/Rouge-L	+1.4%/+1.4%	NA
Human-rating	+7.7%	+0.3%

Table 3: Mitigation: Contrasting performance of GPT-based models on the (automatically detected) subset D_{auto}^{fig} , containing figurative language, against the same subset but with figurative language replaced with *Wiktionary-mitigated* literal counterparts.

offers only small improvements, it is lightweight in terms of time and memory complexity, and provides reasonably fluent and appropriate interpretations for the figurative constructs covered, since these are sourced from the long-term, collaborative editing underlying Wiktionary. Table 4 shows examples where mitigation-based literalization of the figurative context improves model response quality.

Additionally, we observe that Rouge fails to correlate with Meteor in Table 2, but correlates in Table 3. One possible reason for such behavior is that Wiktionary uses dictionary-like, conservative literalizations, adding new words only as necessary. On the other hand, human annotators literalize more freely without regard for word choice fidelity. Meteor is more robust to variation in word choice, being enabled to capture synonymy and other forms of limited surface form variation. Rouge, being more sensitive however, is immediately dampened on account of this.

The proposed approach is based on simple rule based procedures relying on existing resources, and thus there is scope of multiple future extensions. The detection portion of our approach uses an external classifier and a fixed lexicon to detect metaphors & idioms respectively, leading to D_{auto}^{met} . Considering utterances in DD-Fig as gold, we find the recall of this approach, given by $\frac{|D_{auto}^{met} \cap DD-Fig|}{|DD-Fig|} = 0.31$ to be on the lower end. This may be due to the detection 1) missing out on entirely creative *novel* metaphors, rare idioms as well as other unhandled figures of speech like rhetorical questions. 2) being unable to pick out figures of speech from domain shifted contexts. Extending detection to a wider range of figurative constructs and improved ways to address domain shift are thus important avenues for future work.

Furthermore, the proposed mitigation approach

FIGURATIVE CONTEXT (PC)
maybe we can <i>get together</i> sometime if you are not scare of a 30 year old cougar !
MODEL RESPONSE:
i'm not scared of any <i>cats</i> . i've two dogs
LITERALIZED CONTEXT
maybe we can start dating sometime if you are not scare of a 30 year old cougar !
MODEL RESPONSE:
that's a very interesting idea. i've never met one.
FIGURATIVE CONTEXT (DD)
I think it's time for me to meet my admirer and make him <i>face the music</i> .
MODEL RESPONSE
let's <i>rock and roll</i> .
LITERALIZED CONTEXT
I think it's time for me to meet my admirer and make him bear the consequences of his actions .
MODEL RESPONSE:
let's meet.

Table 4: Examples from PersonaChat(PC) and DailyDialog (DD) illustrating how figurative contexts on literalization produce better (GPT-2) model responses. See Appendix F for additional examples.

is deficient in its ability to literalize all figurative text, since it can't handle: (1) Novel metaphors, which, being creative are unseen/ infrequent in extant corpora, as opposed to *conventionalized* metaphors. (2) Complex figurative constructs which lack particular surface forms and are context sensitive or highly creative - e.g sarcasm & hyperbole. Future work could look into improved technique figurative language detection, perhaps including use of resources such as commonsense knowledge resources (Bosselut et al., 2019) which have been separately shown to be useful in dialog domain (Majumder et al., 2020) and in generating figurative language (Chakrabarty et al., 2020).

5 Related Work

Past work has explored fine-grained analysis and understanding of the performance of dialog models (Roller et al., 2020). Saleh et al. (2020) analyze open domain dialog systems for skills such as inferring contradictions and determining the topic of conversation inter alia. Sankar et al. (2019) analyze the change in perplexity when applying certain perturbations in dialog history. Past work has analyzed dialog models from the point of view of safety from toxic language (Xu et al., 2020; Dinan et al., 2019), and gender biases (Dinan et al., 2020). Gao et al.

