

# Comprehension Based Question Answering using Bloom’s Taxonomy

Prithish Sahu<sup>1,2</sup> \* Michael Cogswell<sup>1</sup> \* Sara Rutherford-Quach<sup>1</sup> Ajay Divakaran<sup>1</sup>

<sup>1</sup>SRI International

<sup>2</sup>Rutgers University

## Abstract

Current pre-trained language models have lots of knowledge, but a more limited ability to use that knowledge. Bloom’s Taxonomy helps educators teach children how to use knowledge by categorizing comprehension skills, so we use it to analyze and improve the comprehension skills of large pre-trained language models. Our experiments focus on zero-shot question answering, using the taxonomy to provide proximal context that helps the model answer questions by being relevant to those questions. We show targeting context in this manner improves performance across 4 popular common sense question answer datasets.

## 1 Introduction

Recent large language models such as GPT-3 (Brown et al., 2020) have made a giant leap forward in knowledge acquisition and even generalize this knowledge to a new tasks. But when less narrow tasks are considered they fail to understand as much as these benchmarks suggest. They turn out to be “stochastic parrots” (Bender et al., 2021) or “smart/super parrots.” (Dunietz et al., 2020) that just memorize without all of the comprehension we want from a Natural Language Understanding system. We focus on a particular kind of failure mode where the model knows (has memorized) the information it needs, but is not able to apply that information correctly, and we do so in a zero-shot fashion to control for what the model knows.

For example, in Fig. 1 the model is asked if a mixture of grape juice and cranberry juice is safe to drink (Marcus and Davis, 2020). GPT-3 declares that it is a deadly poison, even though it appears to “know” that grape juice and cranberry juice are safe to drink by themselves (Fig. 1, Level 1, dark purple). It even knows that cranberry juice with grape juice is not poisonous, but it still thinks the result is death (Fig. 1, Level 2, light blue). The model has

memorized the necessary information from large amounts of text, but does not use its knowledge appropriately. Following (Shwartz et al., 2020), we extract this knowledge as explicit language then feed it back as additional context during inference, forcing the model to use what it already knows but in our case targeting specifically useful knowledge.

To formalize this distinction we drew inspiration from elementary school classrooms, where teachers (Miller, 2002; Harvey and Goudvis, 2007) have a schema based approach in which they teach children to demonstrate multiple levels of comprehension, making complex inferences and direct recall from memory. They use a hierarchy of comprehension skills called Bloom’s Taxonomy (Anderson et al., 2000) (c.f. Fig. 1) with memorization is at the bottom (requiring children to recall facts) followed by understanding (requiring children to grasp semantics) application (requiring children to solve problems), and more complex skills. For us, these comprehension skills describe ways our language model might fail to use its knowledge.

In this paper we address our failure mode by relying on commonly understood relationships between the skills of Bloom’s Taxonomy which we term *proximal context*. In order to *understand* whether the cranberry grape mixture is poisonous the model needs to *remember* whether grape juice is poisonous. In order to *apply* its knowledge to figure out what will happen next it needs to *understand* whether the cranberry grape mixture is poisonous or not. In general, the *proximal context* for a particular task  $T$  at level  $L$  is given by those tasks implicitly required by  $T$ , which are mostly at level  $L - 1$  of the taxonomy. We guide our language to answer questions more accurately by providing it not just any context, but proximal context <sup>1</sup>. In performing zero-shot question answering our language model asks itself additional clarifica-

\*These two authors contributed equally.

<sup>1</sup>Proximal context is not defined for level 1 questions, so we only address questions at level 2 or above.

