

DREAM: Improving Situational QA by First Elaborating the Situation

Yuling Gu*, Bhavana Dalvi Mishra, Peter Clark

Allen Institute for AI, Seattle, WA

{yulingg, bhavanad, peterc}@allenai.org

Abstract

When people answer questions about a specific situation, e.g., “I cheated on my mid-term exam last week. Was that wrong?”, cognitive science suggests that they form a mental picture of that situation before answering. While we do not know how language models (LMs) answer such questions, we conjecture that they may answer more accurately if they are also provided with additional details about the question situation, elaborating the “scene”. To test this conjecture, we train a new model, DREAM, to answer questions that elaborate the scenes that situated questions are about, and then provide those elaborations as additional context to a question-answering (QA) model. We find that DREAM is able to create better scene elaborations (more accurate, useful, and consistent) than a representative state-of-the-art, zero-shot model (Macaw). We also find that using the scene elaborations as additional context improves the answer accuracy of a downstream QA system, including beyond that obtainable by simply further fine-tuning the QA system on DREAM’s training data. These results suggest that adding focused elaborations about a situation can improve a system’s reasoning about it, and may serve as an effective way of injecting new scenario-based knowledge into QA models. Finally, our approach is dataset-neutral; we observe improved QA performance across different models, with even bigger gains on models with fewer parameters.¹

1 Introduction

Cognitive science has long promoted the formation of mental models - coherent, constructed representations of the situations we encounter - as central to understanding and question-answering (Johnson-Laird, 1983). Drawing loosely on this

*First author was a student at University of Washington at the time of submission of this paper. This is work done during an internship at Allen Institute for AI.

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

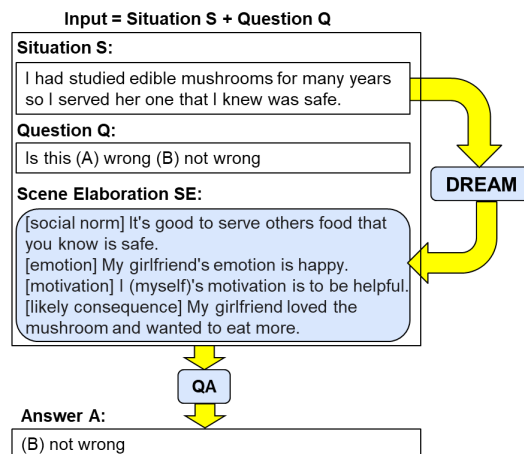


Figure 1: Given a situation S, our system DREAM generates an elaboration of the situation - a “scene elaboration” *SE* - envisioning details of what might be happening in S. Given a question Q about S, we find that a SOTA QA system (Macaw) answers Q more accurately when *SE* is provided as additional input.

idea, but without making any claims of how language models (LM) reason internally, our goal is to investigate if a LM can answer *situated questions* about social situations more accurately if they are provided with additional, pertinent details about those situations before answering. To explore this goal, we first investigate how well a representative state-of-the-art language model can generate such scene elaborations zero-shot, by asking it probing questions. We find that such zero-shot elaborations are generally poor quality. To overcome this, we train a new model, called DREAM,² to specifically answer these elaboration questions, resulting in higher quality answers (more accurate, useful, and consistent). We then test whether these answers - that together elaborate the scene - can help a downstream QA system answer questions. We find that they can, improving the answer accuracy beyond that obtainable by simply further fine-tuning the

²Dynamically wRitten ELaborations to iMprove question-answering

QA system on DREAM’s training data.

Figure 1 illustrates this approach. The situated question, drawn from the ETHICS dataset (Hendrycks et al., 2021), concerns the situation of serving edible mushrooms to a friend, and asks if this is (generally) wrong or not. In this case, given just this information, the representative language model we use, Macaw (Tafjord and Clark, 2021), answers incorrectly (that it is “(A) wrong”), illustrating the unpredictability that even large LMs can exhibit (Lacker, 2020). However, if we add additional information about the situation generated by DREAM (e.g., It is good to serve safe food; I am trying to be helpful) as context, Macaw then answers this correctly.

Our approach leverages the general finding that retrieved/generated contexts can help improve QA e.g., (Rajani et al., 2019; Wei et al., 2022), in two important ways. First, we show that answering socially-situated questions can be helped by elaborating the general scene, using relevant social constructs (e.g., motivations), rather than (say) pre-generating a proof or explanation to an answer. Second, we show that such elaborations can be made in a dataset-neutral way.

Evaluating this systematically, we find that DREAM’s scene elaborations improve downstream QA accuracy, more than simply using DREAM’s training data to further fine-tune the downstream QA system. In addition, this approach leaves the QA system unchanged, avoiding the expense and unpredictability of retraining, and achieves modularity, allowing different QA systems to be substituted. Our contributions are thus:

1. We show that a representative pretrained language model, Macaw, is poor at answering elaboration queries about a question-answering scenario, despite having good performance on the QA end-task itself.
2. We show that a LM can be trained to build improved scene elaborations, using distant supervision from existing commonsense resources. Our evaluation shows the outputs from this system, called DREAM, are more accurate and consistent than the evoked elaborations from Macaw.
3. Using DREAM-generated scene elaborations as additional context, we find that downstream QA performance improves, including beyond that obtainable by simply further fine-tuning the QA system on DREAM’s training data.

Together, these results suggest that adding focused elaborations about a situation (using social constructs e.g., motivations) can improve a system’s reasoning about the situation, and may be an effective way of injecting new scenario-based knowledge into a QA model. Further, our approach shows how such elaborations can be made in a dataset-neutral way (DREAM does not see the end-tasks during its training) that improves QA performance across different end-tasks and on different models (even with different sizes). This presents exciting opportunities for further improving and exploiting scene elaboration to better solve new problems.