(2020) analyze how well dialog models respond to utterances from infrequent sentence function types (e.g *Negative Declarative* utterances like *I feel bad today*). Louis et al. (2020) propose to identify the categorical mapping of an indirect response with respect a polar question in a task oriented dialog setup.

Challenges in handling metaphors and idioms has been explored in prior work on machine translation (Mohammad et al., 2016; Kordoni, 2018; Mao et al., 2018). Mao et al. (2018) propose a method to identify metaphors in English text, and paraphrase them into their literal counterparts before translating to Chinese. Our work on analyzing dialog models against figurative language contexts is along similar direction, though the task setup and scope of figurative language involved are different. Figurative language generation has received reasonable attention such as simile generation (Chakrabarty et al., 2020) and idiom generation (Zhou et al., 2021). Compared to them, our focus is on analyzing capability of popular contemporary dialog models when faced with figurative language.

6 Conclusions

In this work, we demonstrate how existing dialog models fall short in handling figurative language use, and propose a light-weight mitigation technique to ameliorate this lacuna. We encourage future research in dialog models to separately analyze and report model performance

The mitigation techniques used by us are pretty lightweight, but are not able to capture many occurrences of figurative language. Future work could look into improved techniques for figurative language detection. Our work is limited to a couple of open domain dialog datasets in English language. Similar analyses could be conducted on goal oriented dialog setups and datasets in other languages.

Acknowledgements

We thank anonymous EMNLP reviewers for insightful comments and feedback.

Ethics Statement

Our human preference/appropriateness ratings are collected over source content either directly sourced from, or based on typical, off-the-shelf models trained on already existing, publicly available and widely used dialog datasets - namely, DailyDialog (Li et al., 2017) and Personachat (Zhang

et al., 2018) as well as the multiple references dataset from (Gupta et al., 2019). We do collect human evaluation ratings using crowd-sourcing. However, we neither solicit, record or request any kind of personal or identity information from the annotators. Our work primarily performs experiments on dialog in English (Bender and Friedman, 2018). Dialog models are known to suffer from biases learnable from dialog training data, such as gender bias (Dinan et al., 2020). However, our work and contribution does not present or release any new models or model checkpoints, and is primarily concerned with more careful evaluation of a particular phenomena (i.e figurative language), and discussion on lightweight mitigation strategy related to the same.

References

- David Bamman. 2017. Natural language processing for the long tail. In *DH*.
- Emily M Bender and Batya Friedman. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. **COMET: commonsense transformers for automatic knowledge graph construction**. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4762–4779. Association for Computational Linguistics.
- Jaime G Carbonell. 1982. Metaphor: an inescapable phenomenon in natural-language comprehension. *Strategies for natural language processing*, pages 415–434.
- Jaime G Carbonell and Steven Minton. 1983. Metaphor and common-sense reasoning. Technical report, Carnegie Mellon University.
- Tuhin Chakrabarty, Smaranda Muresan, and Nanyun Peng. 2020. Generating similes <effortlessly> like a pro: A style transfer approach for simile generation. *EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Emily Dinan, Angela Fan, Adina Williams, Jack Urbanek, Douwe Kiela, and Jason Weston. 2020. Queens are powerful too: Mitigating gender bias in dialogue generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8173–8188.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. **Build it break it fix it for dialogue safety: Robustness from adversarial human attack**. In *EMNLP-IJCNLP 2019*, pages 4536–4545. Association for Computational Linguistics.
- Varun Gangal, Harsh Jhamtani, Eduard H. Hovy, and Taylor Berg-Kirkpatrick. 2021. **Improving automated evaluation of open domain dialog via diverse reference augmentation**. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, pages 4079–4090. Association for Computational Linguistics.
- Ge Gao, Eunsol Choi, Yejin Choi, and Luke Zettlemoyer. 2018. Neural metaphor detection in context. In *EMNLP*, pages 607–613.
- Yifan Gao, Piji Li, Wei Bi, Xiaojiang Liu, Michael Lyu, and Irwin King. 2020. Dialogue generation on infrequent sentence functions via structured meta-learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 431–440.
- Sharad Goel, Andrei Broder, Evgeniy Gabrilovich, and Bo Pang. 2010. Anatomy of the long tail: ordinary people with extraordinary tastes. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 201–210.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey P Bigham. 2019. Investigating evaluation of open-domain dialogue systems with human generated multiple references. In *Proceedings of the 20th Annual SIG-dial Meeting on Discourse and Dialogue*, pages 379–391.
- Filip Ilievski. 2019. *Identity of Long-tail Entities in Text*, volume 43. IOS Press.
- Ray Jackendoff. 1995. The boundaries of the lexicon. *Idioms: Structural and psychological perspectives*, 133:165.
- Valia Kordoni. 2018. **Beyond multiword expressions: Processing idioms and metaphors**. In *Proceedings of ACL 2018, Melbourne, Australia, July 15-20, 2018, Tutorial Abstracts*, pages 15–16. 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 *IJCNLP*, pages 986–995.
- Annie Louis, Dan Roth, and Filip Radlinski. 2020. **"i'd rather just go to bed": Understanding indirect answers**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7411–7425. Association for Computational Linguistics.

- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. [The Ubuntu Dialogue Corpus: A large Dataset for Research in Unstructured Multi-Turn Dialogue Systems](#). In *Proceedings of the SIG-DIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 2-4 September 2015, Prague, Czech Republic*, pages 285–294. The Association for Computer Linguistics.
- Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian J. McAuley. 2020. [Like hiking? you probably enjoy nature: Personagrounded dialog with commonsense expansions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9194–9206. Association for Computational Linguistics.
- Rui Mao, Chenghua Lin, and Frank Guerin. 2018. [Word embedding and wordnet based metaphor identification and interpretation](#). In *ACL 2018*, pages 1222–1231. Association for Computational Linguistics.
- Saif Mohammad, Ekaterina Shutova, and Peter D. Turney. 2016. [Metaphor as a medium for emotion: An empirical study](#). In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics, *SEM@ACL 2016, Berlin, Germany, 11-12 August 2016*. The *SEM 2016 Organizing Committee.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Stephen Roller et al. 2020. [Open-domain conversational agents: Current progress, open problems, and future directions](#). *CoRR*, abs/2006.12442.
- Ananya B. Sai, Akash Kumar Mohankumar, and Mitesh M. Khapra. 2020. [A survey of evaluation metrics used for NLG systems](#). *CoRR*, abs/2008.12009.
- Abdelrhman Saleh, Tovly Deutsch, Stephen Casper, Yonatan Belinkov, and Stuart M. Shieber. 2020. [Probing neural dialog models for conversational understanding](#). *CoRR*, abs/2006.08331.
- Chinnadhurai Sankar, Sandeep Subramanian, Chris Pal, Sarath Chandar, and Yoshua Bengio. 2019. [Do neural dialog systems use the conversation history effectively? an empirical study](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 32–37. Association for Computational Linguistics.
- I. Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*.
- Gerard J Steen, Aletta G Dorst, J Berenike Herrmann, Anna A Kaal, and Tina Krennmayr. 2010. Metaphor in usage. *Cognitive Linguistics*, 21(4):765–796.
- Oriol Vinyals and Quoc V. Le. 2015. [A neural Conversational Model](#). *CoRR*, abs/1506.05869.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston, and Emily Dinan. 2020. [Recipes for safety in open-domain chatbots](#). *CoRR*, abs/2010.07079.
- Torsten Zesch, Christof Müller, and Iryna Gurevych. 2008. Extracting lexical semantic knowledge from wikipedia and wiktionary. In *LREC*, volume 8, pages 1646–1652.
- 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 *ACL*, pages 2204–2213.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskénazi. 2017. [Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 654–664. Association for Computational Linguistics.
- Jianing Zhou, Hongyu Gong, and Suma Bhat. 2021. From solving a problem boldly to cutting the gordian knot: Idiomatic text generation. *arXiv preprint arXiv:2104.06541*.