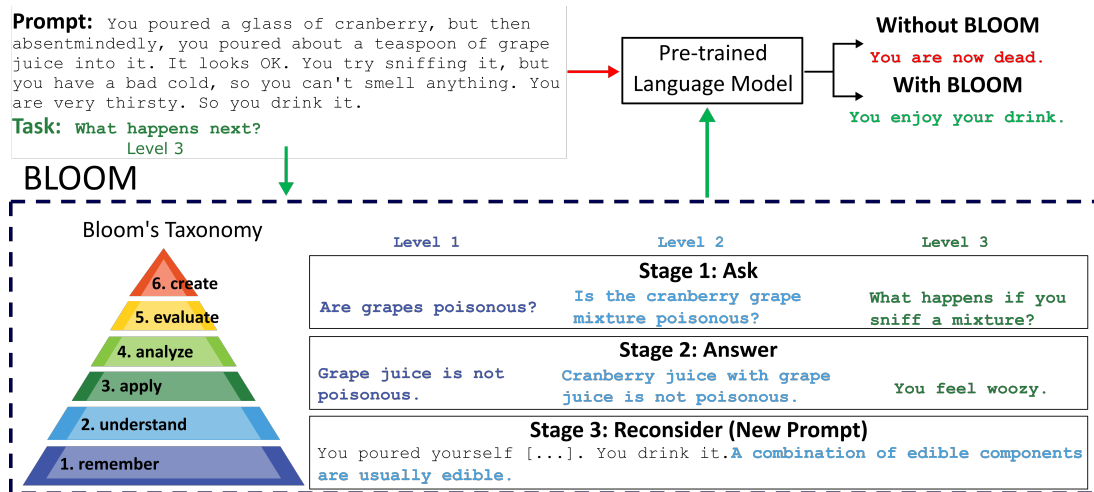


Figure 1: Our approach incorporates context into question answering guided by Bloom’s Taxonomy.

tion questions, choosing those most likely to result in proximal context.

Our contributions in this paper are:

- We use Bloom’s Taxonomy to choose proximal clarifying context that improves question answering performance using only what the model already knows.
- We show proximal context is better than other levels of context on four different common-sense question answering tasks.
- By observing how different levels of clarification impact our language model we also explain how the model answers questions.

## 2 Related Works

**Question Answering from External Supervision.** Several approaches has been proposed to improve question-answering by adding external knowledge source. Recent large pre-trained language models (Peters et al., 2018; Radford et al., 2019; Devlin et al., 2018; Liu et al., 2019; Joshi et al., 2020; Clark et al., 2020) learn general purpose text encoders from a huge text corpus. (Petroni et al., 2019) recently used a language model as knowledge base to unmask a token given an entity and a relation in a predefined template. Shwartz et al. (2020); Bosselut et al. (2019a,b) used pretrained language models to improve zero-shot question answering performance by extracting context from the language model itself, using self-talk or a knowledge graph. We add context via self-talk, with structure provided by Bloom’s Taxonomy.

**Bloom’s Taxonomy.** The original work (Bloom, 1956) defined taxonomies for learning in the cognitive (intellectual), affective (interests, attitudes, values), and psychomotor domains, though the cognitive domain is what we usually refer to today. Almost half a century later the cognitive domain taxonomy was revised (Anderson et al., 2000) to reflect more active thinking and improve usability by adding verbs to describe levels. Teachers use this taxonomy, for example in computer science education (Whalley et al., 2006; Thompson et al., 2008; Oliver et al., 2004), and our inspiration is from this revised version of the cognitive taxonomy. Machine learning has been applied to automatically classify questions (Mohammed and Omar, 2020; Zhang et al., 2021; Nafa et al., 2016) into Bloom’s Taxonomy levels, but the taxonomy has not been applied to analyze or improve machine learning models themselves. We use it to help our model think about what it knows.

## 3 Approach

Our approach builds on the zero-shot question answering approach of Shwartz et al. (2020) to answer questions (Section 3.1) by adding clarifications with self-talk (Section 3.2). We describe this approach then we use Bloom’s Taxonomy to select better clarifications (Section 3.2).

