

D-NLP at SemEval-2024 Task 2: Evaluating Clinical Inference Capabilities of Large Language Models

Duygu Altinok

Deepgram Research, USA
duygu.altinok@deepgram.com

Abstract

Large language models (LLMs) have garnered significant attention and widespread usage due to their impressive performance in various tasks. However, they are not without their own set of challenges, including issues such as hallucinations, factual inconsistencies, and limitations in numerical-quantitative reasoning. Evaluating LLMs in miscellaneous reasoning tasks remains an active area of research. Prior to the breakthrough of LLMs, Transformers had already proven successful in the medical domain, effectively employed for various natural language understanding (NLU) tasks. Following this trend, LLMs have also been trained and utilized in the medical domain, raising concerns regarding factual accuracy, adherence to safety protocols, and inherent limitations. In this paper, we focus on evaluating the natural language inference capabilities of popular open-source and closed-source LLMs using clinical trial reports as the dataset. We present the performance results of each LLM and further analyze their performance on a development set, particularly focusing on challenging instances that involve medical abbreviations and require numerical-quantitative reasoning. Gemini, our leading LLM, achieved a test set F1-score of 0.748, securing the ninth position on the task scoreboard. Our work is the first of its kind, offering a thorough examination of the inference capabilities of LLMs within the medical domain.

1 Introduction

Large language models (LLMs) have brought about a paradigm shift in the field of Natural Language Processing (NLP) (Kojima et al., 2023; Wei et al., 2022). Their exceptional performance across various tasks has led to a surge in real-world applications utilizing LLM-based technology. However, a notable drawback of LLMs is their propensity to generate plausible yet incorrect information,

commonly referred to as "hallucinations" (Ji et al., 2023).

The remarkable breakthrough of LLMs has raised questions regarding their "intelligent" capabilities, particularly in reasoning and inference (Zhao et al., 2023; Chang et al., 2023; Laskar et al., 2023). Two specific areas that have garnered significant attention in relation to LLMs' reasoning abilities are numerical-quantitative reasoning and natural language inference. These areas are considered integral to human intelligence, prompting researchers to establish benchmarks and evaluate LLM performance in these domains (Stolfo et al., 2023; Yuan et al., 2023). LLMs often exhibit limited performance in solving arithmetic reasoning tasks, frequently producing incorrect answers (Imani et al., 2023). Unlike natural language understanding, math problems typically possess a single correct solution, making the accurate generation of solutions more challenging for LLMs. Regarding NLI, performance reduction can be observed due to shortcut learning (Du et al., 2023) and hallucinations (McKenna et al., 2023). These investigations aim to discern whether LLMs are mere memorizers of training data or possess genuine logical reasoning abilities.

The volume of medical publications, including clinical trial data, has experienced a significant upsurge in recent years. The SemEval-2023 Task 7, known as Multi-Evidence Natural Language Inference for Clinical Trial Data (NLI4CT), aimed to address the challenge of large-scale interpretability and evidence retrieval from breast cancer clinical trial reports (Jullien et al., 2023). This task required multi-hop biomedical and numerical reasoning, which are crucial for developing systems capable of interpreting and retrieving medical evidence on a large scale, thereby facilitating personalized evidence-based care. While the previous iteration of NLI4CT resulted in the development of LLM-based models (Zhou et al., 2023;

Kanakarajan and Sankarasubbu, 2023; Vladika and Matthes, 2023) achieving high performance (e.g., F1-score \approx 85%), the application of LLMs in critical domains, such as real-world clinical trials, necessitates further investigation. Consequently, the second iteration of NLI4CT, SemEval-2024 Task 2, titled "Safe Biomedical Natural Language Inference for Clinical Trials" (Jullien et al., 2024) is proposed, featuring an enriched dataset that includes a novel contrast set obtained through interventions applied to statements in the NLI4CT test set. Our work involves the evaluation of various popular open-source and closed-source LLMs on the development and test sets to explore their reasoning capabilities in the domain of medical NLI. We present the results by thoroughly analyzing the performance on the development set, with the best-performing LLM ranking ninth on the task leaderboard. We have made the results on the development set available on our GitHub repository¹.

Another aspect of our work was that we deliberately refrained from investing significant effort into prompting or experimenting with different prompts. Additionally, we aimed to showcase the remarkable development of LLMs, demonstrating their capacity to effectively engage with the task while minimizing dependence on the prompt.

2 Related Work

With the emergence of large language models (LLMs), there has been a growing interest in exploring their capabilities within the clinical domain. Recent studies have delved into both the potential of LLMs and the associated risks when applied in clinical settings. For instance, (Hung et al., 2023) conducted experiments utilizing GPT-3.5 on various medical NLP datasets, assessing metrics such as factuality and safety, ultimately highlighting the high level of safety offered by GPT-3.5². (Pal et al., 2023) focused on the challenges posed by hallucinations in LLMs and proposed a benchmark dataset called Med-HALT (Medical Domain Hallucination Test) to evaluate popular LLMs on this front.

Regarding the reasoning capabilities of LLMs, (Kwon et al., 2024) introduced a diagnostic framework that prioritizes reasoning and employs prompt-based learning. The study specifically fo-

cused on clinical reasoning for disease diagnosis, where the LLMs generate diagnostic rationales to provide insights into patient data and the reasoning path leading to the diagnosis, known as Clinical Chain-of-Thought (Clinical CoT), using GPT-3.5 and GPT-4 (OpenAI, 2024). Notably, none of the previous studies simultaneously examined the performance of both open-source and closed-source LLMs, particularly with a comprehensive focus on inference. Consequently, our work stands as the first of its kind in this regard.

