

# AILS-NTUA at SemEval-2024 Task 9: Cracking Brain Teasers: Transformer Models for Lateral Thinking Puzzles

Ioannis Panagiotopoulos, Giorgos Filandrianos, Maria Lymperaiou, Giorgos Stamou

School of Electrical and Computer Engineering, AILS Laboratory

National Technical University of Athens

yiannispn@gmail.com, {geofila, marialymp}@islab.ntua.gr, gstam@cs.ntua.gr

## Abstract

In this paper, we outline our submission for the SemEval-2024 Task 9 competition: 'BRAIN-TEASER: A Novel Task Defying Common Sense'. We engage in both sub-tasks: Sub-task A-Sentence Puzzle and Sub-task B-Word Puzzle. We evaluate a plethora of pre-trained transformer-based language models of different sizes through fine-tuning. Subsequently, we undertake an analysis of their scores and responses to aid future researchers in understanding and utilizing these models effectively. Our top-performing approaches secured competitive positions on the competition leaderboard across both sub-tasks. In the evaluation phase, our best submission attained an average accuracy score of 81.7% in the Sentence Puzzle, and 85.4% in the Word Puzzle, significantly outperforming the best neural baseline (ChatGPT) by more than 20% and 30% respectively.

## 1 Introduction

In Natural Language Processing (NLP), reasoning serves as the cognitive backbone, enabling systems to transcend mere language comprehension and delve into sophisticated understanding. Despite the excellence of Large Language Models (LLMs) in several linguistic tasks, their reasoning capabilities are still questionable to a non-negligible extent (Floridi and Chiriatti, 2020; Bender et al., 2021; Kauf et al., 2022; Zhang et al., 2023; Shi et al., 2023; Tyen et al., 2024; Giadikiaroglou et al., 2024), often posing the fundamental concerns of whether they can indeed reason or memorize exhaustively (Yuan et al., 2022).

Such limitations can be probed via well-crafted datasets and benchmarks, showcasing varying LLM deficiencies at a time. As the core of the current paper, BrainTeaser (Jiang et al., 2023b, 2024b) incorporates problems that stress models to think "out-of-the-box"; to this end, the key novelty of BrainTeaser is that in order to answer correctly,

models need to defy default senses of concepts and common associations. Surprisingly, state-of-the-art (SoTa) LLMs, such as ChatGPT can only exhibit a maximum accuracy of  $\sim 60\%$  when solving Brain-Teaser riddles, demonstrating an inherently limited reasoning ability in unconventional thinking.

Thus, assuming that large-scale training and prompting may not always serve as universally applicable solutions towards flexible reasoning, we move one step back and leverage transfer learning techniques starting from smaller models based on masked language modelling, such as BERT (Devlin et al., 2019) and consequent BERT-based encoders. Then, we proceed with similar techniques on LLMs, aiming to showcase that significant performance advancements using a small set of in-domain data for parameter updating can be achieved in comparison to merely querying the model's prior knowledge via prompting. Therefore, our contributions are:

1. We perform lightweight tuning on smaller encoder models and LLMs, significantly outperforming the reported baselines.
2. We transform the multiple-choice problem to a binary classification one, aiming to explore diverging reasoning paths for models.
3. We ground final performance on the models' "prior knowledge" in related problems.
4. We delve into models' frequent failures to obtain a deeper understanding of reasoning cues that make models struggle the most.

Our code is available on GitHub <sup>1</sup>.

## 2 Related work

**Reasoning in NLP** has enjoyed several advancements due to the surge of pre-trained language mod-

<sup>1</sup><https://github.com/GiannisPana/AILS-NTUA-at-SemEval-2024-Task-9-Braineater>

els and especially LLMs (Sun et al., 2023). Reasoning challenges incorporate commonsense reasoning (Richardson and Heck, 2023), involving inference regarding everyday situations, mathematical reasoning (Lu et al., 2023), referring to the ability of solving mathematical problems, logical reasoning (Yang et al., 2023), which includes the systematic deduction of conclusions based on established principles and formal rules, causal reasoning (Gendron et al., 2024), which studies cause-and-effect relationships explaining why an event leads to another, and several other sub-tasks (Vashishtha et al., 2020; Wei et al., 2023; Petersen and van der Plas, 2023). In terms of reasoning evaluation, BigBench (Srivastava et al., 2023) comprises 204 reasoning tasks, targeting to explore the related capabilities of recent LLMs. Several dedicated datasets have been developed to tackle different reasoning challenges, including commonsenseQA (Talmor et al., 2019), WinoGrande (Sakaguchi et al., 2019), RiddleSense (Lin et al., 2021) and others; most of these datasets are incorporated in Tasksource (Sileo, 2023). Especially RiddleSense questions aspects of reasoning close to BrainTeaser (Jiang et al., 2023b, 2024b).

### 3 Task and Dataset Description

The BrainTeaser task at SemEval-2024 (Jiang et al., 2023b, 2024b) features lateral thinking puzzles presented as multiple-choice questions (QAs). Each question offers four options, with one being the correct answer and the others serving as distractors. Additionally, the final option is always "None of above". It consists of two sub-tasks, *Task A: Sentence Puzzle* and *Task B: Word Puzzle*. In addition to the original puzzles, the dataset includes adversarial subsets created by manually modifying the original brain teasers while preserving their reasoning paths. The original data were perturbed in two ways: First, there is *semantic reconstruction* of each original question without altering the answers or the distractors. Second, the original data underwent *context reconstruction*, wherein the original reasoning path remains intact, but the brain teaser describes a new situational context. Overall, the dataset used for training and evaluation consists of triplets of data: original, semantic, and context reconstruction. Table 1 provides an example of the triplets of data that constitute the dataset.

**Task A: Sentence Puzzle** In this sub-task, the sentence pairs are crafted in a manner that makes it relatively easy for humans to discern the correct

Question	Choice
<i>Original</i>	
What kind of nut has no shell?	A peanut.
	<b>A doughnut.</b>
	A walnut.
	None of above.
<i>Semantic Reconstruction</i>	
Which nut doesn't have a shell?	<b>A doughnut.</b>
	A walnut.
	A peanut.
	None of above.
<i>Context Reconstruction</i>	
Which type of bell doesn't make a sound?	A fire bell.
	A cow bell.
	<b>A bluebell.</b>
	None of above.

Table 1: Illustration of the structure of each sub-task's dataset, showcasing the original statement along with its two adversarials.

statement, yet challenging for systems, even those equipped with commonsense understanding. Table 2 contains examples of the Sentence Puzzle dataset (on the left). The training data consists of 169 distinct multiple-choice QA sets, each accompanied by its semantic and context reconstructions, resulting in a total of 507 multiple-choice questions ( $3 \times 169$ ).

**Task B: Word Puzzle** involves word-type brain teasers, where the answer defies the default meaning of the word and focuses on the letter composition of the question. The training dataset comprises 132 multiple-choice QAs, each accompanied by its semantic and context reconstructions, resulting in a total of 396 multiple-choice QAs ( $3 \times 132$ ). These brain teaser categories include puns, homophones, ambiguous words, and various other linguistic puzzles, as showcased in the examples provided in Table 2 on the right-hand side. The Word Puzzle sub-task pose challenges not only for systems but also for humans in discerning the correct answer.