A Automatic Detection - Further Details

A.1 Metaphor Detection

A.2 Idiom Detection

Lexicon construction is done through a two step process. First, we curate the lexicon from the mentioned source . Since many entries in this lexicon are templates like *behind someone's back* which could have multiple realizations (e.g *behind her back*), we use a rule-based procedure to expand such instances to all possible realizations by enumerating over all POS combinations (nominative pronoun, nominative pronoun-verb, verb-accusative pronoun etc as applicable). For this, we use the *pyinflect* library.

B Properties of Figurative Utterances - Continued

B.1 Dialog Acts

Since DailyDialog has dialog acts annotated amongst 4 types (*inform*, *question*, *directive*, *commissive*), we analyze how figurative language distributes over these types. Amongst figurative utterances, *inform* is more dominant than overall (58.7% vs 43.2%), while *question* (32.2% vs 22.0%) and *directive* (10.7% vs 17.3%) are underrepresented.

C Wiktionary Mitigation - Resource Description

The final postprocessed version of the resource contains 17743 entries - these are distributed as 88.53% 'idioms', 8.08% euphemisms and 3.39% similes. The 'idioms' here consist of both idioms proper and conventionalized metaphors - it is not easy to provide an exact breakup since Wiktionary does not distinguish between the two.

C.1 Examples

In Table 5, we enlist a few examples from the Wiktionary lexicon for each figurative construct type. The complete resource can be found in our submission materials at *Data/MitigationDictionary*

D Additional Implementation Details

Compute Infrastructure: We used Nvidia RTX 2080 Ti GPU cards to run GPU based models.

Run time: On an average, it took about 5 hours to train GPT2 based models on DD dataset.

Type	Phrase	Literalization
Idiom	a bridge too far	a step or action that is too ambitious
Idiom	a life of its own	an independent existence with some characteristics of life
Idiom	a wild goose never laid a tame egg	most things are inherited and predetermined
Euphemism	aurally challenged	deaf or hard of hearing
Euphemism	between jobs	unemployed
Euphemism	bite the dust	die
Simile	as modern as next week	thoroughly modern
Simile	avoid like the plague	evade or shun
Metaphor	bear the brunt	endure the worst part of something
Metaphor	beat a dead horse	persist or continue far beyond any specific purpose
Metaphor	cede the field	withdraw from any confrontational or potentially confrontational situation

Table 5: Examples from the Wiktionary lexicon resource

Typical inference time ranges from 10-20 minutes.

Number of parameters: GPT2 has approx 120M parameters.

Hyper-parameter search: We varied random seed and the learning rate, when training GPT based models. We use validation loss to pick the best configuration. Best configuration uses initial random seed of 123.

E Additional Data Collection Details

Annotation Framework: A screenshot of the annotation collection task is shared in Figure 2.

Quality Control We restrict to annotators with > 90% HIT acceptance rate. We also perform spot checks and skip the ratings from annotators who seemed not to be adhering to the provided instructions.

F Qualitative Examples

In this section, we enlist qualitative examples from the various stages of analysis performed in our work. Table 7 shows how various models respond

This task requires basic English language understanding.

For each instance, you will have to read the dialog history between two people **A** and **B**. We request you to rate the final utterance from speaker **B** in terms of appropriateness to preceding conversation on a scale of 1-5, with 1 meaning the least appropriate, and 5 being the most appropriate.

1.
Dialog History:
A: \${a1}
B: \${b1}
A: \${a2}
B: \${b2}
A: \${a3}
B: \${b3}

Q. Rate the coherence/appropriateness of the last utterance from B: on a scale of 1-5.

1 2 3 4 5

Figure 2: Annotation Framework to collect human judgement ratings on appropriateness of generated response with respect to dialog context.

to example figurative contexts. Table 6 shows examples where figurative contexts are misinterpreted by models, with the response relying on the unintended literal sense of the construct. Table 8 shows examples where literalization of the figurative context improves quality of the model response.