3 Task and Dataset Description

The clinical trials used to construct the dataset were sourced from ClinicalTrials.gov³, a comprehensive database managed by the U.S. National Library of Medicine. ClinicalTrials.gov contains information on various clinical studies conducted worldwide, both publicly and privately funded. The dataset specifically focuses on clinical trials related to breast cancer and includes a total of 1,000 trials written in English.

- **Eligibility Criteria:** This includes a set of conditions that determine the eligibility of patients to participate in the clinical trial. These criteria may include factors such as age, gender, and medical history.
- **Intervention:** This field provides information about the type, dosage, frequency, and duration of treatments being studied within the clinical trial.
- **Results:** The results section of each CTR reports the outcome of the trial, including data such as the number of participants, outcome measures, units of measurement, and the observed results.
- **Adverse Events:** This field documents any unwanted side effects, signs, or symptoms observed in patients during the course of the clinical trial.

For the task at hand, each CTR may contain one or two patient groups, known as cohorts or arms, which may receive different treatments or have different baseline characteristics.

The dataset consists of a total of 7,400 statements. These statements were divided into a training dataset comprising 1,700 statements, a development dataset containing 200 statements, and a

¹https://github.com/DuyguA/SemEval2024_NLI4CT

²<https://platform.openai.com/docs/models/gpt-3-5>

³<https://clinicaltrials.gov>

Model	Release Date	Params
GPT-3.5	Mar-2022	x
Claude	Mar-2023	x
Gemini Pro	Dec-2023	x
PaLM	Mar-2023	540B
Falcon 40B	May-2023	40B
Mixtral 8x7B	Dec-2023	12B
Llama 2 70B	Jul-2023	130GB

Table 1: Comparison of the LLMs used in our work, indicating the parameter sizes for known closed-source LLMs and denoting unknown parameter sizes with "x".

hidden test dataset consisting of 5,500 statements. The statements can be categorized into two types: those that are solely related to a single CTR and others that involve a comparison between two different reports. Each statement in the dataset is labeled as either "entailment" or "contradiction". Figure 1 shows an example statement from the training set.

The task primarily involves binary classification, aiming to predict whether the label corresponds to entailment or contradiction. The evaluation process encompasses three aspects. Initially, the macro F1-score is computed based on the binary classification results. Subsequently, two semantic evaluations are conducted: faithfulness and consistency. Faithfulness assesses the system’s ability to arrive at accurate predictions for the correct reasons, while consistency measures the system’s ability to produce consistent outputs for semantically equivalent problems. The task organizers evaluate faithfulness by providing semantically altered instances, and consistency by providing preserved instances for comparison.

4 Language Model Performance Evaluation

This section aims to provide a detailed analysis of the performance achieved by each individual LLM. Based on the evaluation of various LLMs, including closed-source models like GPT-3.5 (ChatGPT), Claude (Anthropic, 2023), and Gemini Pro (Gemini Team, 2023), as well as open-source models like Falcon 40B (Almazrouei et al., 2023), Mixtral 8x7B (Jiang et al., 2024), and Llama 2 70B (Touvron et al., 2023), the performance of these models was assessed on the dev and test sets. Table 1 provides comprehensive information regarding the release dates and parameter sizes, measured in token size, for each LLM.

Model	Acc	F1	Prec	Recall
Gemini Pro	0.82	0.81	0.82	0.8
Claude	0.81	0.80	0.81	0.81
PaLM	0.79	0.78	0.79	0.79
Falcon 40B	0.745	0.74	0.74	0.74
GPT-3.5	0.705	0.7	0.711	0.70
Llama 2 70B	0.675	0.67	0.68	0.67
Mixtral 8x7B	0.655	0.64	0.67	0.65

Table 2: Accuracy, macro F1-score, precision and recall results on the development set for each LLM.

Model	F1	Faith	Consist
Gemini Pro	0.75	0.83	0.74
Claude	0.73	0.83	0.72
PaLM	0.72	0.87	0.73
Falcon 40B	0.702	0.569	0.609
GPT-3.5	0.684	0.74	0.66
Llama 2 70B	0.682	0.693	0.638
Mixtral 8x7B	0.604	0.899	0.73

Table 3: Macro F1-score, faithfulness and consistency results on the test set for each LLM.

All conversations took place on the Poe.com platform, providing users with a seamless chat experience. To transmit both the development set and the test set instances, we utilized an API wrapper code in a Python script, which can be accessed in our GitHub repository. As mentioned earlier, we intentionally avoided extensive prompting and instead employed a straightforward, consistent prompt for all instances. Each model’s chat session commenced with a greeting, followed by a brief introductory sentence regarding the task, and subsequently, all instances were dispatched via the Python script. Appendix A provides information regarding the prompts.

Table 2 and Table 3 presents a concise overview of the results obtained on the dev and test sets. Gemini Pro emerged as the best-performing model, ranking first on both the dev and test sets. Following Gemini Pro, Claude and PaLM, two closed-source LLMs, secured the second and third positions, respectively. Falcon 40B, an open-source LLM, achieved the fourth place and outperformed GPT-3.5. The last two positions were occupied by two open-source LLMs, Llama 2 70B and Mixtral 8x7B.