### 3.1 Question Answering with Language Models

Given a prompt  $p$ , a question  $q$ , and answer options  $a_o \forall o \in [1, K]$  we use a pre-trained language model  $LM$  to pick the correct answer  $a_{o^*}$ . This

Dataset	Question Prefix	Sample Clarification Question	Sample Clarification Answer
COPA	What is the definition of : 1	(a) What is the definition of an accident? (b) What is the definition of a flat tire?	(a) The definition of an accident is the crash of collusion caused by the vehicle. (b) The definition of a flat tire is that the tire does not hold air.
	What is the main purpose of : 2	(a) What is the main purpose of this investigation? (b) What is the main purpose of this post?	(a) The purpose of this investigation is to provide information about how and why he was shot. (b) The purpose of this post is to share my thoughts and feelings on his death.
CommonsenseQA	What is the main function of a : 2	(a) What is the main function of a teacher in this area? (b) What is the main function of a farmer?	(a) The main function of a teacher in this area is to teach them about life and love. (b) The main function of a farmer is to provide food for his family and the community.
	What might have caused : 3	(a) What might have caused this problem? (b) What might have caused the animal to flee?	(a) the cause of this problem was that his wife’s husband didn’t have enough money. (b) The cause of the animal to flee was a predator.
Social IQA	What did [NAME] do? : 1	(a) What did Kendall do? (b) What did Kendall do?	(a) What Kendall did was make sure that he wasn’t going anywhere else. (b) What Kendall did was so horrible, that it was hard to believe.
	How would you describe [NAME]? : 3	(a) How would you describe Riley? (b) How would you describe Riley?	(a) Riley is a big brother, he’s an awesome dad. (b) Riley is a very sensitive person and has a lot of anxiety.
Winogrande	What are the properties of a : 1	(a) What are the properties of a diet that is not healthy? (b) What are the properties of a home?	(a) The property of a diet that is not healthy are that it has high cholesterol (a good idea). (b) The properties of a home are that which makes it comfortable and pleasant for the occupants.
	What does it mean to : 2	(a) What does it mean to be an explorer? (b) What does it mean to be sophisticated?	(a) Be an explorer means to explore and make sense of things. (b) Be sophisticated means to be classy, elegant and smart.

Table 1: This table shows some of the question prefixes we used for different datasets in our experiments. We assign each prefix a level in Bloom’s Taxonomy. We show generated clarifications questions and answers for both Distil-GPT2 (a) and GPT-Neo (b) for their corresponding question prefixes.

approach simply concatenates each (prompt, question, answer) tuple into a single string of text  $T_o = [p, q, a_o]$  and feeds this string to the language model to assign each choice a score  $s_o = LM(T_o)$ . The language model’s answer is just the answer with the highest score:  $\hat{o} = \operatorname{argmax}_o s_o$ .

### 3.2 Self-talk Clarifications

Self-talk (Shwartz et al., 2020) has a language model ask itself clarifying questions then answer those questions to generate clarifications.

**Stage 1: Ask clarification questions.** To produce clarifications we start with a set of clarification question prefixes  $r_1, \dots, r_J$  that are designed specifically for each question answering dataset. “What happens if” is a sample prefix for the clarifications, shown in Fig. 1, and in Tab. 1 we present examples for all the datasets we use. In this stage the language model completes each of these prefixes, using its generator function  $LM_G$  to ask one question  $R_j = LM_G(r_j)$  per prefix.

**Stage 2: Answer the questions.** Next we use the model to answer each of these questions, possibly prompted with an answer prefix  $b_j$  corresponding to question prefix  $r_j$ . The results are the clarifications  $c_j = LM_G([R_j, b_j])$ .

**Stage 3: Reconsider with a new prompt.** To use the clarifications we pick one from the list then append it to the original prompt. This approach simply considers all combinations of clarifications questions and answers  $T_{j,o} = [p, q, c_j, a_o] \forall o, j$ , first chooses the clarification which maximizes model score per answer option, then chooses the final answer  $o^* = \operatorname{argmax}_o \max_j LM(T_{j,o})$ . This can improve question answering performance on its own, but in the next section we more carefully choose clarifications using our notion of proximal context and Bloom’s Taxonomy.

