

# DaVinci at SemEval-2024 Task 9: Few-shot prompting GPT-3.5 for Unconventional Reasoning

Suyash Vardhan Mathur\*  
IIIT Hyderabad

Akshett Rai Jindal\*  
IIIT Hyderabad

Manish Shrivastava  
IIIT Hyderabad

suyash.mathur@research.iiit.ac.in akshett.jindal@research.iiit.ac.in

m.shrivastava@iiit.ac.in

## Abstract

While significant work has been done in the field of NLP on vertical thinking, which involves primarily logical thinking, little work has been done towards lateral thinking, which involves looking at problems from an unconventional perspective and defying existing conceptions and notions. Towards this direction, SemEval 2024 introduces the task of BRAINTEASER, which involves two types of questions – Sentence Puzzles and Word Puzzles that defy conventional common-sense reasoning and constraints. In this paper, we tackle both types of questions using few-shot prompting on GPT-3.5 and gain insights regarding the difference in the nature of the two types. Our prompting strategy placed us 26<sup>th</sup> on the leaderboard for the Sentence Puzzle and 15<sup>th</sup> on the Word Puzzle task.

## 1 Introduction

The human brain consists of two hemispheres - left and right. Both of them are responsible for different kinds of thinking strategies. The left hemisphere is involved in vertical thinking, and the right hemisphere is involved in lateral thinking (Waks, 1997). Vertical (linear, convergent, logical) thinking is a more sequential analytical process. In contrast, in Lateral (outside the box, divergent, creative) thinking, we look at the problem from a new point of view, ignoring the expected associations with items.

In the field of NLP, much research has been done around vertical thinking and significant progress has been made. The recent work around Large Language Models (LLMs) (Devlin et al., 2018; OpenAI, 2022) has achieved great performance in solving complex reasoning tasks (Talmor et al., 2018; Bisk et al., 2020; Sap et al., 2019). This performance is consistent in both cases when no task examples have been provided to the model

during inference (zero-shot) (Sanh et al., 2022) and when the model is introduced with the task during inference time (few-shot) (Chung et al., 2022).

However, lateral thinking has been overlooked when training NLP models like LLMs. When creating datasets for various models, texts that involve lateral thinking are mostly considered noise and filtered out from the data because researchers want their models to perform better at traditional reasoning tasks and not get confused by lateral thinking.

The task BRAINTEASER (Jiang et al., 2023, Jiang et al., 2024) tries to bridge this gap that exists between vertical and lateral thinking for LLMs and other NLP models. They formulated a set of Multi-choice Question Answers containing puzzles that can be solved only using lateral thinking. The benchmark dataset contains two types of lateral thinking puzzles - Sentence Puzzles and Word Puzzles. This has been constructed by designing a data collection procedure that crawled relevant puzzles from many websites that were publicly available performing semi-automatic filtering of irrelevant questions.

## 2 Background

### 2.1 Dataset

The dataset being used in this task is BRAINTEASER (Jiang et al., 2023). It was prepared by scraping puzzles from various publicly available websites and then semi-automatically filtering them out. Then *semantic reconstruction* and *context reconstruction* techniques were used to create variants of each puzzle without affecting its out-of-the-box thinking style. This helped in preventing possible memorization by LLMs and the lack of consistency of the puzzles.

The puzzles in this dataset can be divided into two categories:

- **Sentence Puzzles:** These are brain teasers where the puzzle-defying commonsense is

\*Equal Contribution

centered on sentence snippets.

For example, **Question:** *You are running so fast but you're not getting closer. Where are you?* **Answer:** *Treadmill.* **Explanation:** This is because while running on a treadmill, we stay put where we are. The key is understanding that running on a treadmill means you remain stationary despite the motion.

- **Word Puzzles:** These are brain teasers where the answer violates the default meaning of the word and focuses more on the letter composition.

For example, **Question:** *How can you make "ten" out of "net"?* **Answer:** *Just flip it around.* **Explanation:** This is because if we consider the spelling of the word "ten" and we flip the letters of the word around, we get the word "net" which is what we want to make out of "ten".

The training data contains **507 Sentence Puzzles** and **396 Word Puzzles**. Each of these puzzles has 4 options to choose from and only one option is the correct answer.

## 2.2 Related Works

With the recent success of LLMs in various NLP tasks, researchers have also started exploring their use for Multiple Choice Question Answering (MCQA) tasks (Robinson et al., 2022; Zheng et al., 2023).