In the next section, we delve into the detailed performance analysis of the language models on the development set, focusing on specific cases of

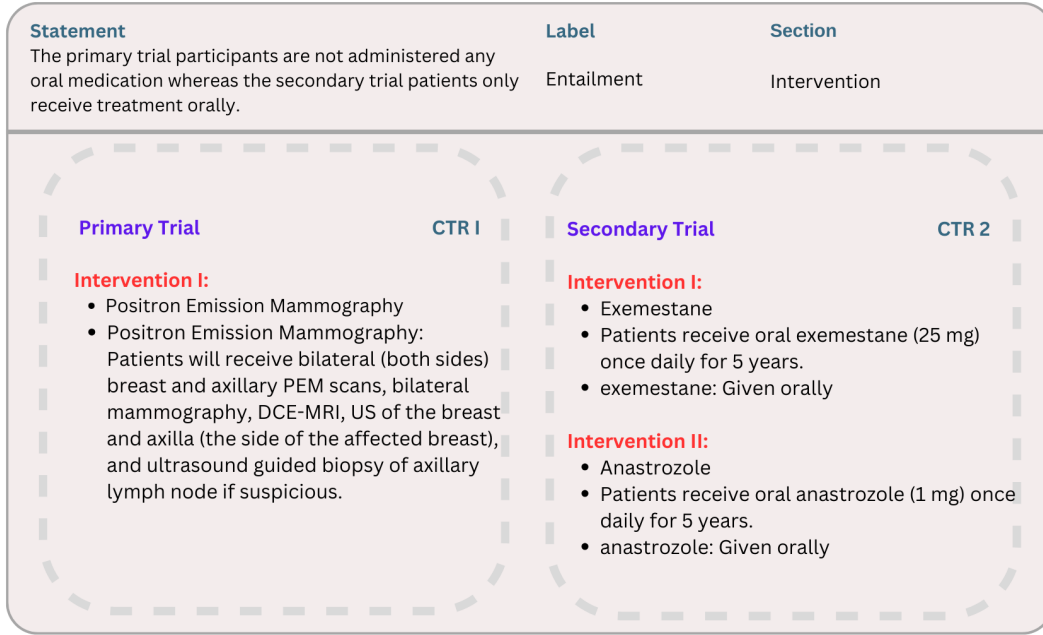


Figure 1: An example comparison task from the training set with two CTRs.

Model	Incorrect
Gemini Pro	36
Claude	38
PaLM	42
Falcon 40B	51
GPT-3.5	59
Llama 2 70B	65
Mixtral 8x7B	69

Table 4: Number of incorrect predictions on the development set of 200 instances for each LLM.

interest.

4.1 General Performance Evaluation

Among the top-ranking LLMs, namely Gemini Pro, Claude, PaLM, and Falcon 40B, their performance on the development set was indeed remarkable. The number of inaccurate predictions made by each LLM on the development set of 200 instances is presented in Table 4. There were only 3 instances in the development set that were incorrectly predicted by all LLMs.

Among the top-performing LLMs, a set of 12 instances emerged as particularly challenging, denoted as "difficult instances". These instances present a significant challenge, as none of the top three performer LLMs in the set - Gemini Pro, Claude and PaLM, were able to accurately predict their outcomes. Within this subset, 2 instances

involved medical abbreviations, 3 featured numerical entities, and the remaining 7 were categorized as "plain instances". The subsequent subsections will delve into a detailed analysis of the difficult instances encompassing numerical entities and abbreviations. For the present discussion, our attention will be directed towards the plain difficult instances.

One particular difficult instance is illustrated in Figure 1, where the top two performing LLMs provided the answers given by Figure 2.

Gemini made a somewhat fair mistake by incorporating knowledge from the training data, suggesting that PET includes orally administered tracer drugs. However, it should be noted that tracer drugs do not qualify as any form of medication for treating cancer or other diseases. In this case, Gemini falters by "assuming too much" and "relying too heavily on prior knowledge." It is important to clarify that this is not a hallucination. On the other hand, Claude correctly deduced that there were no oral medications in the primary trial and that oral medication was administered in the secondary trial. However, Claude failed to synthesize this information and draw a conclusion, indicating a breakdown in reasoning from evidence. Similarly, PaLM and Falcon 40B encountered the same issue as Claude. While they accurately pointed out the evidence, they failed in the reasoning process. In the case of PaLM, it did not provide any indications of en-

GEMINI

Contradiction. The statement “the primary trial participants are not administered any oral medication whereas the secondary trial patients only receive treatment orally” is a contradiction of the given trial information. The primary trial intervention section states that patients will receive Positron Emission Mammography, which involves various PEM views being performed on both the breast and axilla. This indicates that the primary trial participants are administered some form of oral medication, as the PEM views are likely to involve the use of a radioactive tracer that is taken orally. Therefore, it is not accurate to say that the primary trial participants are not administered any oral medication. Hence, the statement is a contradiction of the trial information.

CLAUDE

Contradiction. The primary trial intervention specifies: - Positron Emission Mammography scans. The secondary trial intervention specifies: - Exemestane 25mg oral daily for 5 years.- Anastrozole 1mg oral daily for 5 years. The statement claims:- Primary trial does not administer oral medications\n- Secondary trial only provides oral treatments. However, the primary trial does not mention any oral or other medications. And the secondary trial clearly describes oral aromatase inhibitors as the interventions. As the primary trial information does not validate the negative claim about oral medications, and the secondary trial positively describes oral treatments, the statement is contradicted. Therefore, the answer is contradiction - the trial intervention details conflict with the assertions made in the statement.

