

Just-DREAM-about-it: Figurative Language Understanding with DREAM-FLUTE

Yuling Gu, Yao Fu, Valentina Pyatkin, Ian Magnusson,
Bhavana Dalvi Mishra, Peter Clark

Allen Institute for AI, Seattle, WA
yulingg@allenai.org

Abstract

Figurative language (e.g., “he flew like the wind”) is challenging to understand, as it is hard to tell what implicit information is being conveyed from the surface form alone. We hypothesize that to perform this task well, the reader needs to mentally elaborate the scene being described to identify a sensible meaning of the language. We present *DREAM-FLUTE*, a figurative language understanding system that does this, first forming a “mental model” of situations described in a premise and hypothesis before making an entailment/contradiction decision and generating an explanation. *DREAM-FLUTE* uses an existing scene elaboration model, DREAM, for constructing its “mental model.” In the FigLang2022 Shared Task evaluation, *DREAM-FLUTE* achieved (joint) first place (Acc@60=63.3%), and can perform even better with ensemble techniques, demonstrating the effectiveness of this approach.¹ More generally, this work suggests that adding a reflective component to pretrained language models can improve their performance beyond standard fine-tuning (3.3% improvement in Acc@60).

1 Introduction

Understanding figurative language is a particularly challenging problem in NLP since the underlying meaning of the utterance is very different from the surface meaning of its constituent words (Stowe et al., 2022). In this paper we focus on the task of recognizing and explaining textual entailment between a premise and hypothesis involving figurative language (FigLang 2022 Shared Task in Chakrabarty et al., 2022). We propose *DREAM-FLUTE*,² a system that makes use of scene elaboration for building a “mental model” of the situations

¹We make our code and models publicly available at <https://github.com/allenai/dream>.

²Using DREAM (Gu et al., 2022) on FLUTE: Figurative Language Understanding through Textual Explanations (Chakrabarty et al., 2022).

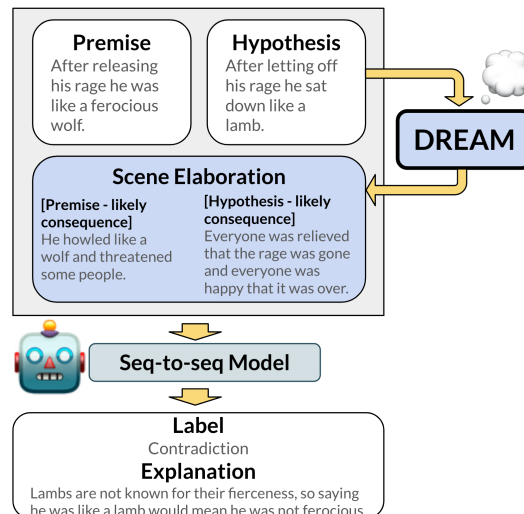


Figure 1: Overview of *DREAM-FLUTE*: It first uses DREAM (Gu et al., 2022) to generate an elaboration of the situation in the premise and hypothesis (separately), then uses this additional context for entailment classification and explanation generation. *DREAM-FLUTE* (consequence), using the “likely consequence” elaboration dimension as additional context, achieved top scores. Such systems also form the building blocks of *DREAM-FLUTE* (ensemble), our best system.

presented in the premise and hypothesis to detect textual entailment between them (see Figure 1).

The design of *DREAM-FLUTE* builds upon the scene elaboration model, DREAM, presented by Gu et al. (2022). DREAM uses a T5-based (Raffel et al., 2020) sequence-to-sequence model to generate additional, pertinent details about each given situation in the input text, along key conceptual dimensions informed by cognitive science, story understanding and planning literature (Minsky, 1974; Dyer, 1983; Mueller et al., 1985; Mueller, 1990). Using such scene elaboration as additional context has been shown to improve question-answering (QA) performance on different models and across different downstream tasks such as ETHICS (Hendrycks et al., 2021), CODAH (Chen et al., 2019) and Social IQA (Sap et al., 2019).

