

Evaluating LLM Reasoning in the Operations Research Domain with ORQA

Mahdi Mostajabdaveh¹, Timothy Tin Long Yu¹, Samarendra Chandan Bindu Dash^{1,2},
Rindranirina Ramamonjison¹, Jabo Serge Byusa¹, Giuseppe Carenini³, Zirui Zhou¹, Yong
Zhang¹,

¹Huawei Technologies Canada, 4321 Still Creek Dr, Burnaby, BC, V5C 6S7, Canada

²University of Toronto, 40 George St, Toronto, ON, M5S 2E4, Canada

³University of British Columbia, 2366 Main Mall, Vancouver, BC, V6T 1Z4, Canada

{mahdi.mostajabdaveh1, timothy.yu, rindranirina.ramamonjison, jabo.byusa, zirui.zhou, yong.zhang3}@huawei.com,
carenini@cs.ubc.ca, dashesam1@cs.toronto.edu.

Abstract

In this paper, we introduce and apply **Operations Research Question Answering (ORQA)**, a new benchmark, to assess the generalization capabilities of Large Language Models (LLMs) in the specialized technical domain of Operations Research (OR). This benchmark is designed to evaluate whether LLMs can emulate the knowledge and reasoning skills of OR experts when given diverse and complex optimization problems. The dataset, crafted by OR experts, presents real-world optimization problems that require multi-step reasoning to build their mathematical models. Our evaluations of various open-source LLMs, such as LLaMA 3.1, DeepSeek, and Mixtral reveal their modest performance, indicating a gap in their aptitude to generalize to specialized technical domains. This work contributes to the ongoing discourse on LLMs' generalization capabilities, providing insights for future research in this area. The dataset and evaluation code are publicly available.

Code and Dataset — <https://developer.huaweicloud.com/develop/aigallery/notebook/detail?id=6b98c56e-913b-47ef-8d9f-3266c8aec06a>

Introduction

The ability of Large Language Models (LLMs) to follow human instructions and perform diverse tasks has made them an exciting area of investigation. Moreover, the considerable interest in adopting LLMs across various complex technical domains (e.g., medicine (Gao et al. 2023; Zhou et al. 2023)) highlights their potential for significant societal impact. A particularly compelling driver for this adoption is the potential of LLMs to automate many tasks, reducing human intervention and improving productivity. However, as LLMs are becoming integrated into the workflow of various industries, it is important to thoroughly understand their capabilities and limitations (Khatun and Brown 2024a; Baktash and Dawodi 2023; Truong et al. 2023). Of particular interest is their ability to reason and perform new challenging tasks across different domains, which are reported critical limitations of LLMs (Arkoudas 2023; Shen and Kejriwal 2023). Our work addresses this need by introducing a new benchmark dataset and applying it to assess these limitations.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To evaluate an LLM's ability to generalize to a new domain, we focus on the field of Operations Research (OR). The choice of this domain is deliberate and significant. First, OR is important for making decisions in various industries (Petropoulos et al. 2024). Second, there are many types of optimization problems applied to real-world applications ranging from production scheduling (Mostajabdaveh, Salman, and Tahmasbi 2022) to creating efficient delivery routes for trucks (Vidal, Laporte, and Matl 2020). Third, optimization modeling presents a unique challenge due to the expert-level knowledge and reasoning skills it requires (Hillier and Lieberman 2015), adding a layer of complexity to the automation of this task. Some recent studies also report modest performances of SOTA LLMs such as GPT-4 and Llama2 for optimization model building tasks (AhmadiTeshnizi, Gao, and Udell 2024; Mostajabdaveh et al. 2024b). Finally, OR is a niche field with limited publicly available text corpora or optimization model code (Xiao et al. 2023; AhmadiTeshnizi, Gao, and Udell 2024), making it an ideal testbed to assess the generalizability of LLMs, reducing the risk of data contamination.

To assess an LLM's knowledge and reasoning skills on unseen, diverse, and complex optimization problems, we propose ORQA (**O**perations **R**esearch **Q**uestion **A**nswering), a new multi-choice Question Answering (QA) benchmark dataset, crafted and verified by OR experts. Each dataset instance (Figure 1; left) presents a natural language description of an optimization problem along with a question that requires multi-step reasoning (Figure 1; right) to answer correctly.