Figure 2: An example comparison task from the training set with two CTRs.

gaging in reasoning. Falcon 40B made a literal mistake by becoming mired in the intricacies of the language. Its answer includes the statement, “the statement cannot be directly entailed from the intervention information provided. Specifically, while the primary trial does not mention oral medication, the secondary trial does not exclusively mention oral medication, contradicting the statement.”

These mistakes range from overthinking, failure to follow the evidence, inadequate reasoning, to becoming excessively focused on minute details—a clear manifestation of the inherent challenges that LLMs face. The less performing LLMs, namely GPT-3.5, LLama 2 70B, and Mixtral 7x8B, demonstrated a decent performance considering the task difficulty. However, they exhibited a relatively higher frequency of failures in reasoning and inference compared to the top-performing LLMs.

Having examined the plain difficult instances, we now turn our attention to evaluating the performance of the LLMs on instances containing medical abbreviations.

4.2 Abbreviated Instances Performance Evaluation

In our development set, we identified 31 instances that contained medical abbreviations. We used the ScispaCy package’s abbreviation detector to extract these instances.

Among the top performers, Gemini, Claude,

PaLM, and Falcon 40B made 4, 6, 7, and 8 mistakes, respectively, in handling these abbreviations. The bottom performers, GPT-3.5, LLama 2 70B, and Mixtral 8x70B, made 10, 8, and 8 mistakes, respectively.

Upon closer examination, we found that all of the LLMs were able to correctly resolve the meanings of the medical abbreviations. However, they made mistakes due to other reasoning problems.

For example, none of Gemini’s four mistakes in handling abbreviations were related to resolving their meanings. Similarly, the other LLMs also failed the task primarily due to quantitative-numerical reasoning failures. Appendix B showcases a more comprehensive example of this particular type of occurrence and the corresponding failure.

Overall, the performance of all LLMs in resolving abbreviations was commendable. However, as mentioned before, the majority of failures stemmed from challenges in numerical-quantitative reasoning.

4.3 Numerical Instances Performance Evaluation

Our development set contained 78 instances with numerical entities. We employed the spaCy package (Honnibal and Montani, 2017) and its NER component to identify these entities. To ensure comprehensive semantic evaluation, we combined

ScispaCy and spaCy models.

Among the top-performing LLMs, Gemini, Claude, PaLM, and Falcon 40B made 13, 14, 18, and 19 mistakes, respectively. Notably, the bottom performers, GPT-3.5, Llama 2 70B, and Mixtral 8x70B, made significantly more mistakes (21, 26, and 24, respectively).

Interestingly, the top performers, Gemini and Claude, made the same mistakes on numerical instances in the development set. Upon examining their responses, we observed that they performed arithmetic operations and reasoned based on the calculated results. In Appendix B, Figure 8 and Figure 9 portray instances of successful outcomes achieved by our LLMs. These figures demonstrate accurate performance in arithmetic operations and logical deduction.

However, even the top performers made occasional errors. For instance, Gemini provided an incorrect answer where there was no evidence of arithmetic operations or reasoning: "The primary trial adverse events section shows that there were 10 patients in cohort1 who suffered adverse events out of a total of 67 patients. Therefore, it is accurate to say that over 1/6 patients in cohort1 of the primary trial suffered adverse events." This suggests that Gemini did not calculate 1/6 of the total number of patients (67).

Among the numerical instances incorrectly predicted by Gemini and Claude, we found no instances where arithmetic calculations were performed. Conversely, correctly predicted instances, such as the one shown in Figure 13 in Appendix B, involved at least one mathematical operation that was logically connected to the rest of the argument. We determined these logical connections by analyzing the dependency tree of the answers, as explained in Appendix B. Our findings indicate that when LLMs demonstrate signs of performing arithmetic operations, their results are generally reliable. Conversely, when there is no evidence of arithmetic operations, the result is likely incorrect.

The other top performers, PaLM and Falcon 40B, exhibited similar behavior to Gemini and Claude. They performed arithmetic operations and made deductions based on those operations. When they failed, they did not provide any numerical clues.

The bottom performer, GPT-3.5 was able to perform arithmetic operations. However, it struggled with simple quantity comparisons, such as $n < m$ for random integers. Mixtral 8x7B also faced similar challenges.

Llama 2 70B performed particularly poorly on numerical instances. For the example in Figure 9, where other LLMs succeeded by performing arithmetic operations, Llama failed completely. It provided an incorrect answer without any evidence of subtraction or comparison. In fact, Llama generally struggled with numerical examples, succeeding primarily in quantitative comparisons where operands were provided directly in the context without requiring mathematical processing.

In conclusion, while other LLMs demonstrated proficiency in handling numerical entities, Llama 2 70B failed to meet expectations.

5 Conclusion

Our detailed analysis of LLMs' performance on various reasoning tasks in the medical domain reveals that they are not merely passive memorizers. They possess the ability to perform numerical-quantitative reasoning, general reasoning, and abbreviation resolution, even in a highly specialized domain with unique vocabulary. Notably, Falcon 40B, an open-source LLM, demonstrated impressive performance, rivaling top closed-source LLMs.

Despite their successes, LLMs are not without limitations. Occasional nonsensical predictions highlight the need for caution when using them in high-stakes domains such as medicine. However, the results of our study are highly promising and suggest that with increased training data and computational power, LLMs have the potential to become invaluable tools in the medical field.