To adapt it for the figurative language understanding shared task, we made three significant extensions to using DREAM that have not been previously explored. First, we incorporate DREAM for elaborating the premise and hypothesis in a natural language inference (NLI) task involving figurative language understanding (Chakrabarty et al., 2021; Stowe et al., 2022). We hypothesize that such additional, pertinent details could also improve a model’s ability to judge whether there is an entailment or contradiction between the premise and hypothesis. This could be especially helpful for the instances that use figurative language, where the underlying meaning might be opaque to the model: further elaborating the context can make certain inferences more explicit. Second, beyond improvements on label prediction accuracy (i.e. choosing from multiple-choice options) shown in Gu et al. (2022), our work uncovers the use of such additional context for improving explanation quality. And lastly, we exploit the dimensions in DREAM to train different models for an ensemble system representing a cognitive continuum (Figure 2), further improving accuracy and explanation quality.

Our approach is easily adaptable to other language models, and task-agnostic in format (e.g. QA or NLI) and domain (e.g. ethical decisions or figurative language understanding). We demonstrate the effectiveness of our single model system in terms of achieving top scores in the task, as well as the flexibility of implementing an ensemble system that not only yields further improvements for this task but also allows customization to suit the requirements of different downstream applications.

2 Approach

We first describe our single model systems in Section 2.1. Next, we present a two-step “classify then explain” pipeline in Section 2.2. In Section 2.3, we take advantage of all information learned by the different models and propose an ensemble approach inspired by cognitive science.

2.1 Single Model Systems

Given an input <Premise, Hypothesis> sentence pair, the task has two goals: (1). first classify the relationship between the premise and hypothesis (*entailment* or *contradiction*); then (2). generate a textual explanation about why the premise entails/contradicts the hypothesis. Figure 1 shows an example. We further consider two additional

pieces of information for performance improvements: (1). the type of the figurative language (*simile*, *metaphor*, *sarcasm*, *idiom*, and *creative paraphrase*) which is provided in the training data (but not the test data); (2). the elaboration of situations in the premise-hypothesis pair provided by DREAM, which gives additional information about the *consequence*, *emotion*, *motivation*, or *social norm* of the input. In Appendix A, we provide intuitive examples showing why such additional information could help this figurative language task.

System 1: Using original data Given the <Premise, Hypothesis, Label, Explanation> in the original data, we first trained a sequence-to-sequence model for the figurative language task using the following input-output format:

Input <Premise> <Hypothesis>

Output <Label> <Explanation>

System 2: Jointly predicting the type of figurative language Using type of figurative language provided as part of the training set (Chakrabarty et al., 2022), one of our models jointly predicts the type of figurative language, together with the target label and explanation:

Input <Premise> <Hypothesis>

Output <Figurative-Language-Type> <Label>
<Explanation>

Systems 3: DREAM-FLUTE - Providing DREAM’s different dimensions as input context

We adapt DREAM’s scene elaborations (Gu et al., 2022) for the figurative language understanding NLI task by using the DREAM model to generate elaborations for the premise and hypothesis separately. This allows us to investigate if similarities or differences in the scene elaborations for the premise and hypothesis will provide useful signals for entailment/contradiction label prediction and improving explanation quality. Figure 1 gives an overview of such systems and the input-output format is:

Input <Premise> <Premise-elaboration-from-DREAM> <Hypothesis> <Hypothesis-elaboration-from-DREAM>

Output <Label> <Explanation>

where the scene elaboration dimensions from DREAM are: *consequence*, *emotion*, *motivation*, and *social norm*. We also consider a system incorporating all these dimensions as additional context.




Figure 2: A cognitive continuum implemented to account for different levels of intuition and analysis.