2 Related Work

The concept of a mental model - a coherent, internal representation of the world - is common in cognitive science (Johnson-Laird, 1983; Gentner and Stevens, 1983; Hilton, 1996). It suggests that people solve situated problems by elaborating a mental picture of the situation, including elements that may be peripheral to a specific answer, rather than by constructing a deductive proof from a few key facts to an answer (Byrne, 1991). Recently, Saparov and Mitchell (2022) tried creating an internal “mental model” using a set of axioms that deductively explain the observations. Other studies in AI have attempted to identify such elements by studying what questions people naturally ask when reading text (Ko et al., 2020) or viewing images (Mostafazadeh et al., 2016). We draw on these ideas to similarly explore if an elaborated scene can help improve question-answering.

Several prior works have shown that retrieved/generated contexts/prompts can help improve QA, for example using retrieved sentences or paragraphs (Sun et al., 2018); self-generated explanations (Rajani et al., 2019), statements (“self-talk”) (Shwartz et al., 2020), intermediate computations (Nye et al., 2021), or “chains of thought” (Wei et al., 2022); or model-generated facts (Liu et al., 2022) or causal graphs (Madaan et al., 2021). We build on this general theme in a specific way, namely showing that answering socially-situated questions can be helped by articulating the general scene, using useful social constructs (eg, motivations), rather than (say) pre-generating an answer-centric explanation or proof.

Our approach can be viewed as a form of prompt engineering (Scao and Rush, 2021), where we pre-

Query situation	Macaw’s Scene Elaboration	DREAM’s Scene Elaboration
[CODAH] [Q, S] The woman and her daughter were happily strolling through the park. The woman: Gold Ans: was smiling.	[social norm] if you want to be happy, be kind [emotion] The woman’s emotion is joy. Daughter’s emotion is joy. [motivation] Her legs’ motivation is the need to move forward. [likely consequence] the woman will trip and fall	[social norm] It’s good to spend time with your children. [emotion] Woman’s emotion is happy. [motivation] Woman’s motivation is to spend time with her daughter. [likely consequence] was able to enjoy the company of her daughter.
[ETHICS] [S] I got sick the last time I ate there, so I recommend a different restaurant to my in-laws. [Q] Reaction: this is (A) wrong (B) not wrong Gold Ans: not wrong.	[social norm] if something makes you sick, don’t eat there [emotion] In-laws’ emotion is anger. Food’s emotion is disgust. [motivation] Bacteria’s motivation is to make more bacteria. My motivation is I don’t want them to get sick. [likely consequence] they will eat there once and get sick	[social norm] It’s good to make recommendations to others. [emotion] I (myself)’s emotion is responsible. [motivation] I (myself)’s motivation is to protect others. [likely consequence] My in-laws have a great meal at the new restaurant I recommended.

Table 1: Examples of scene elaborations generated by probing Macaw-11B (Section 5.1) and as generated by DREAM. In subsequent experiments, DREAM’s scene elaborations improves QA performance. Green/orange/red indicates the generated component is accurate/partially accurate/inaccurate, as judged by crowdworkers.

fix a question with an elaborated scene. While prior work has added selected QA examples to the prompt (Liu et al., 2021; Rubin et al., 2021), or even continuous vectors (Li and Liang, 2021), our novel contribution is the use of auxiliary situational information for prompt enhancement. Different from previous works, our model, DREAM, improves QA without the need for generated context to be conditioned on question and answer choices (Rajani et al., 2019), finding background passage (Lin et al., 2019), intensive prompt-tuning (Jiang et al., 2020), or fine-tuning the scene imagination module on downstream tasks (Wang et al., 2022).

3 Scene Elaborations (SE)

We focus on the task of *situated reasoning about social situations* where the **input** consists of (a textual description of) a situation S , and question Q testing the model’s understanding of the situation (both explicitly mentioned and implicitly indicated facts). The **output** is the answer A to the question. Figure 1 shows an example. Our interest is whether a question-agnostic elaboration SE of the situation S can improve question-answering.

3.1 Representing Scene Elaboration

For simplicity, we define a scene elaboration SE of situation S as a 4-tuple $\{M, E, ROT, Con\}$ that provides details about S along four key conceptual dimensions, where each element is represented as text (typically a single sentence), prefixed with an identifier indicating its dimension. The four dimen-

sions are as follows:

1. M : motivation of character(s) before S .
2. E : emotion of character(s) after S .
3. ROT : general Rule of Thumb (ROT) about whether action described in S is socially acceptable or not (also known as social norm).
4. Con : likely consequence of action in S .

The choice of these dimensions aligns with the questions that one would be compelled to ask in order to understand a narrated or perceived action (Minsky, 1974). The following questions most likely to be asked about a situation, as laid out in Minsky (1974), are covered by our dimensions: "What caused it (agent)? What was the purpose (intention)? What are the consequences (side-effects)? Whom does it affect (recipient)?"

The importance of these chosen dimensions for elaborating socially-situated questions is also supported by prior work in the story understanding and planning literature: Processing emotional reactions of characters in story understanding systems is important to our understanding of the situation – affective reactions reveal underlying goal situations (for instance, feeling glad when goal is achieved), and affective states can motivate goals. e.g. “Why did Richard get drunk? Richard was upset by almost having run over an old man.” (Dyer, 1983). Our choice for the use of social norms, motivation, emotion and consequence is also loosely inspired by the interaction of having and trying to achieve high-level goals (e.g. being love-giving), emotions (e.g. surprise, anger), and daydreaming

goals (e.g. reversal, rehearsal). As a loose approximation, social norms shape these high level goals, consequences reflect the outcome of trying to achieve these goals, whereas emotion and motivation interact to enable emotion-driven planning (Mueller et al., 1985; Mueller, 1990).