The future of LLMs in medicine holds exciting possibilities. As these models continue to evolve, we anticipate that they will play an increasingly significant role in healthcare, transforming the way we diagnose, treat, and prevent diseases.

6 Limitations

As mentioned in earlier sections, we utilized the Poe platform for interacting with LLMs. All the work was accomplished within the confines of a monthly subscription fee of \$20. The results of GPT-4 are not included in this study due to the messaging limit imposed by the platform, which was exceeded by the number of instances in the test set.

References

- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Anthropic. 2023. [Introducing claude](#).
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2023. A survey on evaluation of large language models.
- Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. 2023. [Shortcut learning of large language models in natural language understanding](#). *Commun. ACM*, 67(1):110–120.
- Gemini Team. 2023. [Gemini: A family of highly capable multimodal models](#).
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Chia-Chien Hung, Wiem Ben Rim, Lindsay Frost, Lars Bruckner, and Carolin Lawrence. 2023. [Walking a tightrope – evaluating large language models in high-risk domains](#).
- Shima Imani, Liang Du, and Harsh Shrivastava. 2023. [Mathprompter: Mathematical reasoning using large language models](#).
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Comput. Surv.*, 55(12).
- 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 elio 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. 2024. [Mixture of experts](#).
- Ma el Jullien, Marco Valentino, and Andr e Freitas. 2024. SemEval-2024 task 2: Safe biomedical natural language inference for clinical trials. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*. Association for Computational Linguistics.
- Ma el Jullien, Marco Valentino, Hannah Frost, Paul O’regan, Donal Landers, and Andr e Freitas. 2023. [SemEval-2023 task 7: Multi-evidence natural language inference for clinical trial data](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2216–2226, Toronto, Canada. Association for Computational Linguistics.
- Kamal Raj Kanakarajan and Malaikannan Sankarababu. 2023. [Saama AI research at SemEval-2023 task 7: Exploring the capabilities of flan-t5 for multi-evidence natural language inference in clinical trial data](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 995–1003, Toronto, Canada. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. [Large language models are zero-shot reasoners](#).
- Taeyoon Kwon, Kai Tzu iunn Ong, Dongjin Kang, Seungjun Moon, Jeong Ryong Lee, Dosik Hwang, Yongsik Sim, Beomseok Sohn, Dongha Lee, and Jinyoung Ye. 2024. [Large language models are clinical reasoners: Reasoning-aware diagnosis framework with prompt-generated rationales](#).
- Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Xiangji Huang. 2023. [A systematic study and comprehensive evaluation of chatgpt on benchmark datasets](#).
- Nick McKenna, Tianyi Li, Liang Cheng, Mohammad Javad Hosseini, Mark Johnson, and Mark Steedman. 2023. [Sources of hallucination by large language models on inference tasks](#).
- OpenAI. 2024. [Gpt-4 technical report](#).
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2023. [Med-halt: Medical domain hallucination test for large language models](#).
- Alessandro Stolfo, Zhijing Jin, Kumar Shridhar, Bernhard Sch olkopf, and Mrinmaya Sachan. 2023. [A causal framework to quantify the robustness of mathematical reasoning with language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti 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. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).

Juraj Vladika and Florian Matthes. 2023. [Sebis at SemEval-2023 task 7: A joint system for natural language inference and evidence retrieval from clinical trial reports](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 1863–1870, Toronto, Canada. Association for Computational Linguistics.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#).

Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, and Songfang Huang. 2023. [How well do large language models perform in arithmetic tasks?](#)

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#).

Yuxuan Zhou, Ziyu Jin, Meiwei Li, Miao Li, Xien Liu, Xinxin You, and Ji Wu. 2023. [THiFLY research at SemEval-2023 task 7: A multi-granularity system for CTR-based textual entailment and evidence retrieval](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 1681–1690, Toronto, Canada. Association for Computational Linguistics.

A Prompts

We employed two prompts, the prompts for the individual task and comparison tasks outlined as follows:

"Below find section_name section of the primary trial of a clinical trial. Infer if the following statement entails from the given trial information. Answer should be either entailment or contradiction. Please justify the answer based on numbers. PRIMARY TRIAL section_name: trial_value STATEMENT: statement"

"Below find section_name sections of a primary trial and a secondary trial belonging to same clinical trial. Infer if the following statement entails from the given trial information. Answer should be either entailment or contradiction. Please justify



Figure 3: Initiation of the conversation with PaLM.

the answer based on numbers. PRIMARY TRIAL section_name: trial_value1 SECONDARY TRIAL section_name: trial_value2 STATEMENT: statement".

Figure 3 illustrates the initiation of a chat with PaLM and the method by which instances are provided during the conversation. As evident in the interaction, we maintained minimal prompting and limited additional interactions.

B Example Instances

In this section of the appendix, we present specific instances from the development set to provide readers with a concrete understanding of the performance of LLMs. Firstly, we present a challenging instance, which none of the LLMs in our study were able to correctly predict. Figure 4 depicts

this instance, which involves making an inference about the results section of a single CTR. The inference relates to PFS, a time range spanning from 7.0 to 9.9 months, with an average of 8.4 months. Consequently, the statement presents an entailment. Surprisingly, all the LLMs failed to address this instance. As depicted in Figure 5 and 6, the LLMs struggled to calculate the difference due to various reasons, such as difficulties in numerical deduction or becoming overly focused on linguistic details.