2.2 Two-step System: Classify then explain

In contrast to Systems 1 to 3 where the entailment/contradiction label and associated explanation are predicted jointly, System 4 uses a two-step “classify then explain” pipeline. Previous work on generating explanations have discussed the difference between predicting and generating respective rationalizations in a pipeline vs. jointly. [Wiegreffe et al. \(2021\)](#) showed that for reasoning tasks pipelines work less well than models which jointly predict and explain. [Hase et al. \(2020\)](#) compared rationalizing methods (first predict label and then the explanation) to reasoning methods (predict the explanation first), and showed that rationalization methods perform better. It is therefore of interest to compare such different approaches for explanation generation also for the figurative language task.

2.3 Ensemble System: A cognitive continuum

We take advantage of ensembling to use information learned by Systems 1 to 4 together in *DREAM-FLUTE* (ensemble). For entailment/contradiction label prediction, the top 5 system variants were chosen based on validation Acc@0 (Table 1 *green italicized*) scores, and used for majority voting.

[Brachman and Levesque \(2022\)](#) note that several psychologists claim “there is a *cognitive continuum* between endpoints that they call *intuition* and *analysis*.” Likewise, in rationalizing, our different system variants can be viewed as different points on this continuum. For generating explanations, Systems 1 to 4 were used as building blocks for *DREAM-FLUTE* (ensemble) (excluding the model with social norm due to its low scores on the validation set) to implement such a continuum that includes various levels of intuition and analysis (Figure 2). Specifically, given the entailment label from majority voting, the ensemble looks for the first of the ordered models that agrees with the ensemble label, then uses its explanation.

Our approach first considers more salient factors (Systems 2, 3 (consequence, emotion)) which can

inform the content and style of explanation: likely consequence of the actions and the emotions of characters, which can possibly tease apart whether the sentence pairs entail/contradict,³ as well as type of figurative language which can inform the style of explanation.⁴ Next, we take a step back and look at the bigger picture, in considering all DREAM dimensions ([Gu et al., 2022](#)) (System 3 (all dimensions)). Then we examine some of the less salient dimensions more closely (Systems 3 (motivation), 4). And finally, we use the explanation in the case when there is no context at all (System 1). More details about this ordering and the pseudocode for ensembling can be found in Appendix C.

3 Experiment Settings

Data This shared task has a two-phases timeline: the development phase then the test phase. During the development phase, ~ 7500 samples are provided as the training set. We used a 80-20 split to create our own training (6027 samples) and validation (1507 samples) partitions on which we build our models. Later at the test phase, separate 1500 test samples (without gold labels) are released on which all models are tested. Note that our model is primarily developed during the training phase without having access to the test data.

Model We train all models with a T5-3B backbone using the data formats detailed in Section 2.1. The size of the model is the same as the officially provided fine-tuned T5 baseline. We use the Huggingface implementation ([Wolf et al., 2019, 2020](#)), based on PyTorch ([Paszke et al., 2019](#)). For each system, we fine-tune the 3B version of T5 ([Raffel et al., 2020](#)) for 3 epochs using an Adam Optimizer and a learning rate of $5e-05$, selecting the best checkpoint based on the lowest validation loss.

³E.g. If one situation involves an action leading to good outcome whereas another leads to bad outcome, that is a clear sign (that gives you strong intuition) for contradiction. Whereas, if the premise and hypothesis both describe situations where a person would be happy, that provides intuition for entailment. See Table 2 for examples from task data.

⁴See Appendix A and Table 3.

System		Our validation partition			Official test partition		
		Acc@0	Acc@50	Acc@60	Acc@0	Acc@50	Acc@60
T5-3B (official baseline)		–	–	–	76.7	69.1	44.3
1	Original data	<i>94.8</i>	89.0	66.9	94.7	88.7	60.4
2	+ Figurative language type	<i>94.9</i>	89.8	66.5	94.6	87.8	61.3
3	<i>DREAM-FLUTE</i>						
	emotion	94.2	89.3	65.0	93.9	88.3	61.7
	motivation	<i>95.4</i>	90.2	66.2	94.5	87.7	60.3
	consequence	94.3	90.1	65.8	94.7	88.9	63.3
	social norm	93.1	88.3	64.2	92.3	86.4	60.6
	all 4 dimensions	<i>95.2</i>	89.4	66.6	94.3	87.7	60.0
4	Classify then explain	<i>95.0</i>	90.5	66.6	95.1	89.4	61.1
5	<i>DREAM-FLUTE</i> (ensemble)	96.4	92.1	67.0	95.9	89.8	63.7