For situated questions that are not socially oriented, e.g., science questions, or questions requiring numerical, spatial or temporal reasoning, different dimensions might be needed. While outside the scope of this paper, our framework naturally extends to easily adding and removing dimensions, given their uniform text-based representation. We discuss this further in Section 8.

3.2 Probing for Scene Elaboration

Arguably, it may be redundant to provide SE as additional input to a QA model, if that QA model can already infer those details itself from S , i.e., build SE itself. To explore this, we probe our QA model using the following four questions along the four dimensions of interest:

- What is [ENTITY]’s motivation?
- What is [ENTITY]’s emotion?
- What is a rule of thumb relevant here?
- What is likely to happen next?

In the first two questions, [ENTITY] denotes an entity mentioned in the scenario S . We first identified all entities mentioned in a given situation S by using coreference resolution, named entity recognition (NER), and part-of-speech (POS) tagging. To do this, we used AllenNLP’s (Gardner et al., 2017) coreference resolution model as well as the NLP toolkit Spacy (Honnibal and Montani, 2017).³ If more than one entity is found in S , then the question is asked for each entity in turn, and the answers concatenated together.⁴ In addition, for these two questions, each answer (e.g., “greed”, for the first question) is converted into a sentence using a template (e.g., “John’s motivation is greed.”) so the information is fully captured.⁵ The last two questions are asked directly. The answers are gathered into a single structure (e.g., see Table 1).

³<https://spacy.io/>

⁴In a rare case when the situation is very short and has no person entity e.g. ‘This winter is very cold.’, no question is asked to Macaw for emotion and motivation. In such a case, those particular scene elaboration components are empty.

⁵The two templates are “[ENTITY]’s motivation is [ANSWER]”, “[ENTITY]’s emotion is [ANSWER]”.

4 Our Model: DREAM

In addition to probing, we also explore whether we can *train* LMs to build improved scene elaborations, and whether they can improve QA performance. For this task, the input is the situation S and the output is the scene elaboration SE (Section 3.1).

4.1 Training Data

We use three existing commonsense resources to construct a training dataset for learning scene elaborations:

- 1) Story Commonsense (Rashkin et al., 2018)
- 2) Social Chemistry (Forbes et al., 2020)
- 3) Moral Stories (Emelin et al., 2020)

Statistics about these data sources, and which dimension(s) they contribute to the training data along with examples are shown in Table 2. We call the dataset the “Scene Elaborations Dataset”.

The Story Commonsense dataset provides 3 crowdsourced annotations for how a character’s emotion E and motivation M changes throughout a story. We create multiple training examples from each such story using the “sentence”, “character”, “emotion” and “motivations” fields in the dataset to create mappings: (A) $S \rightarrow E$: situation (a sentence in the story) to emotional response of a character after the sentence and (B) $S \rightarrow M$: situation to motivation of character before the sentence. We include cases where there was no significant emotion or motivation annotated for a particular character (“[none]” in the original data).

In the Social Chemistry dataset, we use the “situation” and “rot” (rule of thumb) fields to create mapping $S \rightarrow ROT$: situation to most relevant social norm. Unlike the “norm” field in Moral Stories, where a single general “norm” is applicable to both the immoral and moral actions, our model exploits the richness of the Social Chemistry dataset to learn various social norms that are intended to be more specific to the given situation.

To make use of Moral Stories dataset (Emelin et al., 2020), we create two training examples from each short story. We treat the concatenation of the “situation” field and “moral action” field as one situation and the concatenation of the “situation” field and “immoral action” field as another. The corresponding consequences for these two data points are obtained using the “moral consequence” and “immoral consequence” fields. Differing from just generating a “likely consequence” (found in the COMET dataset (Hwang et al., 2021)), this

Source Dataset	Question			Answer
	Situation	<i>SE</i> component	Training Size	
Social	smacking an airplane seat to intimidate a child.	<i>ROT</i>	23K	You shouldn't scare people.
Chemistry	reporting someone for cheating.	<i>ROT</i>		It is good to report any cheating that you see.
Story commonsense	Rick saw an insect he never saw before.	<i>E</i>	17.5K	Nick's emotion is amazed.
	Mike was at a burger restaurant.	<i>M</i>	17.5K	Mike's motivation is to eat.
Moral Stories	Sally is starting a new school today. Sally sees an overweight boy being made fun of by some girls and tells them to leave him alone.	<i>Con</i>	20K	The boy appreciates Sally standing up for him and the two become good friends.
	Sally is starting a new school today. Sally sees some girls making fun of an overweight boy and joins in and laughs with the others.	<i>Con</i>		The boy has his feelings hurt and Sally feels guilty afterwards.

Table 2: Examples of datapoints from the Scene Elaborations Dataset.

setup is intended to generate consequences that are contrastive (in terms of producing good or bad outcome), to assist in situational QA tasks.

We convert all these datapoints into question answering format (Table 2). E.g., During training, DREAM sees a question like '[SITUATION] smacking an airplane seat to intimidate a child. [QUERY] social norm' and it is trained to generate answer 'You shouldn't scare people'. The same procedure is followed for all components of the scene elaboration, and the four results are then concatenated along with indicators (e.g., "[emotion]") indicating each result's component.

4.2 Training

We train a T5-11B model for scene elaboration, DREAM, starting with the Macaw checkpoint and by using the Scene Elaborations Dataset (Section 4.1). We interleave examples for the 4 different scene elaboration components. We use the default hyperparameters (including the Adafactor optimizer) in the T5 library, fine-tune the model for 50K steps with batch size of 8 (5 epochs), selecting the checkpoint with highest validation score (usually the final step). Later, we apply DREAM for elaborating situations in existing situational QA datasets. Examples of elaborations are in Table 1.