DIALOG PRE-CONTEXT: sure do . i <i>get a kick out of</i> looking punk .
DIALOG CONTEXT: this one looks really expensive though .
HUMAN WRITTEN RESPONSE: it 's actually quite reasonably priced .
MODEL RESPONSE (GPT2): i'm not one for a kick out.
DIALOG PRE-CONTEXT: not good . my cousin is <i>driving me up the wall</i> .
DIALOG CONTEXT: how so ?
HUMAN WRITTEN RESPONSE: he is always being so loud and annoying .
MODEL RESPONSE (GPT2): you know i'm a good <i>driver</i> .
DIALOG CONTEXT: every time when heroes came back from the battlefield we <i>took our hats off</i> to them .
HUMAN WRITTEN RESPONSE: i think some people still do that today .
MODEL RESPONSE (GPT2): they are just as beautiful as the <i>red haired cap</i> .
DIALOG PRECONTEXT: That dress is very pretty . Why don ' t you like it ?
DIALOG CONTEXT: <i>It ' s too loud</i> .
HUMAN WRITTEN RESPONSE: We've been looking around for many hours.
MODEL RESPONSE (CVAE): I think it 's too noisy for you.

Table 6: Examples from DailyDialog illustrating how various model responses are affected by figurative constructs. DIALOG PRE-CONTEXT here refers to the utterance before the immediate previous utterance.

No	Type	Text
105-0	CONTEXT	A: How come it is slow as a snail today ?
	GPT2	B: i think so. it is a virus.
	CVAE	B: it 's a skill . i have to work out the nerve .
	DUALENC	B: we are basically going to check your heart , lungs , blood sugar levels , and eyes , ears , and nose .
	HRED	B: it 's a good idea .
SEQ2SEQ	B: i 'm afraid you have to .	
1-1	CONTEXT	A: What would the roses cost me ? B: They're only \$ 20 a dozen. A: I like the sound of that
	GPT2	B: i 'll take one.
	CVAE	B: you can always get rid of the foods .
	DUALENC	B: you 'll only have to pay the initial fee .
	HRED	B: that 's a good idea .
SEQ2SEQ	B: i 'll take it .	

Table 7: Response utterances from five different models to various figurative contexts

<p>FIGURATIVE CONTEXT (PC)</p> <p>i am from new york. do i get <i>brownie points</i> for that? haha</p> <hr/> <p>MODEL RESPONSE:</p> <p>no, i don't get <i>points</i> for anything. do you have any hobbies?</p> <hr/> <p>LITERALIZED CONTEXT</p> <p>i am from new york . do i get any favours for that ? haha</p> <hr/> <p>MODEL RESPONSE</p> <p>i've a job at a restaurant.</p>
<p>FIGURATIVE CONTEXT (PC)</p> <p>maybe we can get together sometime if you are not scare of a 30 year old <i>cougar</i> !</p> <hr/> <p>MODEL RESPONSE:</p> <p>i'm not scared of any <i>cats</i>, i've two dogs</p> <hr/> <p>LITERALIZED CONTEXT</p> <p>maybe we can start dating sometime if you are not scare of a 30 year old cougar !</p> <hr/> <p>MODEL RESPONSE:</p> <p>that's a very interesting idea. i've never met one.</p>
<p>FIGURATIVE CONTEXT (DD)</p> <p>I think it's time for me to meet my admirer and make him <i>face the music</i> .</p> <hr/> <p>MODEL RESPONSE</p> <p>let's <i>rock and roll</i>.</p> <hr/> <p>LITERALIZED CONTEXT</p> <p>I think it's time for me to meet my admirer and make him bear the consequences of his actions .</p> <hr/> <p>MODEL RESPONSE:</p> <p>let's meet.</p>

Table 8: Examples from PersonaChat(PC) and Daily-Dialog (DD) illustrating how figurative contexts on literalization produce better (GPT-2) model responses.