Our study is significant for several reasons. First, it contributes to the ongoing dialogue on the generalizability of LLMs. Although many benchmarks claim that LLMs can replicate expert-level knowledge across various technical domains, the actual extent of these models' generalization capabilities remains an open question (Alzahrani et al. 2024). Our benchmark offers a new perspective by focusing on a specialized technical domain that lacks a large-scale, high-quality dataset. To the best of our knowledge, this is the first multi-choice QA dataset in the field of Operations Research. Second, this study has implications for understanding the potential and limitations of LLMs in automating tasks in niche technical fields like OR, where ex-

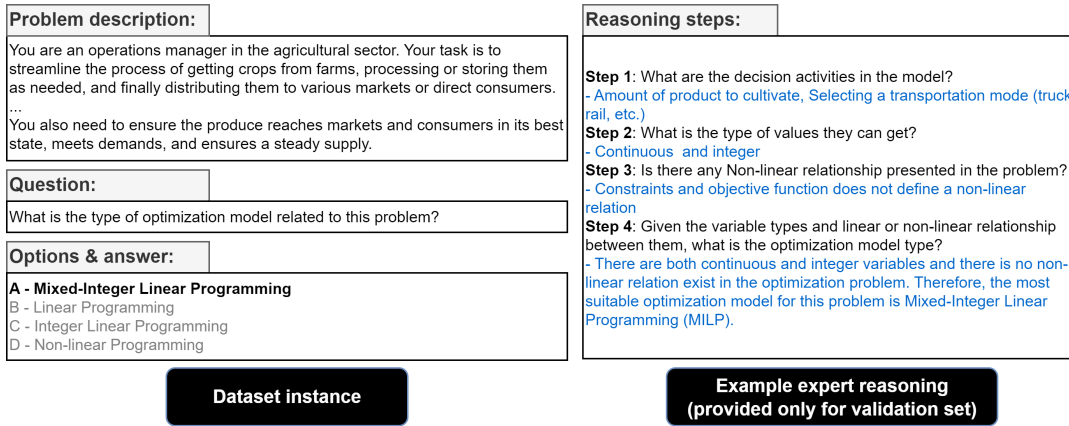


Figure 1: **Left:** Dataset instance containing the description, question, options, and answer (bold). **Right:** Example reasoning steps needed to answer the simple question.

pert knowledge and complex multi-step reasoning are crucial. Optimization modeling requires access to specialized OR experts, which is often impractical for most potential users due to the associated costs.

We choose to focus on translating textual problem descriptions into mathematical optimization models rather than directly solving optimization problems. Even specialized AI models struggle with scalability and generalization when solving simple OR problems (Joshi et al. 2022). We also highlight the challenges of constructing a technical dataset like ORQA. Samples of optimization problems are inherently complex and require significant time and effort to create and verify. Moreover, ensuring correctness demands annotators with graduate-level education or extensive experience in OR modeling, making the process expensive, time-consuming, and labor-intensive.

Related Work

LLMs applied to operations research: Within the field of OR, LLMs are being investigated for their ability to formulate optimization models (Fan et al. 2024). Ramamonjison et al. (2022a) proposed using generative models to automate the formulation of OR problems from natural language. Building on this, Ramamonjison et al. (2022b) introduced methods for recognizing entities and parsing optimization formulations from text. However, these works primarily utilize a toy dataset of elementary linear programming word problems. More recently, AhmadiTeshnizi, Gao, and Udell (2024) proposed the NLP4LP dataset, which includes 332 instances with structured natural language descriptions (SNOP), parameter data values, optimal value, and optimal solution. However, the majority of the problems are still toys, their description is technical, with mathematical notation, and lacks context from real applications. Xiao et al. (2023) introduced the ComplexOR dataset with only 37 optimization problems. This dataset includes NL descriptions, optimization models, and several input data. While their problem descriptions are context-aware, they often mention the related optimization problem by name

(e.g., lot-sizing problem with setup), and only cover a narrow range of application domains. With more than 1.5k instances, our dataset is the largest, offering significant value through its depth, rigor, realism, and diversity.