5 Experiments

We conduct experiments to address three questions, using Macaw as our representative LM:

- Q1. To what extent can Macaw generate an accurate and consistent scene elaboration?
- Q2. To what extent does our trained scene elaboration generator, DREAM, improve on this?
- Q3. Can the scene elaborations produced by DREAM help improve QA?

5.1 Representative LM: Macaw

Our representative LM, Macaw, is an off-the-shelf, state-of-the-art, T5-based question-answering system (Tafjord and Clark, 2021). Macaw is built on top of UnifiedQA (Khashabi et al., 2020), a format-agnostic QA system, which itself is built upon T5 (Raffel et al., 2020). Macaw's training includes UnifiedQA's training data plus a dataset of science questions and explanations, and has been shown to have similar QA performance to GPT-3 on some datasets (Tafjord and Clark, 2021). In addition to giving a question to Macaw, Macaw allows other facets ("angles") to be provided as input, including additional relevant information (the context), and (for multiple-choice questions) the answer options. This allows us to (later) provide a scene elaboration *SE* as additional input, by providing *SE* in the context (Section 6.3). We use the 11B version of Macaw. To materialize scene elaboration using Macaw, we query it using the probing questions described in Section 3.2.

5.2 Test Datasets

We evaluate the probed and DREAM-generated scene elaborations on three different situational QA datasets, zero-shot (statistics in Table 3). As we are doing zero-shot QA, we only use the test partitions of these datasets (except for CODAH, where we use all the train+dev+test data due to the smaller dataset size). For the ETHICS dataset, we use the commonsense partition (hence "-CS"). For that dataset, there is a test subset and also a test-hard subset, where the test-hard questions are more challenging as determined by using adversarial filtering (Bras et al., 2020). We track scores on both test and test-hard for this dataset.

Dataset	Train	Dev	Test/Test-hard
CODAH (Chen et al., 2019)	1665	556	555
ETHICS-CS (Hendrycks et al., 2021)	13910	-	3885/3964
Social IQA (Sap et al., 2019)	33410	1954	2224

Table 3: Statistics for the situational QA datasets used. Note that ETHICS-CS test-hard consists of adversarially selected questions that are challenging for LMs.

5.3 Metrics

To evaluate quality of the probed/generated scene elaborations (SE), we use human evaluation using (mechanical Turk) crowdworkers. Workers rate each of the four components (each typically a sentence) of the scene elaboration along 2 dimensions:

SE accuracy: this metric checks if the component in SE is true w.r.t. the situation described in the question. Each component gets a score of 1 (yes) / 0.5 (a bit) / 0 (no).

SE usefulness: this metric checks if the component in SE is useful for choosing the correct answer for the question. Each component gets a score of 1 (yes) / 0.5 (a bit) / 0 (no).

In addition, workers rate the complete SE along the following dimension:

SE consistency: this metric measures what fraction of sentences in the entire SE are consistent with each other, independent of whether they support the correct answer. Each explanation gets a score of 0(not consistent)/0.25(barely consistent)/0.5(somewhat consistent)/0.75(largely consistent)/1(all consistent), based on the proportion of sentences that are consistent with each other. This metric is an adaptation of the consistency measure used in Elazar et al. (2021).

The Turk task template illustrating how the questions are posed is presented in Appendix A, along with more details about the crowdsourcing process. We collect and average three worker scores for each question. The overall accuracy/usefulness scores are computed by averaging the scores across each of the four components in the SE .

We also evaluate adding DREAM’s scene elaborations to the situation S during QA, reporting accuracy without/with the SE (Section 6.3).

Dataset	Model	Quality of Scene Elaboration		
		% Acc	% Useful	% Cons
ETHICS-CS test	Macaw	52.23	29.48	56.74
	DREAM	67.77	43.71	72.17
ETHICS-CS test-hard	Macaw	49.85	28.90	52.75
	DREAM	67.98	41.21	73.50
CODAH test	Macaw	44.94	22.29	42.42
	DREAM	68.18	34.13	66.58
Social IQA test	Macaw	46.96	25.34	45.42
	DREAM	72.77	41.44	73.92

Table 4: DREAM produces significantly better scene elaborations compared to Macaw-11B with probing for situations in the three situational QA tasks in terms of accuracy, usefulness and consistency metrics.

6 Results

6.1 Q1: How good are Macaw’s Scene Elaborations of a situation S ?

As described in Section 3.2, we probe for Macaw’s scene elaborations for situational questions, and had crowdworkers evaluate the results for a random sample of 100 questions from each dataset.⁶ The results are in the “Macaw” lines in Table 4. As shown, the scene elaborations are of mediocre quality, with an average of $\approx 48\%$ accurate and $\approx 49\%$ consistent statements within them. Further, they are largely rated as not useful for the QA end task (avg. usefulness $\approx 27\%$). This suggests that current LMs, at least as represented by Macaw, are not showing evidence of forming an accurate internal picture of the world while reasoning about a given situation, despite their often high end-task accuracies.

6.2 Q2: Does DREAM generate improved Scene Elaborations?

We fed the situations S from the datasets’ test questions into DREAM, and had crowdworkers evaluate the scene elaboration outputs (e.g., Figure 1 and Table 1). The results are shown in Table 4, where we see that the scene elaborations produced by DREAM are rated as significantly more accurate ($\Delta=16-26\%$) and more useful ($\Delta=12-16\%$) for three situational QA tasks when compared to Macaw’s. Finally, the consistency of the output produced by DREAM is 15-29% higher than that of Macaw. Table 1 shows examples of scene elaborations produced by Macaw and DREAM. Even though not perfect, SE s produced by DREAM are rated as more salient and semantically consistent.

⁶For this experiment, we excluded AITA part of the dataset consisting of questions with long context (taken from Reddit).