Table 1: Results on our validation set and the official test set. Amongst the non-ensemble methods, System 3 with likely consequence, i.e. *DREAM-FLUTE* (consequence), performed the best on the test set in terms of Acc@60 which was used for ranking submissions on the leaderboard. This system was already ranked first, but further gains can still be achieved using ensembling in System 5, *DREAM-FLUTE* (ensemble). *Green italics* indicates systems selected for label prediction in the ensemble system, using validation Acc@0.

A more detailed list of hyperparameters used can be found in Appendix D.

Evaluation There are two major evaluation metrics: (1). *accuracy*, which measures if predicted NLI labels are correct; (2). *explanation score*, which measures if generated explanations are of high quality. The explanation score is computed as the average of BERTScore (Zhang et al., 2020) and BLEURT (Sellam et al., 2020) on the generated explanation against given references. The overall performance metric, Acc@ s (Table 1), is a combination of *accuracy* and *explanation score* where a prediction (label and explanation) counts as correct only when: (a) the label is correct, and (b) the explanation score is at least s (where $s = 0, 50$ and 60). On the official leaderboard, all models are ranked according to Acc@60.

4 Results and Discussion

4.1 Better explanation quality

Table 1 shows the performance of our systems. Based on test Acc@60, the following strategies improve explanation quality compared to the setup with just the original data: predicting figurative language type, using emotion, likely consequence, social norm, two-step “classify then explain” pipeline, and ensembling. Each non-ensemble system can be seen as guiding the model to focus on a particular direction when reasoning about the entailment/contradiction relationship between a sentence pair. Table 2 and Appendix F present examples of how each DREAM dimension helps uncover implicit meaning in the input. *DREAM-FLUTE*

(consequence), by incorporating the likely consequence scene elaboration from DREAM, was already ranked first based on test Acc@60,⁵ which requires explanations to be of high quality. Figure 1 shows another example of how elaborating along this dimension can be useful. On top of that, *DREAM-FLUTE* (ensemble), an ensemble system that makes further use of context achieves further improvements (Acc@60 = 63.7%). The ensemble approach allows for considering these different directions and rationalizing with varying levels of intuition and analysis, then choosing one that fits the current sentence pair, potentially boosting explanation quality.

4.2 Better label prediction accuracy

This ensemble system is also our best submission overall with Acc@0 = 95.9%, Acc@50 = 89.8%, where Acc@0 is equivalent to computing label accuracy alone. The better label prediction accuracy could be attributed to using the different individual systems for majority voting, which mimics arriving at a decision by considering different perspectives, ultimately leading to a more well-thought decision.

4.3 Effect of DREAM generated context

We qualitatively analyze how DREAM-generated context helps in the figurative language understanding task. Table 2 presents examples from our validation set for DREAM dimensions “consequence” and “emotion” comparing predictions from System

⁵See results of shared task at <https://codalab.lisn.upsaclay.fr/competitions/5908#results>.

Premise	Hypothesis	Gold label	T5’s answer + explanation	DREAM context	<i>DREAM-FLUTE</i> ’s answer + explanation
My decision-making skills are not purely based on emotions and gut.	My personal feelings color my judgment in this case.	(C)	(E): My personal feelings color my judgment in this case means that I am influenced by my emotions and gut when making decisions.	[Consequence] Pre: I make more balanced and informed decisions. Hypo: I make a decision that is not in the best interests of the company.	(C): To have personal feelings color one’s judgment means to make decisions based on them, but this context describes making decisions based on facts and not emotions.
We laid in the field of green grass and relaxed.	We laid in fields of gold.	(E)	(C): Gold is a type of metal that is used to make jewelry, while green grass is a type of grass that is used for lawns and gardens.	[Emotion] Pre: I (myself)’s emotion is happy. Hypo: I (myself)’s emotion is happy.	(E): Gold is a color that is associated with happiness, so the fields of gold are associated with happiness.