**Data statistics** The BrainTeaser dataset comprises 3 data splits, namely train, development (used during the practice phase), and the hidden test set, which was used for evaluation. Statistics are provided in Table 3. Throughout the evaluation phase, the leaderboard was kept concealed.

**Evaluation Metrics** Both sub-tasks are assessed via accuracy metrics to gauge the performance of participating systems in two ways. First, instance-based accuracy evaluates each question individually, considering original questions and their seman-

Sentence Puzzle		Word Puzzle	
Question	Choice	Question	Choice
A man shaves everyday, yet keeps his beard long.	<b>He is a barber.</b>	What has toes but no feet or legs?	Cabbages.
	He wants to maintain his appearance.		<b>Tomatoes.</b>
	He wants his girlfriend to buy him a razor.		Onions.
	None of above.		None of above.
You go to the doctor because you're sick, and he gives you three medicines to take every half hour. How long do the drugs keep you going?	One and a half hours.	What did the little lobster get on its math test?	<b>Sea-plus.</b>
	<b>Two hours.</b>		Very-bad.
	An hour.		Very-Good.
	None of above.		None of above.
How many times can you deduct 10 from 100?	<b>Once.</b>	What's the beginning of an argument?	The letter T.
	Infinite time.		<b>The letter A.</b>
	Twice.		The letter U.
	None of above.		None of above.

Table 2: Example questions illustrating both sub-tasks, with correct answers highlighted in bold. Examples on the left pertain to *sub-task A: Sentence Puzzle*, while those on the right correspond to *sub-task B: Word Puzzle*.

Sub-task	Train	Dev	Test
A - Sentence Puzzle	507	120	120
B - Word Puzzle	396	96	96

Table 3: Data statistics.

tic and context adversarials. This metric provides a detailed understanding of a model’s proficiency in reasoning through various scenarios. In contrast, group-based accuracy takes a broader perspective, assessing questions and associated adversarials as cohesive groups. Each group consists of three questions, and a model scores 1 only if it correctly solves *all* questions in a group. This approach evaluates the system’s holistic performance in navigating through lateral thinking challenges. The combined use of instance-based and group-based accuracy metrics provides comprehensive insights into the capabilities of participating systems in tackling the complexities of both sub-tasks.

## 4 Methods

We focus on tuning language models belonging into two categories. First, we fine-tune variations of *encoder* models, namely BERT (Devlin et al., 2019), RoBERTa-large (Liu et al., 2019) and DeBERTaV3-base (He et al., 2023), to assess the impact of transfer learning using various datasets requiring similar reasoning abilities, apart from BrainTeaser. We study the problem using the provided *multi-choice* setup, but we also transform it into a *binary* classification task. Secondly, the encoders’ results are compared with those obtained from *fine-tuned LLMs* using the BrainTeaser dataset. To achieve this, we fine-tune Llama 2 (Touvron et al., 2023b), Phi-2 (Gunasekar et al., 2023) and Mistral-7b (Jiang et al., 2024a), which have already demonstrated enhanced reasoning abilities. In this regard,

we examine the effect of the model size on our task, which has already been reported in the literature to significantly influence the reasoning abilities of the models (Touvron et al., 2023b; Wei et al., 2022), along with other tuning hyperparameters. Model details are presented in App. A.

### 4.1 Encoder models

**Pre-training** First, we evaluate the effects of the pre-training on our task. Thus, we select two variations of each encoder: the *vanilla* one (using the default pre-trained basis and fine-tuned on BrainTeaser data only) and one that has undergone additional pre-training using supplementary commonsense reasoning datasets before fine-tuned on BrainTeaser. In the second case, we use the following pre-trained models: ① BERT-SE: a BERT-base-uncased version pre-trained on the multiple-choice dataset used in SemEval-2020 Task 4b (Wang et al., 2020) ② RoBERTa-WNGRD: a RoBERTa-large version pre-trained on the WinoGrande dataset, and ③ DeBERTaV3-TS: a DeBERTaV3-base model, pre-trained on diverse commonsense reasoning datasets, and fine-tuned with multi-task learning on over 600 tasks from the Tasksource collection.

**Multi-class Classification task** This strategy involves treating the problem as multi-class classification: all four provided options are combined with the given question, and consequently these concatenated inputs are fed into the model, which is fine-tuned to select one of the four options as part of a multi-class classification problem.

**Binary Classification task** Each sample originally consisting of multiple-choice QAs with four available options, underwent the following transformation: each candidate answer (excluding the

"None of above" option) was paired with the question receiving the label 0 if the choice was incorrect, or the label 1 for the opposite. In case all the 3 pairings returned 0, it is directly implied that "None of above" is the correct answer.

## 4.2 LLMs

We demonstrate an in-depth examination of fine-tuning SoTa LLMs (Llama 2, Phi-2, and Mistral-7b) in the context of multi-class classification. Note that during inference, the models prompted to provide an *explanation* along with the label. This experimental step, which we have observed to improve the performance of the model, also provides a qualitative identification of flaws in the models' reasoning process. In our experiments, we explore various combinations of LoRA (Hu et al., 2021)  $\alpha$  and  $r$  hyperparameters, using values of 16, 32, 64, and 128. For the analysis ahead, LLMs are denoted as model\_r\_a, reflecting these hyperparameters. Additional technical information, including prompting details and specifics about QLoRA hyperparameters, is available in App. B, C, D.

## 5 Experimental Results

Our metrics for the Sentence Puzzle sub-task are presented in Table 4 and for the Word Puzzle sub-task in Table 5 along with their baselines. Interestingly, the performance of the binary classification problem is significantly lower than that of the multi-class classification task. Initially, this behavior seemed counterintuitive since it appeared easier to determine whether a question is correct or not than to select the correct answer from four different options. However, this assumption is not accurate. Consider the word riddle: "What is the capital in France?" At first glance, the option 'F' seems incorrect, but when considering the options 'F', 'E', 'A', and 'None of the above', 'F' emerges as the only correct answer, as it becomes apparent that the question refers to the capital *letter* rather than the capital *city*. Therefore, the diverse options provide crucial context to the models, explaining the superior performance of multi-class models. This lack of context is why we refrain from further exploring this methodology across all models in our study.

**Task A: Sentence Puzzle** Table 4 illustrates minimal fluctuations among all instance-based metrics. This consistency extends to the associated group-based metrics for all models, highlighting a systematic behavior towards detecting various rea-

soning paths. This observation holds for both the encoder-based classifiers and LLMs utilized in this sub-task. Sentence puzzles inherently offer more detailed information, enabling models to detect and identify the same reasoning patterns more readily, regardless of changes in context, in contrast to word puzzles, which typically feature shorter contextual statements, presenting a greater challenge for models to discern consistent reasoning patterns.

Initially, it becomes apparent that pre-training encoders across various commonsense reasoning datasets results in substantial performance enhancements, as it enables the system to grasp domain-agnostic features which prove advantageous for the subsequent task. Additionally, several commonsense pre-trained encoders fine-tuned on BrainTeaser data outperform Llama 2 and Phi-2.