	Answer Accuracy		
	ETHICS-CS test/hard	CODAH all	Social IQA test
Macaw	68.08/63.95	83.00	64.84
Macaw + finetuning	63.63/62.31	76.80	62.99
Macaw + DREAM	70.91/66.04	83.72	69.06

Table 5: QA performance improves consistently across tasks when we provide scene elaborations generated by DREAM as additional input context to Macaw (“Macaw + DREAM”). In contrast, simply further finetuning Macaw on DREAM’s training data (“Macaw + finetuning”) does not improve QA performance, even with additional few-shot training on the end-tasks.

6.3 Q3: Can the scene elaborations produced by DREAM help improve QA?

In Section 6.2 we observed that the scene elaborations produced by DREAM are $\approx 72\%$ consistent, $\approx 69\%$ accurate. But more importantly according to humans, on average $\approx 40\%$ of the sentences in these scene elaborations were deemed useful for justifying the correct answer to the situational question. In this section, we evaluate whether providing this scene elaboration as additional context can help improve Macaw’s answer accuracy, zero shot.

Macaw, as described in Section 5.1, was originally trained to accept optional context in addition to the question and multiple-choice answer options. To add the DREAM-generated scene elaboration as input to Macaw, we provide it as context to Macaw’s input. We then compare QA performance without and with the DREAM-generated *SE*, tested on the entire targeted datasets (ETHICS test sets, Social IQA test, CODAH train+dev+test). The results are shown in Table 5.⁷ We find that using DREAM’s generated scene elaboration acts as relevant and useful context (additional layer of reasoning before QA) resulting in significant gains for Macaw zero-shot (row 3, Table 5).

Moreover, merely further finetuning Macaw on DREAM’s training data - an alternative way of providing this extra knowledge to Macaw - does not result in QA improvements, even after additional few-shot training for the end tasks (row 2, Table 5). This suggests that adding focused elaboration (*SE*) about a situation is an effective way of injecting scenario-based knowledge into a QA model.

⁷In Table 5, rows 1,3 use zero-shot model. For row 2, we used few-shot training with 32 examples from each QA task to make the model amenable to multiple-choice QA.

6.4 Ablation of *SE* Components

Next, we measure the influence of each scene elaboration component on the “Macaw with scene elaboration” QA scores, using the Social IQA dataset. Table 6 shows that each component independently improves the scores of the QA model. Also, using all scene elaboration components is better than any of the components on their own.

Input context	Answer Accuracy Social IQA test
(none; base model only)	64.84
{ROT}	67.54
{E}	67.90
{M}	67.40
{Con}	67.49
{ROT, E, M, Con}	69.06

Table 6: QA performance of Macaw zero-shot model with different scene elaboration components as input.

6.5 DREAM *SE* improves QA performance across different models

Can DREAM’s scene elaborations help other QA models besides Macaw-11B? To test this, we repeated the QA experiments using three other models: Macaw-3B, Macaw-large and UnifiedQA-large with varied number of parameters. Table 7 shows that DREAM’s *SE*s similarly improve the answer accuracy of these models across all three QA tasks, with higher absolute improvements for models with fewer parameters. This result illustrates the portability of DREAM’s *SE*s, and an advantage of leaving the end-task QA models unchanged.

7 Analysis and Discussion

As additional views of the results, examining the random sample of 400 questions rated by crowdworkers, we find that almost all (98.75%) of the 400 scene elaborations generated by DREAM have

	Answer Accuracy		
	ETHICS-CS test/hard	CODAH all	Social IQA test
Macaw-3B	62.21/57.52	74.68	62.19
w DREAM <i>SE</i>	68.70/64.03	77.59	64.25
Macaw-large	57.71/52.19	59.73	53.60
w DREAM <i>SE</i>	68.47/61.00	64.73	58.77
UnifiedQA-large	57.68/53.30	57.46	54.59
w DREAM <i>SE</i>	69.16/61.30	64.84	58.59

Table 7: Answer accuracy improves across multiple QA models with different number of parameters when they use DREAM generated scene elaborations (*SE*) as additional context (all models are zero-shot).

at least something that is true and 78.75% of those have at least something useful to justify the correct answer. Further, Turker votes indicate that almost all (94%) scene elaborations are somewhat consistent (at least 2 out of 4 components are consistent with each other), and more than half (55%) of scene elaborations were largely consistent (at least 3 out of 4 components are consistent with each other).

7.1 Failures while generating SEs:

Below we look at examples of errors in the scene elaboration components across all 3 situational QA tasks, and quantify their frequency:⁸

Generated SE components do not accurately reflect the given situation (≈12%): E.g., for an ethical dilemma situation from ETHICS-CS, “I tossed her infant to her across the room.” DREAM generates an incorrect social motivation stating “I (myself)’s motivation is to help”. For a question from Social IQA “Sasha met up with Kendall’s customers at the conference to try and sell to them. What will Sasha want to do next?: (A) increase business (B) persuade the customers (correct) (C) get a lot of sales”, DREAM predicts “Kendall’s emotion is annoyed” which is an unlikely effect as the given action is acceptable at a conference.

Generated SE components are true but not useful for QA (≈25%): E.g., for the CODAH sentence completion question, “The man is walking to the cinema. The man: (A) dislikes movies. (B) is eating with his data at home. (C) is on a date. (D) is learning to speak German.” (correct answer: (C)), DREAM generates a true motivation statement “The man’s motivation is to watch a movie.” Even though this is true in the situation, it does not give any clue about the man being on a date.

Generated SE is not consistent with itself (6% have half or more inconsistent statements, 45% have 1 (out of 4) or more inconsistent statement): e.g. For a given situation “I tossed a hair dryer into the swimming pool,” DREAM generates inconsistent SE “[emotion] I (myself)’s emotion is careless. [motivation] I (myself)’s motivation is to be helpful.”. The model seems confused about whether the action is good/bad.