### 3.3 Using Bloom’s Taxonomy to Choose Clarifications with Proximal Context

To test our idea of proximal context we consider the level  $L$  of task give by each dataset then allow only proximal clarifications of level  $L - 1$ . We label each question prefix with the level of Bloom’s Taxonomy that it falls into, and then force the model to choose from the set  $\mathcal{C}_L$  of clarifications of level  $L$ . This results in a final choice for each level  $o_L^* = \operatorname{argmax}_o \max_{j \in \mathcal{C}_L} LM(T_{j,o})$ . We also provide a *Choice Baseline* that allows the model to choose any level of clarification to show the model would have difficulty choosing proximal clarifica-

tions itself. Note that the annotation of questions along Bloom’s taxonomy requires special skills typically found only among educators. While a layperson can be trained to annotate such questions, our experience was that it takes much more time than we could afford for a preliminary study such as this one. We therefore relied on our co-author, Sara Rutherford-Quach, who is a researcher at SRI’s Education Division and has also worked as a teacher at the kindergarten-elementary level to provide us the annotations. Two other co-authors, Sahu and Cogswell, went through those annotations and made sure that each label had a three way consensus among Rutherford-Quach, Sahu and Cogswell. There might be some ambiguity about which level a particular prefix fits into, but this is also true of other applications of the taxonomy (Thompson et al., 2008). In future work, we plan to carry out a more rigorous annotation with more than one skilled annotator so we can measure inter-annotator agreement through measures such as Kappa scores.

## 4 Experiments

### 4.1 Datasets

We evaluate our study on four datasets that can each be thought of in terms of multiple choice question answering, all measuring some kind of common sense: COPA (Roemmele et al., 2011) measures common sense causal reasoning, CommonsenseQA (Talmor et al., 2019) asks questions that require prior knowledge, Social IQA (Sap et al., 2019) asks about social common sense, and Winogrande (Sakaguchi et al., 2020) adversarially measures semantic common sense. Perhaps surprisingly, all of the datasets we used asked questions that fell into just one level of the taxonomy (Tab. 2). These datasets do focus on very specific problems, but the result is still disappointing because it would be more useful to see variations in both task and clarification level. It may be interesting to develop datasets that can better express the range of abilities described by Bloom’s Taxonomy.

### 4.2 Language Model

We use distil-GPT2 (Sanh et al., 2019) and the publicly released GPT-Neo2.7B (Black et al., 2021) (based on EleutherAI’s replication of the GPT-3 architecture) as the language models throughout our experiments. Our clarification question prefixes and hyperparameter settings for both models are

Table 2: Question answering accuracy and std. dev. using different levels of clarification over multiple clarification samples. Results on the dev sets of each dataset. (\* = level of proximal context \ wrt the dataset)

Task	Model	Level	Accuracy
Winogrande (1267 total) (2: Understand)	Distil-GPT2 (235±5 valid)	0A: Choice Baseline	53.2 ± 1.8
		1A: Remember*	<b>54.7 ± 3.6</b>
		2A: Understand	52.5 ± 3.1
	GPT-Neo (1230±7 valid)	0A: Choice Baseline	54.62 ± 0.5
		1A: Remember*	<b>54.77 ± 0.5</b>
		2A: Understand	54.76 ± 0.3
SocialIQA (1954 total) (3: Apply)	Distil-GPT2 (58±5 valid)	0B: Choice Baseline	44.5 ± 0.1
		1B: Remember	43.7 ± 2.1
		2B: Understand*	<b>48.0 ± 1.1</b>
	GPT-Neo (1334±9 valid)	3B: Apply	44.4 ± 1.8
		0B: Choice Baseline	<b>48.74 ± 0.4</b>
		1B: Remember	47.31 ± 0.1
		2B: Understand*	<b>48.44 ± 0.5</b>
		3B: Apply	48.1 ± 0.1
COPA (100 total) (3: Apply)	Distil-GPT2 (11±2 valid)	0C: Choice Baseline	<b>54.9 ± 0.9</b>
		1C: Remember	46.0 ± 14.7
		2C: Understand*	<b>53.1 ± 12.5</b>
	GPT-Neo (96±0 valid)	3C: Apply	40.8 ± 15.2
		0C: Choice Baseline	<b>70.83 ± 0.0</b>
		1C: Remember	65.62 ± 0.0
		2C: Understand*	<b>70.83 ± 1.4</b>
		3C: Apply	<b>70.83 ± 0.0</b>
CommonsenseQA (1221 total) (3: Apply)	Distil-GPT2 (68±1 valid)	0D: Choice Baseline	<b>29.9 ± 2.7</b>
		1D: Remember	26.5 ± 3.3
		2D: Understand*	<b>28.1 ± 1.2</b>
	GPT-Neo (1118±4 valid)	3D: Apply	25.6 ± 3.4
		0D: Choice Baseline	40.59 ± 3.6
		1D: Remember	38.00 ± 6.0
		2D: Understand*	<b>43.19 ± 0.2</b>
		3D: Apply	42.30 ± 0.8