Another noteworthy observation from Table 4 is that only Mistral-7b from LLMs is able to surpass the encoder-type networks, while both Llama 2 and Phi-2 consistently scored lower. Unlike Llama 2 and Mistral-7b, Phi-2 has not undergone instruction fine-tuning (Gunasekar et al., 2023), which, coupled with the limited number of examples in the BrainTeaser Sentence Puzzle dataset, contributes to its lower performance, as a result of Phi's incapability to capture the complexities of the BrainTeaser data. In this regard, Mistral-7b, which has already demonstrated superior performance compared to every Llama 2 variation when tested in commonsense reasoning benchmarks (Jiang et al., 2023a), is also capable of solving this task more accurately.

**Task B: Word Puzzle** In Table 5, we observe a stark contrast in the models' performance in understanding and detecting reasoning paths when the context changes. There are notable discrepancies in accuracy between original and semantic contexts when compared to context reconstruction, particularly evident in the case of smaller encoder models.

Regarding encoders, it is evident that, especially vanilla RoBERTa-large lacks robust commonsense reasoning and struggles to systematically handle ambiguity; in contrast, RoBERTa-large pre-trained on WinoGrande presents competitive performance. This notable enhancement (over 40%) due to WinoGrande pre-training suggests that this particular dataset effectively equips the model with the ability to understand word puzzle-related reasoning complexities, making its scores competitive with DeBERTaV3 in this sub-task, despite the higher



System	Original	Semantic	Context	Ori. + Sem.	Ori. + Sem. + Con.	Overall
<b>Multi-class classification problem</b>						
Human	.907	.907	.944	.907	.889	.920
ChatGPT	.608	.593	.679	.507	.397	.627
RoBERTa-L	.435	.402	.464	.330	.201	.434
<b>Mistral-7b_128_128</b>	<b>.850</b>	<b>.825</b>	<b>.775</b>	<b>.825</b>	<b>.700</b>	<b>.817</b>
Mistral-7b_64_128	.850	.825	.775	.825	.700	.817
Mistral-7b_16_64	.800	.800	.850	.750	.725	.817
Mixtral-8x7b_128_128	.850	.825	.725	.800	.700	.800
Llama 2-7b_64_128	.725	.650	.700	.575	.475	.692
Llama 2-13b_64_64	.665	.614	.645	.550	.400	.641
Llama 2-7b_64_64	.625	.600	.675	.550	.400	.633
Llama 2-7b_64_32	.250	.250	.425	.075	.000	.308
Phi-2_64_128	.625	.575	.550	.525	.425	.583
Phi-2_128_128	.625	.575	.550	.500	.375	.583
Phi-2_64_64	.525	.425	.550	.375	.300	.500
RoBERTa-WNGRD	.800	.775	.775	.750	.675	.784
DeBERTaV3-TS	.800	.775	.725	.750	.625	.767
DeBERTaV3-base	.725	.750	.675	.725	.625	.717
BERT-SE	.750	.725	.650	.700	.550	.708
RoBERTa-large	.700	.700	.725	.675	.550	.708
BERT	.675	.650	.650	.600	.475	.658
<b>Binary classification problem</b>						
DeBERTaV3-TS	.725	.650	.550	.650	.650	.642
RoBERTa-WNGRD	.575	.600	.500	.550	.550	.558
BERT-SE	.625	.550	.375	.525	.525	.517

Table 4: Model Performance for *sub-task A: Sentence Puzzle*. More results in Table 7.

System	Original	Semantic	Context	Ori.+Sem.	Ori.+Sem.+Con.	Overall
<b>Multi-class classification problem</b>						
Human	.917	.917	.917	.917	.900	.917
ChatGPT	.561	.524	.518	.439	.292	.535
RoBERTa-L	.195	.195	.232	.146	.061	.207
<b>Mistral-7b_16_64</b>	<b>.875</b>	<b>.906</b>	<b>.781</b>	<b>.813</b>	<b>.719</b>	<b>.854</b>
Mistral-7b_128_128	.844	.844	.813	.719	.625	.833
Mistral-7b_8_16	.781	.938	.781	.719	.562	.833
Mixtral-8x7b_128_128	.625	.719	.625	.531	.375	.656
Llama 2-13b_64_64	.354	.344	.438	.125	.031	.379
Llama 2-7b_64_64	.375	.344	.375	.125	.031	.365
Llama 2-7b_64_128	.281	.188	.438	.031	.031	.302
Phi-2_64_64	.688	.625	.688	.562	.438	.667
Phi-2_64_128	.656	.656	.625	.594	.406	.646
Phi-2_16_64	.625	.500	.688	.438	.312	.604
DeBERTaV3-base	.750	.750	.562	.656	.438	.687
DeBERTaV3-TS	.812	.781	.406	.719	.281	.666
RoBERTa-WNGRD	.750	.656	.500	.625	.312	.635
BERT	.562	.594	.469	.562	.312	.542
BERT-SE	.562	.500	.406	.500	.281	.489
RoBERTa-large	.281	.344	.344	.156	.094	.323
<b>Binary classification problem</b>						
DeBERTaV3-TS	.250	.313	.313	.219	.219	.292
RoBERTa-WNGRD	.375	.313	.125	.281	.281	.271
BERT-SE	.000	.000	.000	.000	.000	.000

Table 5: Model performance for *sub-task B: Word Puzzle*. More results in Table 7.

DeBERTaV3-base performance over RoBERTa-large in baseline reasoning benchmarks (He et al., 2023). Other than that, pre-training on other commonsense reasoning datasets does not significantly improve the overall performance for encoders. Conclusively, apart from WinoGrande the rest of the

extra pre-training datasets do not hold reasoning cues close to BrainTeaser’s word puzzles.

Regarding LLMs, Mistral-7b notably outperformed all others by a significant margin, even surpassing the 8 times larger model tuned using the same hyperparameters (*Mixtral-8x7b*). Llama 2

<i>Example of correct prediction with explanation</i>	
Question	Choice
If a dog is tied to a piece of rope that is 6m long how can it reach a bone that is 7m away?	<b>The other end is not tied to anything.</b>
	The rope stretches proportionally, providing the extra length needed for the horse to reach the hay seven meters away.
	The dog dug a 1m tunnel underground to reach the bone that was 7m away while still tied to the 6m long rope.
	None of above.
<b>Explanation:</b> A dog cannot reach a bone that is 7m away if it is tied to a piece of rope that is 6m long. This is because the distance between the dog and the bone is greater than the length of the rope. Therefore, the dog will not be able to reach the bone unless the rope is untied.	
<i>Example of misprediction with explanation</i>	
Question	Choice
How can you drop a raw egg onto a concrete floor and not crack it?	You need to drop the egg gently.
	You need to drop the egg hard.
	<b>No problem, because concrete floors are very hard to crack.</b>
	None of above.
<b>Predicted Answer:</b> None of above.	
<b>Explanation:</b> There is no way to drop a raw egg onto a concrete floor and not crack it. Concrete is a very hard material and will easily break any object that falls on it. The only way to avoid this would be to use some sort of protective covering for the egg before dropping it.	