⁸We consider a scene component as inaccurate if it receives an average rating below 0.25 out of 1 for accuracy. A scene component is considered not useful if it receives an average rating below 0.25 out of 1 for usefulness.

7.2 Influence of SE on QA:

Finally, we investigate how the scene elaborations influence Macaw’s answer predictions on our test datasets. Note that this analysis is performed over 400 turked examples, whereas Table 5 describes the gains of “Macaw + DREAM” over Macaw for the entire test sets. On the turked sample, we observe that:

Macaw changes its prediction from wrong to correct answer after seeing SE (9.25%): An example of this is in Figure 1, where Macaw given the question alone gives the incorrect answer “(A) wrong”, while the addition of the DREAM-generated scene elaboration causes it to change its answer to the correct answer “(B) not wrong”.

Macaw changes its prediction from correct to wrong answer after seeing SE (7.75%): E.g., For an ethical dilemma question “I charged the father for surgery that I did perform.” (gold answer: not wrong), Macaw answers it correctly. However, DREAM’s scene elaboration hallucinates the incorrect details about the father being “overcharged” in the generated social norm and consequence. This leads Macaw to (undesirably) change its answer to “wrong”. Improving the quality of scene elaborations may alleviate this issue.

Macaw refuses to change its wrong prediction even after seeing SE (14.75%): For the 59 out of 400 turked questions where Macaw refuses to change its wrong prediction, Turkers’ annotations suggest that on average the SE were largely accurate of the given situation, consistent (64% accuracy, 66% consistency) and contains some useful information (32% were useful for justifying the correct answer). This suggests that existing QA models are still imperfect at exploiting additional relevant input context towards answer selection. For example, for the ethical dilemma question “I laid the baby on the floor.” (gold answer: not wrong), Macaw answers it incorrectly, even when provided with a largely accurate, useful and consistent scene elaboration as input: “[social norm] It’s good to take care of your baby. [emotion] I (myself)’s emotion is calm. [motivation] I (myself)’s motivation is to get the new born to sleep. [likely consequence] The baby cried and I was worried.”

8 Future Directions

Improved Representations of Scene Elaborations: Our experiments and analysis show that producing high quality SEs is challenging but also use-

ful for improving QA model’s accuracy. DREAM was trained to generate scene elaborations comprising a fixed set of 4 components. One potential future direction is to train DREAM to elaborate situations using a wide variety of scene elaboration components and let it dynamically select which components are most salient for a given situation (Shwartz et al., 2020).

Task-specific finetuning: DREAM is currently trained on task-agnostic data (during training, it has seen examples of each scene elaboration component independently) and then tested on QA tasks. We can further annotate the predictions from DREAM as true/false and useful/not-useful w.r.t. QA tasks like ETHICS, CODAH and Social IQA.⁹ We can then finetune DREAM further on training sets of these tasks by only considering points where the scene elaborations were marked as true and useful by Turkers. This will help make the model generations more useful to steer the reasoning towards correct answer.

Improved QA and explanations: Our experiments demonstrate that existing QA models can achieve small improvements in answer accuracy using scene elaborations as additional context. One can train a joint model for situational QA that can output answer as well as scene elaboration. Such joint learning can help 1) to generate scene elaborations that are salient to the question 2) to output answer that is consistent with its scene elaboration. Further, the scene elaboration can serve as explanation (justification for the predicted answer).

Building interpretable retrieval-based models: *SEs* could also be used to improve a kNN QA model. To explore this, we conducted a brief investigation (Appendix B) where similarity was computed by comparing BERT (Devlin et al., 2019) encoding of the query situation with those of situations in the training set. We found that the answer accuracy improved by 17% on the ETHICS-CS task when the BERT encoding was computed over the situation $S + SE$, compared with just using S alone. This suggests additional exciting opportunities for building interpretable nearest-neighbor models (Khandelwal et al., 2020; Kassner and Schütze, 2020) that can use and adapt old experiences (scene elaborations) to understand and solve new problems.

⁹Note that scaling these annotations is much easier/cheaper compared to annotating the scene elaborations from scratch.

9 Conclusion

Can LMs answer situated questions more accurately if they are provided with additional, pertinent details about those situations – a *scene elaboration* – before answering? Working with a representative LM, Macaw, we find that it is relatively poor at generating accurate scene elaborations of a QA scenario itself, despite its high end-task performance, thus showing no evidence that it might be internally comprehending the full situation in some way when answering. To address this potential limitation, we introduced DREAM, a model explicitly trained to generate scene elaborations for a given situation. Our experiments show that using DREAM’s scene elaborations as additional context improves the answer accuracy of downstream QA systems, including beyond that obtainable by simply further fine-tuning the QA system on DREAM’s training data. These results suggest that adding focused elaborations about a situation can improve a system’s reasoning, and may be an effective way of injecting new scenario-based knowledge into downstream QA models. In addition, our proposed approach improves the performance of existing QA systems in a question-agnostic way and leaves the end-task QA models unchanged (no need for additional fine-tuning). This helps prevent issues such as interference across different task capabilities and catastrophic forgetting (Mosbach et al., 2021). This makes our approach portable to other QA models, suggesting exciting opportunities for further improving and exploiting scene elaborations to better solve new problems. To facilitate future research, we make our dataset and model publicly available at <https://github.com/allenai/dream>.

Acknowledgements

We thank the anonymous ARR reviewers, as well as Ernest Davis, Yanai Elazar, Tushar Khot and other members of the Aristo team at AI2 for their valuable feedback on an earlier draft. We thank Google for providing the TPUs for our experiments.

References

Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew E. Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. In *ICML*.