Table 2: Examples from the validation set where DREAM generated context consisting of consequence and emotion helped our proposed model *DREAM-FLUTE* (System 3) in figurative language understanding. For both the examples a T5-based model that did not have access to additional context (System 1) gave wrong label prediction. DREAM context helped improve both answer accuracy and explanation quality. Labels: (E), (C) refer to Entailment, Contradiction respectively. (Appendix F presents examples where motivation, social norm helped *DREAM-FLUTE*.)

1 (trained using just original data) with those from System 3 (*DREAM-FLUTE*, which uses scene elaboration from DREAM). These examples illustrate that similarities and differences along the scene elaboration dimensions provide useful signals to guide entailment/contradiction label prediction and improve explanation quality.

4.4 More flexibility beyond FigLang2022

The day-to-day mental activities of humans take place on different parts of the cognitive continuum (Brachman and Levesque, 2022). DREAM’s scene elaborations give us the different building blocks to implement to such a continuum, and therefore use various levels of intuition and analysis to better come to a decision and rationalize. This approach also allows customization to suit the requirements of different downstream applications, by changing the order of factors to consider on the continuum (e.g. social norm may be more salient for ethical decisions) and considering different pertinent factors (i.e. in place of the figurative language type).

5 Conclusion

In this work we showed how *DREAM-FLUTE*, a competitive system for the figurative language

understanding NLI task, can be built by utilizing scene elaborations from an existing model, DREAM. Compared to a model without such scene elaborations, *DREAM-FLUTE* makes use of scene elaboration for building a “mental model” of situations in the premise and hypothesis to make inferences more explicit, thus improving label prediction accuracy and explanation quality. *DREAM-FLUTE* (ensemble) uses different elaborations to form building blocks for implementing a continuum with varying levels of intuition and analysis, modeling deriving answers and rationalizing by considering different positions on a cognitive continuum. This novel use of DREAM not only obtained the highest scores for the figurative language understanding shared task, but could also easily be applied to the situational QA tasks in Gu et al. (2022), and beyond. Our approach is easily adaptable to other language models, and task-agnostic in format (e.g. QA or NLI) and domain (e.g. ethical decisions or figurative language understanding). More generally, our work demonstrates that adding a reflective component helps to improve answer accuracy and explanation quality in pretrained language models.

Limitations

Our approach is designed for applications involving natural language understanding for short text (around 1-3 sentences), e.g. in the figurative language NLI task and situational QA tasks tackled in the original DREAM paper. Building on a better understanding for short text, we hope our work can inspire future efforts towards extending the approach for long text too. The current approach presented also requires the use of GPU resources for model training. However, we also demonstrate that using DREAM scene elaboration as additional context yields improvements on label prediction accuracy for an off-the-shelf NLI model, without any training (Table 4 in Appendix E).

Ethics Statement

Like any other large-scale language model, despite the best intentions, there is a risk of our models producing biased or offensive statements as part of the free-form rationalization. We release our models for research purposes only.

Acknowledgements

We would like to thank the entire Figurative Language Understanding Shared Task organizing committee for organizing this shared task. We thank the anonymous reviewers for their helpful comments. This work was done as part of a Hackathon project during AI2’s 2022 Hackathon. We are grateful to the Hackathon organizers, Caitlin Wittlif and Carissa Schoenick, for the great 3-day Hackathon that led to this work.