It is worth noting that this instance also includes an abbreviation, PFS, which is fully explained in the body of the CTR. Despite the LLMs demonstrating some understanding of this abbreviation, they ultimately failed due to their inability to perform the necessary numerical inference.

Subsequently, we present a numerical case study depicted in Figure 7 to showcase the numerical reasoning capabilities of the Language and Logic Models (LLMs). Impressively, almost all LLMs accurately predicted this particular case. However, Llama 2 70B exhibited a complete failure, displaying no signs of any numerical inference whatsoever. Figures 8 and 9 illustrate how other LLMs meticulously explained their reasoning step by step. They initiated the process by performing the subtraction $89\% - 88\% = 1\%$ and subsequently compared the result to the claimed amount of 13.2%.

To process numerical instances, we adopted the following approach: firstly, we utilized spaCy's Matcher component to extract all numerical expressions⁴. This component, being part of the pretrained spaCy pipelines, is incredibly helpful in extracting expressions based on patterns. These patterns can involve characteristics such as token shape, POS tags, and even entity types if the token forms part of an entity. By leveraging spaCy's built-in NER component, we could extract various numerical entity types, including cardinal numbers, ordinal numbers, percentages, and quantities. We formulated two general Matcher patterns, namely *NUMERIC OP NUMERIC* and *NUMERIC OP NUMERIC = NUMERIC*, and then generated all possible combinations of numerical entities and mathematical expressions by taking the cross product between numeric entity types and mathematical operator tokens. This comprehensive approach facilitated the extraction of all numerical expressions from the LLM answers. For identifying medical entities, we utilized the ScispaCy package, as medical

entities are not included in spaCy's general-purpose NER models.

Following this, we parsed the dependency tree of the answer and determined the syntactic head of the numerical expression. We then examined whether the numerical expression attached meaningfully to the rest of the answer. For a detailed explanation of the reasoning process, refer to Figure 13.

Moving on to our list of examples, we encounter an intriguing case worth mentioning. Figure 10 presents a CTR with an empty adverse events section, making it a particularly interesting example. As depicted in Figures 11 and 12, all LLMs, except for Falcon 40B, demonstrate impressive intelligence by correctly interpreting "0/0" as indicating the absence of any adverse events, thereby resulting in an empty adverse events section. This example highlights the remarkable general language understanding and common sense reasoning abilities of LLMs, transcending the boundaries of the medical domain.

⁴<https://spacy.io/api/matcher>

Statement	Label	Section
The shortest PFS in cohort 1 of the primary trial was 1.4 months below average.	Entailment	Results

Primary Trial	CTR I
<p>Outcome Measurement:</p> <ul style="list-style-type: none"> • Progression Free Survival • PFS is defined as the period from the day of randomization until the first observation of lesion progression or death from any cause. • <p>Results I:</p> <ul style="list-style-type: none"> • Arm/Group Title: NK105 • Arm/Group Description: received NK105 (65 mg/m²) on days 1, 8 and 15 of a 28-day cycle • Median (95% Confidence Interval) • Unit of Measure: months 8.4 (7.0 to 9.9) 	

Figure 4: A challenging instance that was incorrectly predicted by the top-performing LLMs.

GEMINI
 Contradiction. The statement "The shortest PFS in cohort 1 of the primary trial was 1.4 months below average" is a contradiction of the given trial information. The primary trial results section shows that the median PFS in cohort 1 was 8.4 months (95% CI: 7.0 to 9.9 months). Therefore, it is not accurate to say that the shortest PFS in cohort 1 of the primary trial was 1.4 months below average. Hence, the statement is a contradiction of the trial information.

CLAUDE
 Contradiction. The primary trial results reported: - Outcome Measurement: Progression Free Survival (PFS) - Arm 1 Participants Analyzed: 211 - Arm 1 PFS Median: 8.4 months (7.0 to 9.9) - Arm 2 Participants Analyzed: 211 - Arm 2 PFS Median: 8.5 months (6.9 to 11.5). The statement claims the shortest PFS in cohort 1 was 1.4 months below average. However, the results reported PFS as medians with confidence intervals, not averages. There is no information provided about specific individual PFS durations. Since the results do not include data to validate the claim of any patient being 1.4 months below average, this contradicts the statement. Therefore, the trial results do not substantiate the assertion made in the statement. The answer is contradiction, as the reported outcomes conflict with the statement.

PALM
 Contradiction. The text states that "Median (95% Confidence Interval) 8.4 (7.0 to 9.9)" in the primary trial results 1 and "Median (95% Confidence Interval) 8.5 (6.9 to 11.5)" in the primary trial results 2. This means that the shortest PFS in cohort 1 of the primary trial was not 1.4 months below average.

Figure 5: Responses of the top-performing LLMs to the selected challenging instance, where all models failed to exhibit any signs of numerical inference.

FALCON 40B

Median PFS reported is 8.4 months with a 95% CI of 7.0-9.9 months. Results 2: - Paclitaxel arm analyzed 211 participants- Median PFS reported is 8.5 months with a 95% CI of 6.9-11.5 months. Based on the numbers directly reported: - No cohort or arm assignments are specified - Median values reported are 8.4 and 8.5 months - No PFS of 1.4 months below any average is stated. Therefore, upon my precise numerical analysis, the statement cannot be logically entailed from the PFS results explicitly shared and appears to contradict them.