Another disadvantage of existing mentioned datasets is they require the optimization model to be solved for their evaluation. LLMs must generate a model code from the NL description, which is then fed to a solver with data to obtain the optimal value or solution. The evaluation focuses solely on the correctness of the optimal value or solution, presenting two key limitations: (1) it is an end-to-end evaluation that cannot differentiate between minor notation errors and entirely incorrect structures; (2) it cannot distinguish between errors in code generation and errors in model formulation. (AhmadiTeshnizi, Gao, and Udell 2024) show that coding errors account between 21% to 31%. In contrast, ORQA tackles more complex optimization tasks across diverse application scenarios. Our problem descriptions are context-aware, relevant to real applications, and free of OR jargon and mathematical notations. Additionally, ORQA is multiple-choice question answering dataset offers a straightforward evaluation that is independent of model code and does not require solvers.

Multi-choice question answering: The multi-choice QA task involves receiving a question along with several candidate options and selecting the correct answer. Many such datasets have been made publicly available (Talmor et al. 2019; Mihaylov et al. 2018; Clark et al. 2020). However, the complexity of current reasoning benchmarks has been called into question (Khatun and Brown 2024b; Valmeekam et al. 2023; Sawada et al. 2023), motivating the creation of more challenging benchmarks that surpass basic common-sense reasoning (Kweon et al. 2024; Sawada et al. 2023). While multi-choice QA is a well-studied NLP task, ORQA stands out as a handcrafted, expert-curated dataset in the technical field of OR. This domain is notably underrepresented in current benchmarks and demands deep optimization knowledge. Moreover, tasks in ORQA require identifying optimization model components and their interrelation-

ships, which necessitates multi-step reasoning.

LLM reasoning capabilities and limitations: Zhou et al. (2024) demonstrated that LLMs can be prompted to reason, leading to improved performance and insights into how decisions are made. Techniques for this purpose include multi-step chained prompts (Yoran et al. 2023; Kojima et al. 2022), single-step chain-of-thought (Kojima et al. 2022), tree-of-thought (Yao et al. 2024), and chain-of-thought with self-consistency and verification (Zhao et al. 2023). While these methods are promising, limitations persist in leveraging and evaluating LLMs’ reasoning capabilities. Evaluating the faithfulness of reasoning, as highlighted by Lanham et al. (2023), is a significant challenge. Most existing reasoning benchmarks are overly simplistic, indicating a need for more complex benchmarks (Valmeekam et al. 2022; Laban et al. 2023). Given this need, our ORQA benchmark dataset is designed to test the reasoning abilities of LLMs in the demanding OR context.

Task: Identifying Optimization Model Characteristics

Task background and motivation

An optimization model is a mathematical representation of a decision-making problem. Optimization models are constructed using various components. These components include elements, decision activities, data attributes, calculations, objective criteria, and specifications (Sánchez et al. 2021). The first and most crucial step in formulating an optimization model is to identify the components of the model and understand their relationships, as any error in this step will result in an incorrect model. ORQA focuses on identifying these components and their relations from the natural language description of the optimization problem by asking questions such as “What are the decision activities of the optimization problem?” and “Which data parameters are participating in the objective criterion?”.

To illustrate these components and their relationships, we refer to the parking spot assignment example problem described in Figure 2. The figure demonstrates how the extraction of optimization problem components and their relationships can be directly used to formulate the optimization problem mathematically. Unique apartment units and parking spots are the elements of this example problem and their data attributes directly map to sets (i.e., apartment groups, parking spots) and data parameters (i.e., parking need, number of vehicles, etc.). The decision activities are direct actions in the system and define the optimization variables (i.e., assign vehicles to parking spots). These three model components are combined to form the utility/-cost function to be maximized/minimized (i.e., minimize total distance). They also form the specifications that define business rules or system limitations, which lead to optimization constraints.

Task definition

We propose a multi-choice QA task to identify the components of the optimization problem, their attributes, and

relationships, from a given natural language problem description. This is a highly complex task requiring multi-step reasoning as these components have multiple layers of interaction and dependencies (Hillier and Lieberman 2015). For example, identifying the objective criteria involves not only recognizing the objective measure and sense, but also determining the specific data attributes and decision activities that influence it. Figure 7 in Appendix of Mostajabdaveh et al. (2024a) illustrates these complex relationships between problem components.

Task characteristics

The task we propose is complex not only due to the heavily mathematical nature of the field of OR, but also the complexity of the optimization models the dataset is built upon. The complexity is directly related to the number of components in the corresponding mathematical model. We describe in Section 4.2 our approach to ensure a standard level of complexity during our dataset creation process.