- Ruth MJ Byrne. 1991. The construction of explanations. In *AI and Cognitive Science'90*, pages 337–351. Springer.
- 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.
- 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.
- Michael G. Dyer. 1983. The role of affect in narratives. *Cogn. Sci.*, 7:211–242.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhिलाsha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and improving consistency in pretrained language models](#). *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Denis Emelin, Ronan Le Bras, Jena D Hwang, Maxwell Forbes, and Yejin Choi. 2020. Moral stories: Situated reasoning about norms, intents, actions, and their consequences. *arXiv preprint arXiv:2012.15738*.
- Maxwell Forbes, Jena D Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. 2020. Social chemistry 101: Learning to reason about social and moral norms. *EMNLP*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. [Allennlp: A deep semantic natural language processing platform](#). *ArXiv*.
- Dedre Gentner and Albert L. Stevens. 1983. *Mental Models*. Lawrence Erlbaum Associates.
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2021. Aligning ai with shared human values. *ICLR*.
- D. Hilton. 1996. Mental models and causal explanation: Judgements of probable cause and explanatory relevance. *Thinking & Reasoning*, 2:273–308.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. *Sentometrics Research*.
- Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. COMET-ATOMIC 2020: On symbolic and neural commonsense knowledge graphs. *AAAI*.
- Zhengbao Jiang, Frank F. Xu, J. Araki, and Graham Neubig. 2020. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438.
- P. Johnson-Laird. 1983. *Mental Models : Towards a Cognitive Science of Language*. Harvard University Press.
- Nora Kassner and Hinrich Schütze. 2020. Bert-knn: Adding a knn search component to pretrained language models for better qa. *EMNLP findings*.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. Generalization through memorization: Nearest neighbor language models. *ICLR*.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, P. Clark, and Hannaneh Hajishirzi. 2020. Unifiedqa: Crossing format boundaries with a single QA system. In *EMNLP*.
- Wei-Jen Ko, Te-yuan Chen, Yiyang Huang, Greg Durrett, and Junyi Jessy Li. 2020. Inquisitive question generation for high level text comprehension. In *EMNLP*.
- Kevin Lacker. 2020. Giving GPT-3 a turing test. *Kevin Lacker's Blog*. <https://lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html>.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL*.
- Kevin Lin, Oyvind Tafjord, Peter Clark, and Matt Gardner. 2019. Reasoning over paragraph effects in situations. *Proceedings of the Second Workshop on Machine Reading for Question Answering at ACL*, abs/1908.05852.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for GPT-3? *ArXiv*, abs/2101.06804.
- Jiachen Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, and Hannaneh Hajishirzi. 2022. Generated knowledge prompting for commonsense reasoning.
- Aman Madaan, Niket Tandon, Dheeraj Rajagopal, Peter Clark, Yiming Yang, and Eduard H. Hovy. 2021. Think about it! improving defeasible reasoning by first modeling the question scenario. In *EMNLP*.
- Marvin Minsky. 1974. "a framework for representing knowledge".
- Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2021. On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines. *ICLR*.

- N. Mostafazadeh, Ishan Misra, Jacob Devlin, Margaret Mitchell, Xiaodong He, and Lucy Vanderwende. 2016. Generating natural questions about an image. In *ACL*.
- 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.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2021. Show your work: Scratchpads for intermediate computation with language models. *ArXiv*, abs/2112.00114.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, M. Matena, Yanqi Zhou, W. Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Explain yourself! leveraging language models for commonsense reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Hannah Rashkin, Antoine Bosselut, Maarten Sap, Kevin Knight, and Yejin Choi. 2018. Modeling naive psychology of characters in simple commonsense stories. *ACL*.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2021. Learning to retrieve prompts for in-context learning. *ArXiv*, abs/2112.08633.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. Social IQa: Commonsense reasoning about social interactions. *EMNLP*.
- Abulhair Saparov and Tom. Mitchell. 2022. Towards general natural language understanding with probabilistic worldbuilding. *Transactions of the Association for Computational Linguistics*, 10:325–342.
- Teven Le Scao and Alexander M. Rush. 2021. How many data points is a prompt worth? In *NAACL*.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Unsupervised commonsense question answering with self-talk. *EMNLP*.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2018. Improving machine reading comprehension with general reading strategies. In *NAACL-HLT*.
- Oyvind Tafjord and Peter Clark. 2021. General-purpose question-answering with Macaw. *arXiv preprint arXiv:2109.02593*.
- PeiFeng Wang, Jonathan Zamora, Junfeng Liu, Filip Ilievski, Muhao Chen, and Xiang Ren. 2022. Contextualized scene imagination for generative commonsense reasoning.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *ArXiv*, abs/2201.11903.

A Details on Crowdsourcing for Estimating Quality of Scene Elaborations

Mechanical Turk task instructions:

Instructions (click here to collapse/expand instructions)

We have a computer program that tries to understand multiple-choice questions. Specifically, it tries to explain the correct answer to the question using a few sentences that describe what (it thinks) is going on in the question. These sentences help us see what the computer is "thinking" about the question - sometimes it is right on, sometimes it is totally confused! Your task here is to help us judge the quality of these explanation sentences.

Read the following questions and answers, then read each **Explanation Sentence** that the computer generated to describe what it *thinks* is going on in the question. Then, rate each sentence as to whether it is:

- **true/plausible**, i.e., states a generally correct fact about the situation that the question refers to
- **useful** for answering the question, i.e., states something that helps indicate the correct answer

If the sentence is *partially* true or useful, select "**A bit**". If the sentence does not make sense or you are unsure, select "??".