Table 6: An example of explanations generated by Mistral-7b, while addressing the multi-class classification problem. The explanations, regardless of their validity, are detailed and showcase a high level of comprehension and reasoning. Correct answers are indicated in bold, while mispredictions reference the predicted answer by the model.

exhibited the worst results regardless of size (7/13 billion) and LoRA hyperparameters ( $r$  and  $a$ ). Conversely, Phi-2 demonstrated relatively better performance, particularly considering its smaller parameter count (2.7 billion) compared to the other LLMs. However, both models performed worse compared to most fine-tuned encoders. This observation strongly confirms that word puzzles possess a distribution that diverges from the analytical commonsense reasoning required for sentence puzzles, entailing a unique set of cognitive demands.

Mistral-7b exhibits a trend where higher quality explanations were generated with higher values of lora rank  $r$ . However, the top-performing model showcased a configuration with  $r=16$  and  $a=64$ . The QLoRA method (Hu et al., 2021) explains why our top model has a rank of 16 instead of 128, contrary to common expectations (more details regarding QLoRA hyperparameters in App. C). Drawing from the widespread presence of low-rank structures, as highlighted by prior studies (Li et al., 2016, 2019; Grasedyck et al., 2013), we leverage the intrinsic low-rank structure in our problem, as emphasized in Hu et al. (2021). It is well-established that many tasks, particularly involving heavily overparametrized models, exhibit low-rank properties post-training (Oymak et al., 2019).

Overall, our systems demonstrate remarkably

high overall accuracy, being less than 10% lower than human performance and more than 30% greater than ChatGPT. This suggests our methods' proficiency in understanding and detecting word-play patterns, consistently addressing ambiguity irrespective of contextual and semantic variations in brain teasers. Upon reviewing the short explanations provided with each prediction (Table 6), we note thorough justifications even for incorrect answers. Errors typically adhere to specific word-play patterns across original, semantic, and context multiple-choice questions (details in App. E).

## 6 Conclusion

In this study, we systematically evaluate pre-trained and fine-tuned encoders, along with instruction-tuned Large Language Models (LLMs), against two multi-class classification sub-tasks within the "BRAINTEASER: A Novel Task Defying Common Sense". We achieve competitive performance in both sub-tasks, accompanied by a plethora of insights regarding the influence of leveraging in-domain data, the variability model scale and architecture introduce, as well as the examination of diverging reasoning paths. As future work, we will delve into further reasoning patterns LLMs tend to follow with regard to lateral thinking challenges.

## References

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- L. Floridi and Massimo Chiriatti. 2020. [Gpt-3: Its nature, scope, limits, and consequences](#). *Minds and Machines*, 30:681–694.
- Gaël Gendron, Michael Witbrock, and Gillian Dobbie. 2024. [A survey of methods, challenges and perspectives in causality](#).
- Panagiotis Giadikiaroglou, Maria Lymperaioi, Giorgos Filandrianos, and Giorgos Stamou. 2024. [Puzzle solving using reasoning of large language models: A survey](#).
- Lars Grasedyck, Daniel Kressner, and Christine Tobler. 2013. [A literature survey of low-rank tensor approximation techniques](#).
- Sylvain Gugger, Lysandre Debut, Thomas Wolf, Philipp Schmid, Zachary Mueller, Sourab Mangrulkar, Marc Sun, and Benjamin Bossan. 2022. [Accelerate: Training and inference at scale made simple, efficient and adaptable](#). <https://github.com/huggingface/accelerate>.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#).
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#).
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023a. [Mistral 7b](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2024a. [Mixtral of experts](#).
- Yifan Jiang, Filip Ilievski, and Kaixin Ma. 2024b. [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. 2023b. [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.
- Carina Kauf, Anna A. Ivanova, Giulia Rambelli, Emanuele Chersoni, Jingyuan S. She, Zawad Chowdhury, Evelina Fedorenko, and Alessandro Lenci. 2022. [Event knowledge in large language models: the gap between the impossible and the unlikely](#). *ArXiv*, abs/2212.01488.
- Yuanzhi Li, Yingyu Liang, and Andrej Risteski. 2016. [Recovery guarantee of weighted low-rank approximation via alternating minimization](#).
- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. 2019. [Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations](#).
- Bill Yuchen Lin, Ziyi Wu, Yichi Yang, Dong-Ho Lee, and Xiang Ren. 2021. [Riddlesense: Reasoning about riddle questions featuring linguistic creativity and commonsense knowledge](#).
- 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](#).

- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. 2023. [A survey of deep learning for mathematical reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14605–14631, Toronto, Canada. Association for Computational Linguistics.
- Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. [Peft: State-of-the-art parameter-efficient fine-tuning methods](https://github.com/huggingface/peft). <https://github.com/huggingface/peft>.
- Samet Oymak, Zalan Fabian, Mingchen Li, and Mahdi Soltanolkotabi. 2019. [Generalization guarantees for neural networks via harnessing the low-rank structure of the jacobian](#).
- Molly R. Petersen and Lonneke van der Plas. 2023. [Can language models learn analogical reasoning? investigating training objectives and comparisons to human performance](#).
- Christopher Richardson and Larry Heck. 2023. [Commonsense reasoning for conversational ai: A survey of the state of the art](#).
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Winogrande: An adversarial winograd schema challenge at scale](#).
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Huai hsin Chi, Nathanael Scharli, and Denny Zhou. 2023. [Large language models can be easily distracted by irrelevant context](#). *ArXiv*, abs/2302.00093.
- Damien Sileo. 2023. [tasksource: Structured dataset preprocessing annotations for frictionless extreme multi-task learning and evaluation](#). *arXiv preprint arXiv:2301.05948*.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askell, Amanda Dsouza, Ambrose Slone, Ameet Rahane, Anantharaman S. Iyer, Anders Andreassen, Andrea Madotto, Andrea Santilli, Andreas Stuhlmüller, Andrew Dai, Andrew La, Andrew Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash Gholamidavoodi, Arfa Tabasum, Arul Menezes, Arun Kirubarajan, Asher Mullokandov, Ashish Sabharwal, Austin Herrick, Avia Efrat, Aykut Erdem, Ayla Karakaş, B. Ryan Roberts, Bao Sheng Loe, Barret Zoph, Bartłomiej Bojanowski, Batuhan Özyurt, Behnam Hedayatnia, Behnam Neyshabur, Benjamin Inden, Benno Stein, Berk Ekmekci, Bill Yuchen Lin, Blake Howald, Bryan Orinion, Cameron Diao, Cameron Dour, Catherine Stinson, Cedrick Argueta, César Ferri Ramírez, Chandan Singh, Charles Rathkopf, Chenlin Meng, Chitta Baral, Chiyu Wu, Chris Callison-Burch, Chris Waites, Christian Voigt, Christopher D. Manning, Christopher Potts, Cindy Ramirez, Clara E. Rivera, Clemencia Siro, Colin Raffel, Courtney Ashcraft, Cristina Garbacea, Damien Sileo, Dan Garrette, Dan Hendrycks, Dan Kilman, Dan Roth, Daniel Freeman, Daniel Khashabi, Daniel Levy, Daniel Moseguí González, Danielle Perszyk, Danny Hernandez, Danqi Chen, Daphne Ippolito, Dar Gilboa, David Dohan, David Drakard, David Jurgens, Debajyoti Datta, Deep Ganguli, Denis Emelin, Denis Kleyko, Deniz Yuret, Derek Chen, Derek Tam, Dieuwke Hupkes, Diganta Misra, Dilyar Buzan, Dimitri Coelho Mollo, Diyi Yang, Dong-Ho Lee, Dylan Schrader, Ekaterina Shutova, Ekin Dogus Cubuk, Elad Segal, Eleanor Hagerman, Elizabeth Barnes, Elizabeth Donoway, Ellie Pavlick, Emanuele Rodola, Emma Lam, Eric Chu, Eric Tang, Erkut Erdem, Ernie Chang, Ethan A. Chi, Ethan Dyer, Ethan Jerzak, Ethan Kim, Eunice Engefu Manyasi, Evgenii Zheltonozhskii, Fanyue Xia, Fatemeh Siar, Fernando Martínez-Plumed, Francesca Happé, Francois Chollet, Frieda Rong, Gaurav Mishra, Genta Indra Winata, Gerard de Melo, Germán Kruszewski, Giambattista Parascandolo, Giorgio Mariani, Gloria Wang, Gonzalo Jaimovitch-López, Gregor Betz, Guy Gur-Ari, Hana Galijasevic, Hannah Kim, Hannah Rashkin, Hannaneh Hajishirzi, Harsh Mehta, Hayden Bogar, Henry Shevlin, Hinrich Schütze, Hiromu Yakura, Hongming Zhang, Hugh Mee Wong, Ian Ng, Isaac Noble, Jaap Jumelet, Jack Geissinger, Jackson Kernion, Jacob Hilton, Jaehoon Lee, Jaime Fernández Fisac, James B. Simon, James Koppel, James Zheng, James Zou, Jan Kocoń, Jana Thompson, Janelle Wingfield, Jared Kaplan, Jarema Radom, Jascha Sohl-Dickstein, Jason Phang, Jason Wei, Jason Yosinski, Jekaterina Novikova, Jelle Bosscher, Jennifer Marsh, Jeremy Kim, Jeroen Taal, Jesse Engel, Jesujoba Alabi, Jiacheng Xu, Ji-aming Song, Jillian Tang, Joan Waweru, John Burden, John Miller, John U. Balis, Jonathan Batchelder, Jonathan Berant, Jörg Frohberg, Jos Rozen, Jose Hernandez-Orallo, Joseph Boudeman, Joseph Guerr, Joseph Jones, Joshua B. Tenenbaum, Joshua S. Rule, Joyce Chua, Kamil Kanclerz, Karen Livescu, Karl Krauth, Karthik Gopalakrishnan, Katerina Ignatyeva, Katja Markert, Kaustubh D. Dhole, Kevin Gimpel, Kevin Omondi, Kory Mathewson, Kristen Chifullo, Ksenia Shkaruta, Kumar Shridhar, Kyle McDonell, Kyle Richardson, Laria Reynolds, Leo Gao, Li Zhang, Liam Dugan, Lianhui Qin, Lidia Contreras-Ochando, Louis-Philippe Morency, Luca Moschella, Lucas Lam, Lucy Noble, Ludwig Schmidt, Luheng He, Luis Oliveros Colón, Luke Metz, Lütfi Kerem Şenel, Maarten Bosma, Maarten Sap, Maartje ter Hoeve, Maheen Farooqi, Manaal Faruqui, Mantas Mazeika, Marco Baturan, Marco Marelli, Marco Maru, Maria Jose Ramírez Quintana, Marie Tolkiehn, Mario Giulianelli, Martha Lewis, Martin Potthast, Matthew L. Leavitt, Matthias Hagen, Mátyás Schubert, Medina Orduna Baitemirova, Melody Arnaud, Melvin McElrath, Michael A. Yee, Michael Co-



hen, Michael Gu, Michael Ivanitskiy, Michael Starritt, Michael Strube, Michał Śwędrowski, Michele Bevilacqua, Michihiro Yasunaga, Mihir Kale, Mike Cain, Mimeo Xu, Mirac Suzgun, Mitch Walker, Mo Tiwari, Mohit Bansal, Moin Amninsari, Mor Geva, Mozhdah Gheini, Mukund Varma T, Nanyun Peng, Nathan A. Chi, Nayeon Lee, Neta Gur-Ari Krakover, Nicholas Cameron, Nicholas Roberts, Nick Doiron, Nicole Martinez, Nikita Nangia, Niklas Deckers, Niklas Muennighoff, Nitish Shirish Keskar, Niveditha S. Iyer, Noah Constant, Noah Fiedel, Nuan Wen, Oliver Zhang, Omar Agha, Omar Elbaghdadi, Omer Levy, Owain Evans, Pablo Antonio Moreno Casares, Parth Doshi, Pascale Fung, Paul Pu Liang, Paul Vicol, Pegah Alipoormolabashi, Peiyuan Liao, Percy Liang, Peter Chang, Peter Eckersley, Phu Mon Htut, Pinyu Hwang, Piotr Miłkowski, Piyush Patil, Pouya Pezeshkpour, Priti Oli, Qiaozhu Mei, Qing Lyu, Qinlang Chen, Rabin Banjade, Rachel Etta Rudolph, Raefer Gabriel, Rahel Habacker, Ramon Risco, Raphaël Millière, Rhythm Garg, Richard Barnes, Rif A. Saurous, Riku Arakawa, Robbe Raymaekers, Robert Frank, Rohan Sikand, Roman Novak, Roman Sitelew, Ronan LeBras, Rosanne Liu, Rowan Jacobs, Rui Zhang, Ruslan Salakhutdinov, Ryan Chi, Ryan Lee, Ryan Stovall, Ryan Teehan, Rylan Yang, Sahib Singh, Saif M. Mohammad, Sajant Anand, Sam Dillavou, Sam Shleifer, Sam Wiseman, Samuel Gruetter, Samuel R. Bowman, Samuel S. Schoenholz, Sanghyun Han, Sanjeev Kwatra, Sarah A. Rous, Sarik Ghazarian, Sayan Ghosh, Sean Casey, Sebastian Bischoff, Sebastian Gehrmann, Sebastian Schuster, Sepideh Sadeghi, Shadi Hamdan, Sharon Zhou, Shashank Srivastava, Sherry Shi, Shikhar Singh, Shima Asaadi, Shixiang Shane Gu, Shubh Pachchigar, Shubham Toshniwal, Shyam Upadhyay, Shyamolima, Debnath, Siamak Shakeri, Simon Thormeyer, Simone Melzi, Siva Reddy, Sneha Priscilla Makini, Soo-Hwan Lee, Spencer Torene, Sriharsha Hatwar, Stanislas Dehaene, Stefan Divic, Stefano Ermon, Stella Biderman, Stephanie Lin, Stephen Prasad, Steven T. Piantadosi, Stuart M. Shieber, Summer Mishner, Svetlana Kiritchenko, Swaroop Mishra, Tal Linzen, Tal Schuster, Tao Li, Tao Yu, Tariq Ali, Tatsu Hashimoto, Te-Lin Wu, Théo Desbordes, Theodore Rothschild, Thomas Phan, Tianle Wang, Tiberius Nkinyili, Timo Schick, Timofei Kornev, Titus Tunduny, Tobias Gerstenberg, Trenton Chang, Trishala Neeraj, Tushar Khot, Tyler Shultz, Uri Shaham, Vedant Misra, Vera Demberg, Victoria Nyamai, Vikas Raunak, Vinay Ramasesh, Vinay Uday Prabhu, Vishakh Padmakumar, Vivek Srikumar, William Fedus, William Saunders, William Zhang, Wout Vossen, Xiang Ren, Xiaoyu Tong, Xinran Zhao, Xinyi Wu, Xudong Shen, Yadollah Yaghoobzadeh, Yair Lakretz, Yangqiu Song, Yasaman Bahri, Yejin Choi, Yichi Yang, Yiding Hao, Yifu Chen, Yonatan Belinkov, Yu Hou, Yufang Hou, Yuntao Bai, Zachary Seid, Zhuoye Zhao, Zijian Wang, Zijie J. Wang, Zirui Wang, and Ziyi Wu. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models.](#)