from (Shwartz et al., 2020). For each question prefix, we generate 5 clarification questions using nucleus sampling threshold probability  $p = 0.2$  and adding at most 6 words to the clarification question prefix. We then generate 10 answers to each clarification question using  $p = 0.5$  and maximum answer length 10. Some changes were necessary to accurately measure the impact of clarification level. Instead of always including *no clarification* as a choice we do not allow this option as it defeats our goal of measuring clarification level impact. Furthermore, we do not use the clarification questions which were manually completed without input from the model (as in COPA and Winogrande).

In order to compare performance across different levels of clarifications we only consider examples where the model was able to generate at least one clarification from each level. To increase the number of viable examples we found it necessary to remove some restrictions relative to the implementation of (Shwartz et al., 2020). In particular, we kept all clarifications that had no overlapping words with the context and did not allow the model to chose the “no clarification” option. Even with these constraints it was still often the case that distil-GPT2 could not generate a short clarification



sentence that was plausible enough to use whereas GPT-Neo was able to generate clarifications for almost the entire dataset. This indicates larger scale models may be more able to take advantage of clarifying questions. The number of examples with valid clarifications for all levels is indicated for each model in column 2 of Tab. 2. These changes help us more accurately measure the impact of Bloom’s Taxonomy, but mean our approach is not directly comparable to Shwartz et al. (2020).

### 4.3 Results

Table 2 reports the performance of our Bloom’s Taxonomy infused zero-shot question answering method. Each row shows question answering accuracy for a particular dataset and level of clarification. If our hypothesis is correct then the level of available clarifications should matter and clarifications that provide proximal context –one level below the dataset level– should be most helpful.

**Clarification Level Makes a Difference.** All levels of clarification questions and answers provide some amount of extra information that changes how a language model processes the entire string it is presented with. This is often helpful information, but it may be that all levels of Bloom’s Taxonomy provide equally useful information. We find that is not the case. Different levels of clarification help more or less, as evidenced by the large gap between minimum and maximum accuracy for each dataset. Furthermore, when the model can choose any clarification (rows 0A/B/C/D) it either does a worse job than proximal context or its performance similar to proximal context, so enforcing a particular kind of context should be helpful.

**Proximal Context Helps Most.** Proximal context, as we’ve defined it with respect to Bloom’s Taxonomy is context from the clarification level directly below the dataset question level. The proximal clarification level for each dataset is marked by a \* in Tab. 2. In all cases proximal clarifications are better than using clarifications of a lower level. For the datasets that ask level 3 questions the proximal (level 2) clarifications also outperform level 1 clarifications (2B/C/D greater than 1B/C/D). Proximal clarifications are also about as good as or better than using clarifications of a higher level. You can see this for Winogrande by noting row 1A is greater than 2A and for the other datasets by noting rows 2B/C/D usually have greater performance than 3B/C/D. Overall, proximal context is most

consistent in efficacy.