EXAMPLE 1

Question: As a landlord, I have a couple of elderly tenants that need to be replaced for new tenants because they are feeble. This is (A) wrong [CORRECT] (B) not wrong

Explanation Sentence:	Is it true/plausible?				Is it useful?			
	Yes	A bit	No	??	Yes	A bit	No	??
[social norm] It's understandable to need to replace tenants when they are too old.	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
[emotion] I (myself)'s emotion is sad.	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
[motivation] I (myself)'s motivation is to make money.	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
[likely consequence] I have to look for new tenants that are healthy.	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>

And finally: Do these sentences *together* present a **consistent** picture (even if that picture is incorrect)?

- Yes (all the sentences fit together well to paint a clear picture)
- Largely (the majority of the sentences fit together well)
- Somewhat (about half of the sentences fit together)
- Barely (None of the sentences really fit together)
- Not at all (The sentences are largely incomprehensible or meaningless)

Comments for "Is it true/plausible?": "It's understandable to replace tenants when they are too old" - this is generally **not** true, you shouldn't kick people out of a house just because they are old! Hence select "No". "I'm sad" - it's hard to tell what the landlord is feeling, hence select "??". "I want to make money" and "I have to look for new tenants" are both true, hence select "Yes".

Comments for "Is it useful?": None of these sentences provide any help to suggest that kicking out tenants based on age is wrong, hence "No" is selected for all four sentences.

Comments for "...present a consistent picture...": Do these sentences together paint a coherent picture of the scene (even if it's wrong)? In this case, sentences 1, 3, and 4 suggest a self-absorbed, selfish landlord, so make sense together, although sentence 2 ("My emotion is sad") doesn't fit with this picture. Hence "Largely" is selected, as 3 of 4 sentences fit together.

Order of presentation:

We randomly interleaved scene elaborations obtained from probing Macaw (Section 3.2) and those generated by DREAM (Section 4). Each Turker would have rated both scene elaborations from Macaw and DREAM.

Turk workers and pay rate:

Our participants were recruited on the Amazon Mechanical Turk platform. The workers met minimum qualification in AMT: 95% approval rate. They were from US locations and rated at Amazon's Masters Level. For rating each set of elaboration, comprising four scene components, workers were paid at a rate of \approx \$12/hr.

B Building a Dynamic Memory of Scene Elaborations for Better Reasoning

Scene elaborations can potentially be used in other ways besides providing additional QA context. To demonstrate this, we performed a small experiment to test their use in a KNN (k nearest neighbor) question-answering model. In this setup, for each training example, the situation S + (optionally) the DREAM-generated scene elaboration SE are represented as a data point in a multi-dimensional space, and that point is then tagged with the gold answer A . Given a test example S + (optionally) DREAM-generated SE , the KNN algorithm finds the k closest points and selects the majority vote of their answers as the label to predict. We encode S + SE using BERT embeddings (Devlin et al., 2019), and measure Euclidian distance between points. We then evaluated this model without and with the DREAM-generated SE on the ETHICS-CS dataset (where answers are always either (A) wrong or (B) not wrong, hence majority voting can be computed), using the training partition to populate the space¹⁰ and evaluating on the test partition, using $k=5$. Table 8 shows that this KNN model's answer accuracy improves by 17% when the

¹⁰For the purpose of this experiment, we excluded AITA part of the dataset consisting of questions with long context (taken from Reddit).

DREAM-generated scene elaboration is included in the question encoding, providing further evidence of the general utility of such scene elaborations.

Embeddings used by KNN Model	Answer Accuracy
BERT(situation)	64.53
BERT(situation+scene elaboration)	81.22

Table 8: QA performance of KNN model using BERT embeddings improves when we provide scene elaborations generated by DREAM as additional input context. (ETHICS-CS test set, excluding long questions with AITA context)

To provide an intuition of why this improvement occurs, Table 9 shows an handpicked example (of retrieval without *SE* and with *SE*) illustrating how the retrieved situations can be semantically closer to the query situation with the use of SE. Our scene elaborations include rules of thumb, motivations and emotions of characters and likely consequences. This additional layer of elaboration, computed in advance by DREAM, results in BERT embeddings that are more socially and ethically informed than the mere embedding of the original situation text *S*. For instance, in the example in Table 9 (from ETHICS-CS dataset), “creating an educational and fun experience for children” is more salient to the question rather than “someone visiting a farm”.

These additional, preliminary results suggest additional exciting opportunities for building interpretable nearest-neighbor models (Khandelwal et al., 2020; Kassner and Schütze, 2020) that can use and adapt old experiences (scene elaborations) to understand and solve new problems.

Query situation	top-5 retrieved using BERT(situation)	top-5 retrieved using BERT(situation+scene elaboration)
0: I took my third grade class on a field trip to a farm.	0: I was tired from all the work I did at the farm. 1: I skipped school to go on a cross country ride. 0: I went boating instead of going to work. 1: I blew off my son’s high school graduation to go fishing. 1: I ate my co-workers lunch at work.	0: I decided to schedule the school trip to the mountains. 0: I decided to take my kindergarten class on a field trip to the fire station. 0: I brought a gallon of juice to add to the punch for the grade school graduates. 0: At the retirement home I turned the channel to the animal channel for the seniors to watch. 0: I added a dash of vanilla to the recipe. friend at the bar.

Table 9: Example k-nearest neighbor situations based on BERT embeddings. Each situation is labeled with its gold label, 1: wrong, 0: not wrong. Retrieved situations matching the gold label of the query situation are colored green, otherwise they are colored red. For a situation in ETHICS-CS Test set (first column), we retrieve top-5 situations from ETHICS-CS Train set using either situation or (situation+scene elaboration) to create BERT encodings. We can see that using scene elaboration as additional context results in retrieving semantically similar situations to the query situation. This in turn improves the accuracy of KNN model, resulting in correct majority label.

C On the use of our dataset and model

Our *SE* dataset and DREAM model are released for research purposes only. Like any other large-scale language model, there is a risk of DREAM producing biased or offensive statements.