GPT-3.5

The statement "The shortest PFS in cohort 1 of the primary trial was 1.4 months below average" contradicts the given trial information. Justification: Based on the given trial information, the median PFS for cohort 1 (NK105 arm) is reported as 8.4 months (with a confidence interval of 7.0 to 9.9), while the median PFS for cohort 2 (Paclitaxel arm) is reported as 8.5 months (with a confidence interval of 6.9 to 11.5). Neither of these results indicates the shortest PFS being 1.4 months below the average. Therefore, the statement contradicts the given trial information as it incorrectly suggests that the shortest PFS in cohort 1 was 1.4 months below the average.

LLAMA 2 70B

Contradiction. The statement "The shortest PFS in cohort 1 of the primary trial was 1.4 months below average" implies that the shortest PFS in cohort 1 was 1.4 months less than the average PFS in cohort 1. However, the results provided in the PRIMARY TRIAL RESULTS section show that the median PFS for cohort 1 (NK105) was 8.4 months, and the median PFS for cohort 2 (Paclitaxel) was 8.5 months. Since the median PFS for both cohorts is above 8 months, it is not possible for the shortest PFS in cohort 1 to be 1.4 months below average.

MIXTRAL 8x7B

The statement "The shortest PFS in cohort 1 of the primary trial was 1.4 months below average" cannot be inferred from the given trial information. The results show that the median PFS for Arm/Group Title: NK105 was 8.4 (7.0 to 9.9) and for Arm/Group Title: Paclitaxel was 8.5 (6.9 to 11.5). Neither of these results indicate that the shortest PFS was 1.4 months below average.

Figure 6: Responses of the low-performing LLMs to the selected challenging instance, which were not significantly different from the answers provided by the top LLMs.

Statement	Label	Section
there is a 13.2% difference between the results from the two the primary trial cohorts.	Contradiction	Results

Primary Trial	CTR I
Results I: <ul style="list-style-type: none">• Arm/Group Title: Exemestane• Unit of Measure: percentage of participants 88 (87 to 89)	
Results II: <ul style="list-style-type: none">• Arm/Group Title: Anastrozole• Unit of Measure: percentage of participants 89 (88 to 90)	

Figure 7: A numerical example extracted from the development set, focusing on inferring information from a single CTR's results section.

GEMINI

Contradiction. The statement "there is a **13.2%** difference between the results from the two the primary trial cohorts" is a contradiction of the trial information provided. To calculate the difference between the results from the two primary trial cohorts, we can subtract the event-free survival rate in the exemestane arm from the event-free survival rate in the anastrozole arm: **89% - 88% = 1%**. Therefore, the difference between the results from the two primary trial cohorts is **1%**, not **13.2%**. Hence, the statement is a contradiction of the trial information.

CLAUDE

Contradiction. The primary trial results provide event-free survival percentages of 87-89% for the exemestane cohort and 88-90% for the anastrozole cohort. These ranges show overlap between the results of the two cohorts. There is no information given that directly states or implies a 13.2% difference.

PALM

Contradiction. The text states that the event-free survival rate in the exemestane arm was **88% (87 to 89)**, while the event-free survival rate in the anastrozole arm was **89% (88 to 90)**. This means that the difference between the two arms was only **1%**, not **13.2%**.

FALCON 40B

The statement "there is a 13.2% difference between the results from the two primary trial cohorts" would be a contradiction. The results reported a range of **87-89%** for the exemestane cohort and **88-90%** for the anastrozole cohort. This shows only a difference of **1-2%** between the lowest and highest estimates provided. **13.2%** is much greater than the actual difference observed between these two cohorts for the 5 year event-free survival endpoint.

Figure 8: Responses of the top-performing LLMs to the selected numerical instance, showcasing impressive performance in numerical inference across all models.

GPT-3.5

To determine if there is a **13.2%** difference between the results from the two primary trial cohorts, let's compare the percentages provided in the results. In Results 1 (Exemestane group), the percentage of participants analyzed is **88% (87 to 89)**, and in Results 2 (Anastrozole group), the percentage of participants analyzed is **89% (88 to 90)**. The difference between these two percentages can be calculated as follows: **89% - 88% = 1%**. Based on the provided numbers, the calculated difference is **1%**, not **13.2%**. Therefore, the statement "there is a 13.2% difference between the results from the two primary trial cohorts" contradicts the information given in the primary trial results.

LLAMA 2 70B

ENTAILMENT. The statement mentions a **13.2%** difference between the results from the two primary trial cohorts, which is within the margin of error for the confidence interval provided in the results (**87-89%** for exemestane and **88-90%** for anastrozole). Therefore, the statement entails from the given trial information.

MIXTRAL 8x7B

Contradiction. The statement is not correct. The difference between the two groups is actually **1 percentage** point (**89%** for anastrozole minus **88%** for exemestane), which is a relative difference of about **1.12%** ($1/88 * 100$). This is not the same as a **13.2%** difference as stated

Figure 9: Responses of the bottom-performing LLMs to the selected numerical instance, where all models, except for Llama 2 70B, successfully performed the subtraction operation and made the corresponding numerical inference.

Statement	Label	Section
The adverse events section in the primary trial is empty.	Entailment	Adverse Events
<p>Primary Trial</p> <p>Adverse Events 1:</p> <ul style="list-style-type: none"> Total: 0/0 		
<p>CTR I</p>		

Figure 10: A CTR from the training set that is both interesting and uninteresting. It is interesting because the value of 0 denotes emptiness, but uninteresting because humans can understand 0 as representing null.