References

- R.J. Brachman and H.J. Levesque. 2022. *Machines like Us: Toward AI with Common Sense*. MIT Press.
- Tuhin Chakrabarty, Debanjan Ghosh, Adam Poliak, and Smaranda Muresan. 2021. [Figurative language in recognizing textual entailment](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3354–3361, Online. Association for Computational Linguistics.
- Tuhin Chakrabarty, Arkadiy Saakyan, Debanjan Ghosh, and Smaranda Muresan. 2022. [Flute: Figurative language understanding through textual explanations](#).
- Michael Chen, Mike D’Arcy, Alisa Liu, Jared Fernandez, and Doug Downey. 2019. [CODAH: An adversarially-authored question answering dataset for common sense](#). In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 63–69, Minneapolis, USA. Association for Computational Linguistics.
- Michael G. Dyer. 1983. The role of affect in narratives. *Cogn. Sci.*, 7:211–242.
- Yuling Gu, Bhavana Dalvi, and Peter Clark. 2022. [DREAM: Improving situational QA by first elaborating the situation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1115–1127, Seattle, United States. Association for Computational Linguistics.
- Peter Hase, Shiyue Zhang, Harry Xie, and Mohit Bansal. 2020. Leakage-adjusted simulatability: Can models generate non-trivial explanations of their behavior in natural language? In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4351–4367.
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2021. [Aligning ai with shared human values](#). *ICLR*.
- 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](#).
- Marvin Minsky. 1974. *A framework for representing knowledge*.
- Erik T Mueller. 1990. *Daydreaming in humans and machines: a computer model of the stream of thought*. Intellect Books.
- Erik T Mueller, Michael G Dyer, et al. 1985. Daydreaming in humans and computers. In *IJCAI*, pages 278–280.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An Imperative Style, High-Performance Deep Learning Library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.

- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). 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 4463–4473, Hong Kong, China. 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.
- Kevin Stowe, Prasetya Utama, and Iryna Gurevych. 2022. [IMPLI: Investigating NLI models’ performance on figurative language](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5375–5388, Dublin, Ireland. Association for Computational Linguistics.
- Sarah Wiegrefe, Ana Marasović, and Noah A Smith. 2021. Measuring association between labels and free-text rationales. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10266–10284.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv*, abs/1910.03771.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. 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.
- 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*.

A Examples from training set

We randomly sampled around 100 examples from the training set and manually looked at the targeted explanations to get a sense of how explanations for this task look like. We observed that the explanation style may depend on the type of figurative language involved. Table 3 shows some of these examples. For instance, when the type of figurative language is sarcasm, the explanation often starts by describing what is usually the case and then goes into how one of the sentences describes an unusual or unexpected situation. Whereas, if the type is idiom, then the explanation often involves elucidating what the idiom means. This motivated the design of System 2.

Further, we noticed that the gold explanations often involve elements like emotion and motivation of characters. In the first example in Table 3, for example, identifying the emotions in the premise and hypothesis directly helps us identify the contradiction — in that the person’s emotion is scared in one case and fearless in another. Therefore, we explored elaborating the situations in the given premise and hypothesis along such dimensions using DREAM (Gu et al., 2022). By using DREAM to generate scene elaborations and using that as additional context to the input, we have the different variations of *DREAM-FLUTE* (System 3).

B Details of input prompt

In training our T5 based sequence-to-sequence models, whenever the target output is the entailment/contradiction label and explanation, we append the question “Is there a contradiction or entailment between the premise and hypothesis?” to the input to prompt the model for the NLI task. In the case of System 2, where the model jointly predicts the type of figurative language then the label and explanation, we first append the question “What is the type of figurative language involved?” to the input, then append the usual contradiction or entailment question.

C Algorithm for ensembling

The order of systems used in rationalizing when implementing the cognitive continuum described in Section 2.3 is as follows: likely consequence, emotion, type of figurative language, all DREAM dimensions, motivation, two-step “classify then explain,” no context. Algorithm 1 shows more

Algorithm 1: Ensemble - a cognitive continuum

```
Input: Individual systems’ predicted label and explanation
Output: Ensemble label; Ensemble explanation
ensemble_label =
  majority_vote(top5_Acc@0_systems_labels)
ensemble_explanation = None
// ordered_systems takes an order
  described in Section C
for system_prediction ∈ ordered_systems do
  if system_prediction.label == ensemble_label
  then
    ensemble_explanation =
      system_prediction.explanation
    break
  end
end
```

details on how to obtain the ensemble label and explanation from the individual systems.