Researchers have also started employing the technique of **few-shot prompting** (Liu et al., 2023; Ma et al., 2024; Lu et al., 2021) for various tasks and it has shown improvements when compared with **zero-shot prompting**.

LLMs like GPT-3.5 have been trained on vast amounts of human-generated text. The main features around which such models are trained are **Pattern Recognition**, **Creative Reasoning** and **Wide Knowledge Range**.

Thus, we decided to employ few-shot prompting on LLMs for this task.

## 3 System Overview

Our architecture uses GPT-3.5 (Brown et al., 2020) (specifically *gpt-3.5-turbo*) with few-shot prompting to answer the question.

### 3.1 GPT-3.5

In NLP, the architecture of Generative Pre-trained Transformer (GPT) 3.5 (GPT-3.5) stands as a significant advancement, which is the culmination of iterative improvement over its predecessors. The architecture of the model is based upon the Transformer model (Vaswani et al., 2017), which uses self-attention to enhance performance over the prior sequential models. GPT-3.5 scales this Transformer architecture to over hundreds of billions of parameters, which have been trained by exposing and training the model on hundreds of billions of tokens.

In particular, due to the autoregressive nature of GPT-3.5 and due to being trained on extremely large data, it has enough knowledge about the language and the real world to perform tasks in a Zero-shot setting (Sanh et al., 2022). This Zero-shot setting allows the model to understand and execute a task it hasn't been explicitly trained for. These capabilities have been reflected in GPT-3.5 being used in Summarization (Liu and Healey, 2023), Question Answering (Bahak et al., 2023), Natural Language Inference (Ye et al., 2023), etc.

### 3.2 Few-shot prompting

While zero-shot prompting works well for simple tasks, tasks like BrainTeaser are a bit more complex in nature, and in such cases providing explicit instructions to the LLM about the nature of the task along with few examples of the task (*few shots*) becomes extremely helpful for the model (Chung et al., 2022). Here, the few-shot technique involves providing GPT-3.5 some examples, allowing GPT-3.5 to generalize from the few examples, drawing on its large pre-trained knowledge about the language and the real-world.

Thus, 2 different sets of prompts are created for the task, one for the Sentence Puzzle Task and another for the Word Puzzle task, since the 2 tasks are fundamentally different and need different instructions and examples.

### 3.3 Experimental Setup

We provide 2-shot prompts to GPT-3.5 for our leaderboard submission. We also try out 5-shot prompt in the post evaluation phase to test if providing more examples helps the model perform better.

The prompt used for the Sentence Puzzle task is shown in Listing 1. As we can see, the prompt first

You are given a question with multiple choices that you need to answer.

- ↪ The answer would only be one index of the multiple choices available.
- ↪ Such a question would involve brain teaser questions where the puzzle defying commonsense is centered on sentence snippets.

**IMPORTANT:** It's crucial to analyze the question from an unconventional perspective, focusing on the literal or alternative meanings of the words used, rather than relying on common sense. You must not use commonsense, but look at meaning from a different perspective than what would commonly be done. For example,

Example 1:  
 Question: You are running so fast but you're not getting closer. Where are you?

Option 0: Country road.  
 Option 1: Treadmill.  
 Option 2: High way.  
 Option 3: None of above.

Answer: 1  
 Reason: This is because while running on a treadmill, we stay put where we are. The key is understanding that running on a treadmill means you remain stationary despite the motion. This is not valid for Country road or High way. Thus, the answer is 1 - Treadmill.

Example 2:  
 Question: From elementary school to collage, how many "first day of school" does the average person have in their lifetime?

Option 0: They technically only have one first day of school in their lifetime. That's the very first day they started attending school as a child.  
 Option 1: Average people have 4: elementary school, middle school, high school, and college.  
 Option 2: Average people have "first day of school" in each semester, so it will be more than 10!  
 Option 3: None of above.

Answer: 0  
 Reason: First day of school can only be one day in a person's lifetime. Here, it is important to understand that first day of middle school, high school, college won't be first day of school. Similarly, each semester's first day is not **TECHNICALLY** first day of school. This, the answer is 0 - They technically only have one first day of school in their lifetime. That's the very first day they started attending school as a child. Thus, the key here is the term 'first day of school' technically refers to the very first day a person attends school, making all subsequent 'first days' at different educational levels irrelevant to the specific question.