GEMINI
 Entailment. The statement "The adverse events section in the primary trial is empty" is an entailment of the given trial information. The primary trial adverse events section shows a total of 0 adverse events. Therefore, it is accurate to say that the adverse events section in the primary trial is empty

CLAUDE
 Entailment. The adverse events section provided reports: Adverse Events 1: Total: 0/0. The statement says "The adverse events section in the primary trial is empty". Since the only information given in the adverse events section is: Adverse Events 1: Total: 0/0. And there are no other details listed, this directly matches the statement that the adverse events section is empty. Therefore, the statement can be directly substantiated and entailed by the adverse events data shown, which contains no reported events. The answer is entailment, as the content or lack of content in the adverse events section supports the claim made in the statement.

PALM
 Entailment. The text states that "Total: 0/0". This means that the adverse events section in the primary trial is empty.

Figure 11: Responses from the top-performing LLMs, demonstrating high intelligence and deliberate reasoning.

FALCON 40B

Here, the statement is that "The adverse events section in the primary trial is empty". Since the section provided does contain information (Adverse Events 1 and a total of 0/0), even if no specific adverse events are listed, it cannot logically be considered empty. Therefore, based on reviewing the details shared, the statement that "The adverse events section in the primary trial is empty" would contradict the content presented in the PRIMARY TRIAL ADVERSE EVENTS.

GPT-3.5

Based on the given adverse events section from the primary trial, the statement "The adverse events section in the primary trial is empty" entails from the trial information. Here's the justification: In the adverse events section provided for the primary trial, the data is presented as follows: Adverse Events 1: Total: 0/0. The total number of adverse events reported is indicated as 0/0. This suggests that no adverse events were reported in the primary trial. Since there are no adverse events mentioned or reported, it can be inferred that the adverse events section in the primary trial is indeed empty.

LLAMA 2 70B

The Adverse Events section in the primary trial is shown to be empty, with a total of 0/0 adverse events reported. Therefore, the statement is true and entails from the given information.

MIXTRAL 8x7B

Entailment. The adverse events section of the primary trial does not list any adverse events, indicating that no adverse events were recorded during the trial.

Figure 12: Responses from the low-performing LLMs, where all models, except for Falcon 40B, achieved success comparable to the top performers.

To calculate the difference between the results from the two primary trial cohorts, we can subtract the event-free survival rate in the **exemestane CHEMICAL** arm from the event-free survival rate in the **anastrozole CHEMICAL** arm: **89% - 88% = 1% NUMERIC**

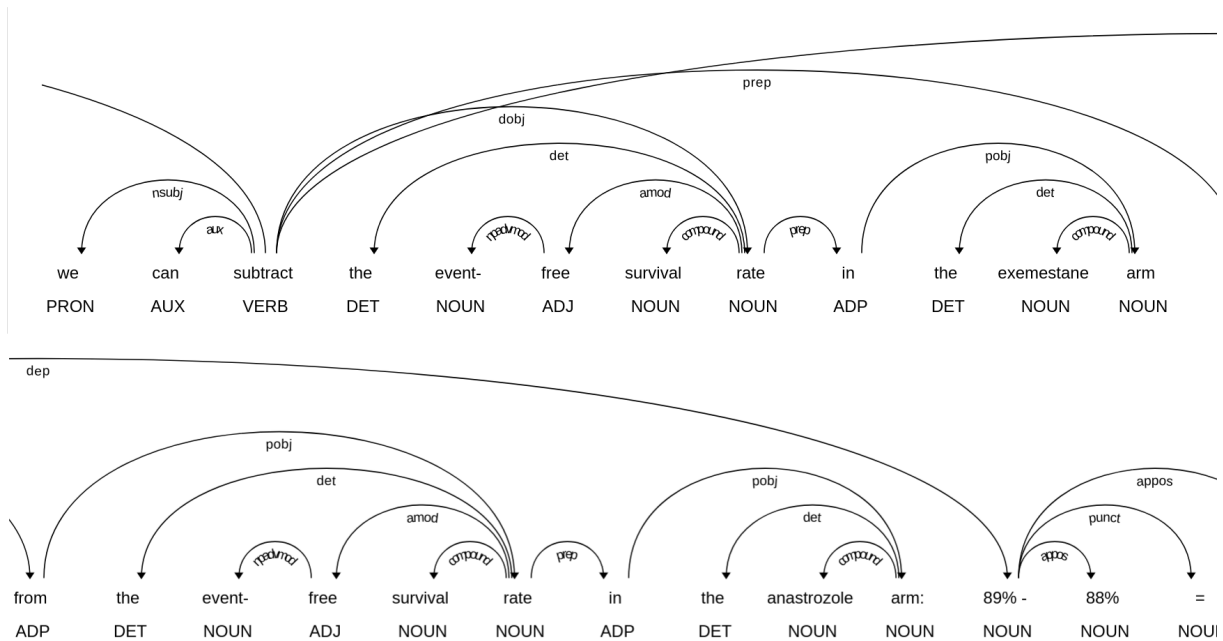


Figure 13: Semantic parse of a successful answer by Gemini. Named entities are highlighted in the above picture, where dependency tree of the sentence is exhibited in the below pictures. In the dependency tree, head token of the numerical expression **89% - 88% = 1%** is **89%** and syntactic head of **89%** is **subtract**, which is the mathematical operation. By following the syntactic parent of the numerical expression, we reach the explanation of the chain of mathematical operations, hence we can deduce that Gemini put down a valid argument and numerical reasoning.