Note that beyond the figurative language understanding task, this ensembling approach representing a cognitive continuum could be applied to other tasks, with the possibility of modifying the order of component systems to better suit different applications.

D Hyperparameters used during training

The following hyperparameters were used during training:

- learning_rate: 5e-05
- train_batch_size: 1
- eval_batch_size: 1
- seed: 42
- distributed_type: multi-GPU
- num_devices: 2
- total_train_batch_size: 2
- total_eval_batch_size: 2
- optimizer: Adam with betas=(0.9,0.999) and epsilon=1e-08
- lr_scheduler_type: linear
- num_epochs: 3.0

Type of figurative language	Premise	Hypothesis	Gold label	Gold Explanation
Sarcasm	Yesterday two gangs were fighting just in front of my home.	Yesterday I saw two gangs fighting right in front of my house and it totally didn't make me scared at all.	Contradiction	The sight of two gangs fighting is often very violent and can invoke fear in people, so someone who saw it and wasn't scared is not being truthful.
Idiom	If you want fresh food, just go with your gut feeling and you will find villagers happy to sell or trade what they have.	If you want fresh food, just follow your noses and you will find villagers happy to sell or trade what they have.	Entailment	To follow your nose means to trust one's instinct, which is what you would need to do in order to find fresh food.

Table 3: Examples from Chakrabarty et al. (2022)'s training set. Text in bold illustrate how the style of explanation may depend on the type of figurative language involved.

E Baseline: Off-the-shelf MNLI model

Without any training on the task data, we can similarly achieve better label prediction accuracy if we provide additional context from DREAM as input. Table 4 shows that with the off-the-shelf RoBERTa MNLI model (Liu et al., 2019), we achieve improvements in accuracy when providing the emotion of characters, and even more improvements if we provide all 4 dimensions generated by DREAM. Since this model is unable produce any explanations, we measure only Acc@0 scores.

System	Our validation partition		
	Acc@0	Acc@50	Acc@60
RoBERTa MNLI	73.9	–	–
+ DREAM emotion	77.4	–	–
+ DREAM 4 dimensions	79.3	–	–
<i>DREAM-FLUTE</i> (ensemble) (our model)	96.4	92.1	67.0

Table 4: Comparing off-the-shelf RoBERTa MNLI model (Liu et al., 2019) to our proposed model on our validation set.

F Analysis: Effect of DREAM generated context

Table 5 presents validation set examples for DREAM dimensions “motivation” and “social norm” where this additional context helped the model correct its prediction.

Premise	Hypothesis	Gold label	T5's answer + explanation	DREAM context	<i>DREAM-FLUTE</i> 's answer + explanation
After years of service & contribution to the company, he was finally promoted.	The company released him after many years of service.	(C)	(E): The company released the employee after many years of service.	[Motivation] Pre: Company's motivation is to recognize his hard work. Hypo: Company's motivation is to get someone else to work.	(C): To release someone means to let them go from a position, while to promote someone means to give them a higher position.
Sure, he snorted just to make me feel even better about the already great situation.	Sure, he snorted, just rub it in.	(C)	(E): To rub it in means to make someone feel better about an already great situation, which is what he did by snorting.	[Social Norm] Pre: It's good to make people feel better about a situation. Hypo: It's rude to rub something in someone's face when they don't want to.	(C): To rub it in means to make someone feel bad about themselves, but in this sentence he is making the speaker feel better about the already great situation.

Table 5: Examples from the validation set where DREAM generated context consisting of motivation and social norm helped our proposed model *DREAM-FLUTE* (System 3) in figurative language understanding. For all these examples a T5-based model that did not have access to additional context (System 1) gave wrong label prediction. DREAM context helped improve both answer accuracy and explanation quality. Labels: (E), (C) refer to Entailment, Contradiction respectively.