Jiankai Sun, Chuanyang Zheng, Enze Xie, Zhengying

Liu, Ruihang Chu, Jianing Qiu, Jiaqi Xu, Mingyu Ding, Hongyang Li, Mengzhe Geng, Yue Wu, Wenhai Wang, Junsong Chen, Zhangyue Yin, Xiaozhe Ren, Jie Fu, Junxian He, Wu Yuan, Qi Liu, Xihui Liu, Yu Li, Hao Dong, Yu Cheng, Ming Zhang, Pheng Ann Heng, Jifeng Dai, Ping Luo, Jingdong Wang, Ji-Rong Wen, Xipeng Qiu, Yike Guo, Hui Xiong, Qun Liu, and Zhenguo Li. 2023. [A survey of reasoning with foundation models.](#)

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, Minneapolis, Minnesota. Association for Computational Linguistics.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models.](#)

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models.](#)

Gladys Tyen, Hassan Mansoor, Victor Cărbune, Peter Chen, and Tony Mak. 2024. [LLMs cannot find reasoning errors, but can correct them!](#)

Siddharth Vashishtha, Adam Poliak, Yash Kumar Lal, Benjamin Van Durme, and Aaron Steven White. 2020. [Temporal reasoning in natural language inference.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4070–4078, Online. Association for Computational Linguistics.

Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan

Lambert, and Shengyi Huang. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.

Cunxiang Wang, Shuailong Liang, Yili Jin, Yilong Wang, Xiaodan Zhu, and Yue Zhang. 2020. SemEval-2020 task 4: Commonsense validation and explanation. In *Proceedings of The 14th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.

Jingxuan Wei, Cheng Tan, Zhangyang Gao, Linzhuang Sun, Siyuan Li, Bihui Yu, Ruifeng Guo, and Stan Z. Li. 2023. [Enhancing human-like multi-modal reasoning: A new challenging dataset and comprehensive framework](#).

Zonglin Yang, Xinya Du, Rui Mao, Jinjie Ni, and Erik Cambria. 2023. [Logical reasoning over natural language as knowledge representation: A survey](#).

Zhangdie Yuan, Songbo Hu, Ivan Vulic, Anna Korhonen, and Zaiqiao Meng. 2022. Can pretrained language models (yet) reason deductively? *ArXiv*, abs/2210.06442.

Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. 2023. [How language model hallucinations can snowball](#).

## A Model Selection

### A.1 Encoder

BERT (Devlin et al., 2019): Bidirectional Encoder Representations for Transformers, is a pretrained deep bidirectional transformer model producing context representations. Using a fine-tuning setting, BERT has advanced state-of-the-art performances on a wide range of NLP tasks.

RoBERTa-large (Liu et al., 2019): Robustly Optimized BERT pre-training Approach (RoBERTa) is an adaptation of BERT architecture trained with larger batches on 160 GB data from various domains. RoBERTa-large was trained by dynamically modifying language masking while the next sentence prediction loss used in BERT was dropped. Other improvising techniques like larger input text sequences, byte pair encoding are used in training which seemingly improved the model performance in downstream tasks.

DeBERTaV3 (He et al., 2023): Decoding-enhanced BERT with disentangled attention is an extension of the original DeBERTa model. It builds

upon the BERT (Bidirectional Encoder Representations from Transformers) architecture, aiming to enhance its decoding capabilities and overall performance across various natural language processing (NLP) tasks. DeBERTaV3 further improves the efficiency of DeBERTa (He et al., 2021) using ELECTRA-Style pre-training with Gradient Disentangled Embedding Sharing. Compared to DeBERTa, V3 significantly improves the model performance on downstream tasks. It incorporates a disentangled attention mechanism to allow the model to focus on different aspects of input independently, improving its ability to capture diverse linguistic patterns. The model also features enhancements in the decoding process, enabling more accurate text generation and sequence classification.

### A.2 LLMs

Mistral-7b (Jiang et al., 2023a): Developed by EleutherAI, is a language model tailored for large-scale natural language processing tasks. With its 7 billion parameters, it excels in handling complex language understanding and generation tasks. Designed to perform exceptionally well across various NLP applications such as text generation, comprehension, and summarization, Mistral-7b surpasses the best open 13b model, Llama 2 (Touvron et al., 2023b), and the best released 34b model, Llama 1 (Touvron et al., 2023a), in reasoning, mathematics, and code generation tasks. Leveraging grouped-query attention (GQA) and sliding window attention (SWA), Mistral-7b ensures efficient inference and can handle sequences of arbitrary length with reduced inference cost. Its performance across a wide range of benchmarks makes it a promising solution for our sub-tasks, given its extensive task capabilities and superior performance in baseline benchmarks compared to similar or larger language models. While we considered experimenting with its larger variant, Mixtral-8x7b (Jiang et al., 2024a), limitations on available resources forced us to deal in depth only with the small variant, Mistral-7b.

Llama 2 (Touvron et al., 2023b): A language model that represents a significant advancement in natural language processing. It is a collection of pre-trained and fine-tuned large language models (LLMs) ranging in scale from 7 billion to 70 billion parameters. With its large parameter count and advanced architecture, Llama 2 is designed to tackle complex language understanding and generation tasks effectively. It outperforms many other mod-

els, including its predecessor, Llama 1, in various benchmarks, demonstrating superior capabilities in reasoning, mathematics, and code generation. Leveraging its extensive parameterization and innovative techniques, Llama 2 offers state-of-the-art performance across a wide range of NLP applications, making it a notable contender in the field. For our experiments we were able to experiment with various configurations with the 7 billion and the 13 billion models. Our involvement with the 70 billion parameter model has been restricted due to limitations associated with the extensive parameter count, particularly during the fine-tuning process.