### 4.4 Qualitative Results

In Tab. 1 we show samples of question answer pairs generated for each model and in Tab. 5 of the appendix we show complete examples (with context and choices) for each model and dataset. GPT-Neo is much larger than distil-GPT2 and is expected to generalize to slightly new tasks like the clarification generation task better than the smaller model. This expectation is clearly met by the observed quality of clarifications. Distil-GPT2 clarification questions and answers often do not have meaningful semantics, are not correct, or are not relevant. GPT-Neo is much more likely to generate questions and answers which are meaningful, correct, and relevant. This suggests the greater number of valid clarifications generated by GPT-Neo may be due to an increase in clarification quality. Furthermore, it fails in an intuitive fashion: when it fails to generate meaningful answers it often has also failed to generate a meaningful clarification question in the first place.

Also note that the performance differences observed for distil-GPT2 occur despite its relatively poor interpretability. This indicates that context which is somewhat relevant to the topic even if it does not precisely make sense can still be useful.

## 5 Conclusion

Large pre-trained language models sometimes have the right information, but they just do not know how to use it. We used Bloom’s taxonomy to pick questions with the right amount of proximal context. This helped the language models use their knowledge to more effectively answer questions. In the future we would like to extend our work on tasks that present a wide range of questions that fall under different levels of the taxonomy. Similarly, we also would like to study and improve upon the current limited set of prefix questions used.

### Acknowledgments

The authors thank Yunye Gong, Stephanie Nunn, and the anonymous reviewers for the helpful discussions and comments.

### References

L. Anderson, D. Krathwohl, and B. Bloom. 2000. A taxonomy for learning, teaching, and assessing: A

- revision of bloom’s taxonomy of educational objectives.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? . *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.
- Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large scale autoregressive language modeling with mesh-tensorflow](#).
- B. Bloom. 1956. Taxonomy of educational objectives: The classification of educational goals.
- Antoine Bosselut, Ronan Le Bras, and Yejin Choi. 2019a. Dynamic neuro-symbolic knowledge graph construction for zero-shot commonsense question answering. *arXiv preprint arXiv:1911.03876*.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019b. Comet: Commonsense transformers for automatic knowledge graph construction. *arXiv preprint arXiv:1906.05317*.
- T. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, J. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. Henighan, R. Child, A. Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, J. Clark, Christopher Berner, Sam McCandlish, A. Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. *ArXiv*, abs/2005.14165.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- 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*.
- Jesse Dunietz, Greg Burnham, Akash Bharadwaj, Jennifer Chu-Carroll, Owen Rambow, and D. Ferrucci. 2020. To test machine comprehension, start by defining comprehension. In *ACL*.
- Stephanie Harvey and Anne Goudvis. 2007. *Strategies that work: Teaching comprehension for understanding and engagement*. Stenhouse Publishers.
- Mandar Joshi, Kenton Lee, Yi Luan, and Kristina Toutanova. 2020. Contextualized representations using textual encyclopedic knowledge. *arXiv preprint arXiv:2004.12006*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Gary Marcus and Ernest Davis. 2020. [Gpt-3, bloviator: Openai’s language generator has no idea what it’s talking about](#). *MIT Technology Review*.
- Debbie Miller. 2002. *Reading with meaning: Teaching comprehension in the primary grades*. Stenhouse Publishers.
- Manal Mohammed and N. Omar. 2020. Question classification based on bloom’s taxonomy cognitive domain using modified tf-idf and word2vec. *PLoS ONE*, 15.
- Fatema Nafa, S. Othman, and J. Khan. 2016. Automatic concepts classification based on bloom’s taxonomy using text analysis and the naïve bayes classifier method. In *CSEDU*.
- D. Oliver, Tony Dobebe, Myles Greber, and Tim S. Roberts. 2004. This course has a bloom rating of 3.9. In *ACE*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- 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.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*, pages 90–95.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavathula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34-05, pages 8732–8740.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108.