Now, using these examples, answer the question below. It is **IMPORTANT** that you just provide the index of the answer in the response. DO NOT output the reason behind choosing the answer:

Question: In a small village, two farmers are working in their fields - a diligent farmer and a lazy farmer. The hardworking farmer is the son of the lazy farmer, but the lazy farmer is not the father of the hardworking farmer. Can you explain this unusual relationship?

Option 0: The lazy farmer is his mother.  
 Option 1: The lazy farmer is not a responsible father as he is lazy.  
 Option 2: The diligent farmer devoted himself to the farm and gradually forgot his father.  
 Option 3: None of above.

Answer:

Listing 1: Prompt for the Sentence Puzzle

details the task, and *IMPORTANT* keyword is used to express to GPT-3.5 that commonsense must not be used in the task, but instead it should look at meaning from an unconventional sense. Then, 2 examples are given, along with **reasoning** behind the answers too. This was important, as this gave the model more knowledge to be able to generalize the task from the examples. Further, the output format was clearly specified in the prompt so as to avoid getting extra information in the model output.

Similar prompt for Word Puzzle can be seen in Listing 2. The prompt clarifies that the struc-

You are given a question with multiple choices that you need to answer.

- ↪ The answer would only be one index of the multiple choices available. The question demands an unorthodox approach, focusing on the spellings or structural aspects of words, rather than their standard meanings. Your task is to choose the correct answer from the given multiple-choice options by analyzing the words in a literal or unconventional way.

**IMPORTANT:** It's crucial to analyze the question from an unconventional perspective, focusing on the spellings of certain words, rather than relying on common sense. You must not use commonsense, but look at meaning from a different perspective considering arrangement of the letters in certain words than what would commonly be done. For example,

Example 1:  
 Question: How can you make "ten" out of "net"?

Option 0: Just flip it around.  
 Option 1: Remove the letter "e".  
 Option 2: Move the letter "t" to the end.  
 Option 3: None of above.

Answer: 0  
 Reason: This is because if we consider the spelling of the word 'ten' and we flip the letters of the word 'ten' around, we get the word 'net', which is what we want to make out of 'ten'. The answer focuses on the literal rearrangement of the letters, disregarding the typical meanings of the words. Thus, the answer is 0 - Just flip it around.

Example 2:  
 Question: What is the most fast city?

Option 0: Urban city.  
 Option 1: Inner city.  
 Option 2: Velocity.  
 Option 3: None of above.

Answer: 2  
 Reason: The term 'fast' in the question prompts an unconventional interpretation. All options contain the word "city", but "velocity" stands out as it directly relates to speed or 'fastness'. The question cleverly uses the term 'city' as a red herring, while the actual focus is on the concept of speed.

Now, using these examples, answer the question below. It is **IMPORTANT** that you just provide the index of the answer in the response. DO NOT output the reason behind choosing the answer:

Question: What sort of cheese is made in reverse?

Option 0: Cheddar cheese..  
 Option 1: Edam cheese.  
 Option 2: Blue cheese.  
 Option 3: None of above.

Answer:

Listing 2: Prompt for the Word Puzzle

tural aspect of the words should be focused on, emphasizing that unconventional meaning should be looked at. Then, 2 examples that exhibit structural aspect are given along with reasoning behind their answers as well as constraints for the output format.

## 4 Results and Analysis

The results are detailed in Tab.1. For comparison, we also list the zero-shot prompting results reported in Jiang et al., 2023. As we can see, the two-shot performance on Word puzzle improved over the zero-shot setting for all the categories, while the same worsened in case of the Sentence puzzle.

This is because of the very nature of the two problems. Sentence puzzle involves deeper non-

Table 1: Results of zero-shot and few-shot prompting on GPT-3.5 for the two BRAINTEASER subtasks. Ori = Original, Sem = Semantic, Con = Context.

Model	Instance-based			Group-based		Overall
	Original	Semantic	Context	Ori & Sem	Ori & Sem & Con	
<i>Sentence puzzle</i>						
GPT-3.5 (zero-shot) <b>Baseline</b>	60.7	59.3	<b>67.9</b>	50.7	39.7	<b>62.6</b>
GPT-3.5 (two-shot)	57.5	55.0	42.5	50.0	30.0	51.7
GPT-3.5 (five-shot)	<b>62.5</b>	<b>65.0</b>	55.0	<b>62.5</b>	<b>42.5</b>	60.8
<i>Word puzzle</i>						
GPT-3.5 (zero-shot) <b>Baseline</b>	56.1	52.4	51.8	43.90	29.3	53.5
GPT-3.5 (two-shot)	71.9	71.9	62.5	59.4	46.9	68.6
GPT-3.5 (five-shot)	<b>78.1</b>	<b>90.6</b>	<b>84.4</b>	<b>78.1</b>	<b>68.8</b>	<b>84.4</b>