**Phi-2** (Gunasekar et al., 2023): An advanced language model designed to address complex natural language processing tasks efficiently. It is part of the small language models (SLMs) released by Microsoft Research team. With its innovative architecture and extensive parameter count, Phi-2 surpasses its predecessor, Phi-1, in various benchmarks, showcasing superior performance in reasoning, comprehension, and text generation. Leveraging cutting-edge techniques and a comprehensive understanding of language patterns, Phi-2 demonstrates remarkable capabilities across a diverse range of NLP applications, solidifying its position as a prominent model in the field. Given its 2.7 billion-parameter architecture, which exhibits exceptional reasoning and language understanding abilities in comparison to various Llama 2 iterations and Mistral-7b, we are confident that this model will deliver noticeable performance for both of our sub-tasks.

## B Experimental Setup

In our experiments, we employed the Google Colab platform and Kaggle, leveraging various open-source Python packages such as Transformers, TRL (Transformer Reinforcement Learning) (von Werra et al., 2020), PEFT (Parameter-Efficient Fine-Tuning) (Mangrulkar et al., 2022), BitsAndBytes, Accelerate (Gugger et al., 2022), and SentenceTransformers.

**Encoders** BERT-SE<sup>2</sup>: During fine-tuning, a learning rate of  $3e^{-5}$  was used, with a batch size of 16 samples processed in each iteration, over the course of 3 epochs. This process aimed to adapt the pre-trained model to better suit our sub-task. Our optimizer was AdamW and our learning scheduler

<sup>2</sup><https://huggingface.co/JazibEijaz/bert-base-uncased-finetuned-semantic-similarity-2020-task4b-append-e3-b32-l4e5>

was linear. Same setup was used for the fine-tuning of the BERT encoder.

RoBERTa-WNGRD<sup>3</sup> underwent fine-tuning on the train split of each dataset, utilizing a learning rate of  $3e^{-5}$ , a batch size of 16, and running for 3 epochs. The optimizer was also AdamW and the learning scheduler was linear. RoBERTa-large was fine-tuned on the train split of each sub-task’s specific dataset using identical configurations.

DeBERTaV3-TS<sup>4</sup>, like DeBERTaV3-base, underwent a fine-tuning process similar to the RoBERTa-WNGRD system, differing only in the batch size, which was set to 4.

**LLMs** Phi-2<sup>5</sup> underwent fine-tuning using the prompt format outlined in Section *Prompting Details*. The fine-tuning process involved setting a learning rate of  $2e^{-5}$  and a batch size of 2, with the model trained for 250 steps. We conducted experiments with different configurations of  $r$  and  $\text{lora\_alpha}$ , encompassing combinations such as  $r = 64, 128$  and  $\text{lora\_alpha} = 64, 128$ . The dropout rate was consistently set to 0.1 across all experiments. We used an AdamW optimizer and a constant learning scheduler. Despite promising benchmarks accompanying its release, the model’s performance during inference on the test split of both sub-tasks’ datasets was subpar, scoring lower compared to the encoders mentioned above. This discrepancy raises the possibility, supported by various reports, that the model’s training process using methods like quantization and LoRA may not be fully optimized yet, particularly given its recent introduction.

Both variations of Llama 2<sup>6</sup>, with 7 billion and 13 billion parameters, underwent the same fine-tuning pipeline described earlier, utilizing the QLoRA technique. The fine-tuning process followed the prompt format outlined in Section D (*Prompting Details*), employing a learning rate of  $2e^{-5}$  and a batch size of 1, with each model trained for 250 steps. Despite experimenting with various combinations of values for  $r$  and  $a$  (32, 64, 128), while the dropout rate was consistently set to 0.1, the results were disappointing. As a text generation model, Llama 2 provided explanations for each multiple-choice prompt. However, even when

<sup>3</sup><https://huggingface.co/DeepPavlov/roberta-large-winogrande>

<sup>4</sup><https://huggingface.co/sileod/deberta-v3-large-tasksource-nli>

<sup>5</sup><https://huggingface.co/microsoft/phi-2>

<sup>6</sup>[https://huggingface.co/docs/transformers/en/model\\_doc/llama2](https://huggingface.co/docs/transformers/en/model_doc/llama2)



incorrectly predicting a choice as correct, the generated explanations often lacked logical coherence. Many explanations produced during the inference phase were irrelevant to the context of the brain teaser, indicating a failure to capture the reasoning path of most multiple-choice questions. In summary, both variations of Llama 2, despite their large scale, proved incapable of effectively understanding and reasoning through the multiple-choice questions provided.

The Mistral-7b<sup>7</sup> model outperformed all others significantly. Prior to fine-tuning, we applied the QLoRA technique. Using a learning rate of  $2e^{-5}$  and a batch size of 2, each model underwent fine-tuning for 250 steps using the train split of the sub-tasks' dataset. The initial results were promising. During experimentation with the  $r$  and  $a$  parameters, while maintaining a dropout of 0.1, certain patterns emerged. Specifically, we observed higher quality explanations and scores when using higher rank values, ranging from (16, 32, 64, 128). This outcome was expected, as higher rank values correspond to higher precision weight changes, resulting in superior weight tuning and overall model performance. Interestingly, when the ratio of  $a/r$  was low (0.5 - 1), explanations maintained high quality irrespective of predictions, implying a coherent reasoning path even if the predicted choice was incorrect. However, setting the  $a/r$  ratio to 2 or 4 potentially enhanced results, signifying a stronger influence from QLoRA layers on the base model. However, this adjustment led to a decline in the quality of explanations. The improvement could be attributed to the model's low intrinsic dimensionality. Despite having many parameters, the effective dimensionality of the model's learned representations is low. Consequently, after conducting several experiments, the best-performing model regarding word puzzles aligns with this concept. After conducting numerous tests, we achieved our best performances with the first model using  $r=128$  and  $\alpha=128$ , and the third best using  $r=64$  and  $\alpha=32$ . These models are denoted as Mistral-7b\_lora\_r\_lora\_a, representing Mistral-7b\_128\_128 and Mistral-7b\_64\_32 configurations, respectively.

Our exploration of Mistral-8x7b<sup>8</sup> was constrained, yet initial results were promising, despite

<sup>7</sup><https://huggingface.co/mistralai/Mistral-7B-v0.1>

<sup>8</sup><https://huggingface.co/mistralai/Mistral-8x7B-v0.1>

the limited configurations. Further experimentation with various hyperparameter settings may yield improved performance. In our single attempt with this system, we employed a learning rate of  $2e^{-5}$  and a batch size of 2, fine-tuning the models for 250 steps using the train split of the sub-task's dataset. Both  $r$  and  $a$  were set to 128, accompanied by a dropout rate of 0.1. This configuration was selected based on the  $r$  and  $a$  values of the best-performing model across both sub-tasks, Mistral-7b. Despite its larger scale, Mistral-8x7b achieved the second-best accuracy during inference on the test split regarding the first subtask, trailing behind its smaller variation, Mistral-7b. This model is referenced in the results table of both sub-tasks as Mistral-8x7b\_128\_128. Further experimentation with various configurations may yield improvements, particularly when leveraging the low intrinsic dimensionality and redundancy inherent in the model.