- 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 4453–4463.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Unsupervised commonsense question answering with self-talk. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4615–4629.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158.
- E. Thompson, Andrew Luxton-Reilly, J. Whalley, M. Hu, and Phil Robbins. 2008. Bloom’s taxonomy for cs assessment. In *ACE ’08*.
- J. Whalley, R. Lister, E. Thompson, T. Clear, Phil Robbins, P. Kumar, and C. Prasad. 2006. An australasian study of reading and comprehension skills in novice programmers, using the bloom and solo taxonomies.
- James Zhang, Casey Wong, Nasser Giacaman, and Andrew Luxton-Reilly. 2021. Automated classification of computing education questions using bloom’s taxonomy. *Australasian Computing Education Conference*.

## A Prefixes and Examples

In the appendix we provides more details about the question prefixes we used in Tab. 3 and provide more examples of outputs from our models in Tab. 5.

Table 3: All the prefix questions with its corresponding taxonomy level used in our zero shot question answering evaluation.

Question Prefix	Answer Prefix	Bloom's Taxonomy Level
<b>CommonsenseQA &amp; COPA</b>		
What is the definition of	The definition of _ is	1
What is the main purpose of	The purpose of _ is to	2
What is the main function of a	The main function of a _ is	2
What are the properties of a	The properties of a _ are that	1
What is a	_ is	1
What happened as a result of	As a result of _,	3
What might have caused	The cause of _ was	3
<b>SocialQA</b>		
What will [NAME] want to do next?	[NAME] wanted	3
What will [NAME] want to do after?	[NAME] wanted	3
How would [NAME] feel afterwards?	[NAME] felt	3
How would [NAME] feel as a result?	[NAME] felt	3
How would [NAME] feel after?	[NAME] felt	3
How would you describe [NAME]?	[NAME] is a	2
What kind of person is [NAME]?	[NAME] is a	2
How would you describe [NAME] as a person?	[NAME] is a	2
Why did [NAME] do that?	[NAME] did this because they wanted	3
Why did [NAME] do this?	[NAME] did this because they wanted	3
Why did [NAME] want to do this?	[NAME] did this because they wanted	3
What does [NAME] need to do beforehand?	Before doing that, [NAME] first had to	2
What does [NAME] need to do before?	Before doing that, [NAME] first had to	2
What does [NAME] need to do before this?	Before doing that, [NAME] first had to	2
What did [NAME] need to do before this?	Before doing that, [NAME] first had to	2
What will happen to [NAME]?	[NAME]	3
What will happen to [NAME] next?	[NAME]	3
What will [NAME] do next?	[NAME]	3
What did [NAME] do?	What [NAME] did was	1
<b>Winogrande</b>		
What is the definition of	The definition of _ is	1
What is the main purpose of	The purpose of _ is to	2
What is the main function of a	The main function of a _ is	2
What are the properties of a	The properties of a _ are that	1
What is	_ is	1
What does it mean to	_ means	2



Table 4: Example contexts, questions, choices, clarification questions and clarification answers for each dataset. We present results for both Distil-GPT2 and GPT-Neo.