conventional semantic understanding of the question and the choices, which despite conveying reasoning behind the answers in the few-shot examples, cannot be generalized as easily with just 2 examples. On the other hand, the only tricky component of the Word Puzzle is that the structural aspect of certain words needs to be taken instead of the actual surface meaning of the said words. This can be much more easily generalized through just as few as two examples in the prompt. Further, adding the examples in the Sentence puzzle that don't generalize very well for other questions in the testing set might have acted as noise for the model, which led to poorer performance.

We also note that using five-shot prompt instead of two-shot prompt hugely increases the performance. This is to be expected, as providing more examples would help the model generalize even better towards solving the task. This is specially true for Word Puzzle questions, where adding more examples allows the model to generalize the task much better.

However, in Sentence Puzzle we still notice a drop in the overall performance as compared to the zero-shot model. This is because of a drop in the performance of the context reconstruction questions, and a marginal increase in comparison to zero-shot in other types of questions. However, group based accuracy increases in five-shot, which might indicate that with five examples, the model is able to handle the variations in reconstructions better, albeit with performance of Contextual Reconstruction taking a dip. These observations are in line with the drop observed in two-shot prompt in comparison to the zero-shot prompt, highlighting the difficult nature of the task of Sentence Puzzle questions and the inability of the model to generalize using few Sentence Puzzle examples. However, we do note that the performance on Sentence Puzzle also does improve with additional examples

between two-shot and five-shot prompting.

## 5 Conclusion

In conclusion, we explored the effectiveness of few-shot prompting for LLMs for complex and unconventional tasks. Further, it demonstrates that few-shot prompting is helpful only in scenarios where the examples convey enough information that can be better generalized, as the results worsened in the Sentence Puzzle while improved in the Word Puzzle.

In future, better prompting strategies like Chain of Thought prompting (Wang et al., 2023) can be utilized to improve the performance. Additionally, finetuning the pre-trained LLMs might also help in the task further. Also, increasing the number of training examples might help in further improving the model performance, as observed in the gains of performance in the five-shot prompt in comparison to the two-shot prompt.

## References

- Hossein Bahak, Farzaneh Taheri, Zahra Zojaji, and Arefeh Kazemi. 2023. [Evaluating chatgpt as a question answering system: A comprehensive analysis and comparison with existing models](#).
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish,

- Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).
- 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*.
- Yifan Jiang, Filip Ilievski, and Kaixin Ma. 2024. [Semeval-2024 task 9: Brainteaser: A novel task defying common sense](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 1996–2010, Mexico City, Mexico. Association for Computational Linguistics.
- Yifan Jiang, Filip Ilievski, Kaixin Ma, and Zhivar Sourati. 2023. [BRAINTEASER: Lateral thinking puzzles for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14317–14332, Singapore. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35.
- Sengjie Liu and Christopher G. Healey. 2023. [Abstractive summarization of large document collections using gpt](#).
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*.
- Huan Ma, Changqing Zhang, Yatao Bian, Lemao Liu, Zhirui Zhang, Peilin Zhao, Shu Zhang, Huazhu Fu, Qinghua Hu, and Bingzhe Wu. 2024. Fairness-guided few-shot prompting for large language models. *Advances in Neural Information Processing Systems*, 36.
- OpenAI. 2022. Chatgpt: A language model by openai. <https://www.openai.com>.
- Joshua Robinson, Christopher Michael Rytting, and David Wingate. 2022. Leveraging large language models for multiple choice question answering. *arXiv preprint arXiv:2210.12353*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Tali Bers, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#).
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Shlomo Waks. 1997. Lateral thinking and technology education. *Journal of Science Education and Technology*, 6:245–255.
- Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun. 2023. [Towards understanding chain-of-thought prompting: An empirical study of what matters](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2717–2739, Toronto, Canada. Association for Computational Linguistics.
- Junjie Ye, Xuanting Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, Zeyang Zhou, Chao Gong, Yang Shen, Jie Zhou, Siming Chen, Tao Gui, Qi Zhang, and Xuanjing Huang. 2023. [A comprehensive capability analysis of gpt-3 and gpt-3.5 series models](#).
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. On large language models’ selection bias in multi-choice questions. *arXiv preprint arXiv:2309.03882*.