## C QLoRA hyperparameters

Initially, we employed the QLoRA technique (Dettmers et al., 2023) for optimization. The QLoRA technique entails the following steps. First we quantized the models using 4-bit precision to reduce memory usage and computational requirements. The quantization process was facilitated by the BitsAndBytes library. Following quantization, we implemented the LoRA technique (Hu et al., 2021) using the PEFT library. LoRA, applied to the quantized model, resulted in the creation of Quantized LoRA (QLoRA). This pipeline effectively addresses the challenges posed by memory-intensive models on hardware with limited capabilities, ensuring optimized performance and resource utilization. Regarding the hyperparameters of the QLoRA, the rank ( $r$ ) determines the dimensionality of the low-rank approximation used in the adapter layers, while alpha ( $a$ ) is the scaling factor that determines the magnitude of the newly learned weights compared to the original model's weights. The choice of alpha influences how much emphasis is given to the task-specific information compared to the pre-trained knowledge encoded in the original model.

In our experiments, we observed that lower values of  $r$  occasionally yielded slightly superior results. This phenomenon can be attributed to the regularization effect introduced by lower-rank approximations. Essentially, lower-rank approximations act as a form of regularization, discouraging



the model from memorizing the training data and instead promoting the learning of more generalizable patterns. This regularization effect becomes particularly significant when dealing with small datasets, as the risk of overfitting is heightened in such scenarios. By limiting the model’s capacity through lower-rank approximations, we encourage it to focus on learning essential features and avoid capturing noise or idiosyncrasies present in the training data. Therefore, in our case where the dataset size is small, the regularization provided by lower-rank approximations becomes crucial. It helps prevent overfitting and encourages the model to generalize better to unseen data, ultimately leading to improved performance in certain cases.

Table 7 depicts further analysis of LoRA hyperparameters for Mistral and Mixtral models, which have exhibited the best results among all other models and across the two tasks. Due to computational restrictions, we trained the Mixtral model, which is eight times larger, only for the best performing hyperparameters of Mistral, as a proxy for the performance difference.

## D Prompting Details

Here, we provide a comprehensive overview of the prompt utilized consistently throughout the fine-tuning process of the LLMs, which ultimately led to optimal performance across both sub-tasks. Prompt:

```
### Instructions:
Below is an instruction that describes a multiple choice task. Answer the following multiple choice question by giving the most appropriate response. Answer should be one among options provided after the question. Select the most suitable answer while making the necessary assumptions. Give only answer and a short explanation of two or three sentences. Nothing else.
```

```
### Input:
Question: {question}
1) {a}
2) {b}
3) {c}
4) {d}
```

```
### Answer:
The correct answer is: {label}) {answer}
In the Instructions section, we define the task and
```

provide detailed steps for the system. Results varied depending on the content of the *Instructions* section. It’s important to note that our model isn’t just tasked with selecting the most appropriate choice from the given options; it’s also instructed to generate a brief explanation. This additional step aims to assess the model’s ability to identify and comprehend a logical reasoning path that can justify its chosen answers for each multiple-choice problem. Given that the questions are brain teasers that challenge common sense, this approach helps us gauge the model’s understanding and reasoning capabilities more effectively. In the *Input* section, we structure the provided dataset into a multiple-choice question format. Each component serves a specific purpose:

**Question {question}** This section contains the main question extracted from the dataset.

**Choices ({a}, {b}, {c}, {d}):** These represent the options provided as answers for the question within the dataset.

**Correct Answer {label}) {answer}** This section indicates the correct label and its corresponding answer from the dataset.

This structured format enables the model to comprehend and process each question along with its associated choices and correct answer during the fine-tuning training process. During the *inference phase*, the same prompt is reproduced, with the sole distinction of a blank space within the Answer section. This deliberate inclusion of a blank space aims to support the model’s text generation process. In inference, the model is tasked with generating the correct answer using the information presented in the prompt. This setup enables the model to dynamically generate responses, utilizing its comprehension of the question and the contextual details provided within the prompt.

## E Assessment and Insights on Dataset Quality

Upon reviewing our incorrect predictions across both sub-tasks, subsequent to the task organizer releasing the labels for the test split of the datasets, we reached several conclusions. Across all triplets, encompassing original, semantic, and context reconstruction statements, we observe a considerable degree of ambiguity in various patterns. This ambiguity often leads to inconsistent selection of correct answers, even when answered by humans. This underscores the need for clearer formulation of ques-

System	Original	Semantic	Context	Ori. + Sem.	Ori. + Sem. + Con.	Overall
<b>Task A</b>						
Mistral-7b_64_128	<b>.850</b>	<b>.825</b>	<b>.775</b>	<b>.825</b>	<b>.700</b>	<b>.817</b>
Mistral-7b_16_64	.800	.800	.850	.750	.725	.817
Mixtral-8x7b_128_128	.850	.825	.725	.800	.700	.800
Mistral-7b_128_64	.850	.800	.725	.775	.625	.792
Mistral-7b_64_32	.850	.775	.725	.750	.675	.783
Mistral-7b_8_16	.800	.800	.700	.750	.625	.767
Mistral-7b_128_32	.825	.775	.725	.750	.600	.775
<b>Task B</b>						
Mistral-7b_128_128	<b>.844</b>	<b>.844</b>	<b>.813</b>	<b>.719</b>	<b>.625</b>	<b>.833</b>
Mistral-7b_8_16	.781	.938	.781	.719	.562	.833
Mistral-7b_16_16	.812	.812	.875	.688	.625	.833
Mistral-7b_8_8	.875	.812	.812	.750	.688	.833
Mistral-7b_16_32	.875	.812	.781	.750	.594	.823
Mistral-7b_64_32	.844	.875	.719	.750	.562	.812
Mistral-7b_128_64	.844	.812	.781	.688	.531	.812
Mistral-7b_64_64	.719	.812	.625	.625	.406	.719
Mixtral-8x7b_128_128	.625	.719	.625	.531	.375	.656

Table 7: The performance of various LoRA hyperparameters for Mistral and Mixtral in both sub-tasks.

tions and unambiguous expression to enhance the accuracy of model predictions. Another notable pattern we identified pertains to the quality control of semantic reconstruction in certain questions. In these instances, some words were not replaced with accurate synonyms, resulting in a shift in the definition of the brain teaser presented by the question. While this may not inherently be problematic, the dataset’s correct answers remained unchanged compared to the original version of the question. This discrepancy suggests that the alteration in question definition went unnoticed by the task organizers, leading to some erroneous predictions by our model, when in reality the correct context of the provided multiple-choice statement was captured by our system. The two observations above highlight the inherent difficulty in generating clear and precise brain teasers, as well as the challenge that models face in understanding them. In the above scenarios, our top-performing model either detects the presence of a contradiction in the questions and opts to select "None of above," as elucidated in its brief and explanatory justification, or it provides an incorrect answer based on the dataset’s answer but correctly reflects the problem context, which may have been altered due to inadvertent synonym usage.