COPA	Generated Clarifications	
<p>(Context, Question, Choice)            C: The man turned on the faucet.            Q: effect            Ch: [The toilet filled with water.,  <b>Water flowed from the spout.</b>            C: The woman retired.            Q: effect            Ch: [<b>She received her pension.</b>,            She paid off her mortgage.]            C: I wanted to conserve energy.            Q: effect            Ch: [I swept the floor in the unoccupied room.,  <b>I shut off the light in the unoccupied room.</b>]</p>	<p><b>Distil-GPT2</b>            Q: What might have caused it to be so hot?            A: The cause of it to be so hot was that he would have a big fire in his face and then immediately had an explosion, which killed all but one person.            Q: What happened as a result of this?            A: As a result of this, the government has announced that it will pay Rs 7 crore to her family for compensation and legal expenses related thereto as part-time payments.            Q: What is the main function of a plant?            A: The main function of a plant is the growth and maintenance of nutrients, so that we can increase our food intake by 10% or more per year, he said.</p>	<p><b>GPT-Neo</b>            Q: What happened as a result of that?            A: As a result of that, he was forced to drink water.            Q: What happened as a result of her retirement?            A: As a result of her retirement, she had been receiving a pension for life.            Q: What might have caused this?            A: The cause of this was the lack of electricity in my home.</p>
Winogrande	Generated Clarifications	
<p>(Sentence, Context, Choice)            S: The home that my parents had when I was in school was a lot nicer than my house now because the _ was sophisticated.            C: home            Ch: [<b>home</b>, house]            S: He had enough time between classes to go to a cafe or to the library. He went to the _ because his paper could wait.            C: classes            Ch: [<b>cafe</b>, library]            S: They had to eat a lot to gain the strength they had lost and be able to work, the _ was too much.            C: strength            Ch: [<b>work</b>, strength]</p>	<p><b>Distil-GPT2</b>            Q: What is the main function of a building?            A: The main function of a building is to keep people happy and healthy.            Q: What is the definition of a book?            A: The definition of a book is an article written in order for it to be read.            Q: What is the main function of a good diet?            A: The main function of a good diet is not for food or health but also on exercise.</p>	<p><b>GPT-Neo</b>            Q: What are the properties of a home?            A: he properties of a home are that which makes it comfortable and pleasant for the occupants.            Q: What is the definition of cafe?            A: cafe is defined as : having power gather for refreshment.            Q: What is the definition of strength?            A: strength is defined as : having power or ability to do something.</p>
SocialQA	Generated Clarifications	
<p>(Context, Question, Choice)            C: Kendall got a new sports car and could not wait to show friends.            Q: What will Kendall want to do next?            Ch: [drive that sports car,  <b>show off his new sports car.</b>            clean and wax her legs]            C: Kendall was really into house plants and often pour water in to jars so the fluoride would evaporate out.            Q: What will Kendall want to do next?            Ch: [<b>watered a few of the plants.</b>            gather jars,            watered a few of her pets]            C: Remy gave birth to a baby girl on a Saturday afternoon.            Q: What does Remy need to do before this?            Ch: [leave her husband,            did this to start a family,  <b>go to the hospital]</b></p>	<p><b>Distil-GPT2</b>            Q: Why did Kendall want to do this?            A: Kendall did this because they wanted her, too.            Q: Why did Kendall do that?            A: Kendall did this because they wanted him at home, he said.            Q: What did Remy need to do before this?            A: Before doing that, Remy first had to tell her he was too old.</p>	<p><b>GPT-Neo</b>            Q: What will Kendall want to do next?            A: Kendall wanted to impress them with his new ride."            Q: Why did Kendall do this?            A: Kendall did this because they wanted to get a lot of water for their plants.            Q: What does Remy need to do before?            A: Before doing that, Remy first had to go to the hospital to deliver her baby.</p>
CommonsenseQA	Generated Clarifications	
<p>(Context, Question, Choice)            C: In what Spanish speaking North American country can you get a great cup of coffee?            Q: cup of coffee            Ch: [mildred's coffee shop, <b>mexico</b>, diner,            kitchen, canteen]            C: Unlike a spider and his many sight seers, people only have what?            Q: people            Ch: [tongues, names, brains, feelings, <b>two eyes</b>]            C: The artist was sitting quietly pondering, then suddenly he began to paint when what struck him?            Q: sitting quietly            Ch: [sadness, anxiety, <b>inspiration</b>, discomfort, insights]</p>	<p><b>Distil-GPT2</b>            Q: What is the definition of coffee?            A: The definition of coffee is the word that comes from Latin for espresso.            Q: What is the main purpose of this article?            A: The purpose of this article is to explain the basic characteristics that are characteristic of spiders.            Q: What is a man who can do that?            A: Man who can do that is a really great painter.</p>	<p><b>GPT-Neo</b>            Q: What is the definition of mexico?            A: mexico is defined as the country that produces the best coffee.            Q: What is the definition of two eyes?            A: two eyes is defined as the organs by which we see.            Q: What is the definition of inspiration?            A: inspiration is defined as a sudden flash of inspiration that comes from within.</p>