Additionally, the task is difficult due to the underrepresentation of open-sourced OR data during LLM training. The findings from Kandpal et al. (2023) and Mallen et al. (2023) align with our claim that scarcity in optimization modeling-related data would make this task challenging for LLMs. In this regard, we consider optimization modeling as a task that requires long-tail knowledge.

ORQA Dataset

Dataset overview

Each dataset instance contains the following (see Figure 1, left):

1. A *CONTEXT* describing an optimization problem as a case study using natural language,
2. A *QUESTION* asking about the problem specifications (e.g., objective criterion or constraints), the underlying components of the model (e.g., the elements participating in the optimization), or the structure and logic of the optimization model (e.g., the logical relation between two components),
3. A list of *OPTIONS* for the answer, which was created by OR experts to make the question challenging. The LLM must select the correct answer from a list of four options.
4. The correct *TARGET ANSWER*.

Table 1 presents the characteristics of the dataset. A wide range of application domains are represented within ORQA ranging from common problems (e.g., Traveling Salesman Problem) to niche problems (e.g., multi-period production planning of a drone manufacturing company). ORQA is comprised of a total of 20 application domains each represented by at least three problems and 60 to 90 multi-choice questions. Some of these domains include healthcare, urban design, human resources, petroleum, and sales. ORQA is comprised of 1513 data instances with 45 instances allocated as the validation set. The validation set provides the in-context learning (ICL) examples used for few-shot prompting. We include expert-written reasoning steps for the instances in the validation set.

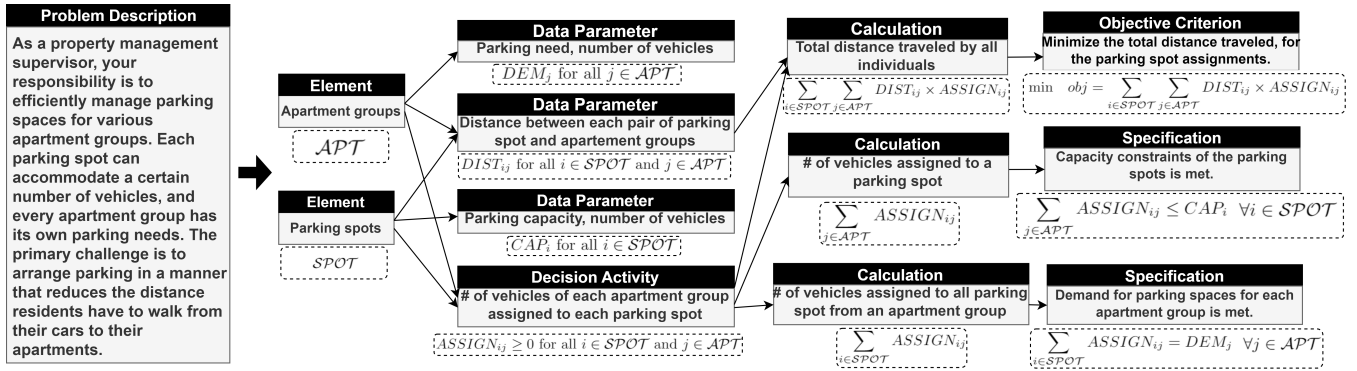


Figure 2: An example of optimization problem components, their relationships, and corresponding mathematical formulations.

Characteristics	ORQA
Number of instances	1513
Test/validation split	1468/45
Average input length (words)	231
Number of domains	20

Table 1: ORQA dataset statistics.

Question type. The questions in this QA task can be split into 11 question types that have been derived from three critical skills of optimization modeling. Specifically, (A) understanding the high-level problem specifications, (B) identifying entities of the corresponding optimization model, and (C) identifying relationships between components. For examples and additional details of the 11 question types, please refer to Table 8 in Appendix of Mostajabdaveh et al. (2024a).

Dataset creation

The dataset was carefully created and verified by a team of five experts with extensive experience in optimization modeling: a Bachelor’s graduate, two Master’s graduates, and two PhDs. An example of a data instance is shown in Figure 1. The experts went through a joint training session that provided them with step-by-step instructions created by the lead OR expert for the dataset creation process. In the training session, all experts annotated the same three data instances to ensure that they were aligned. Then, each OR expert was assigned a subset of problems to create independently. The OR experts used available optimization models from OR textbooks, academic journals, and online code repositories as the initial source and they substantially modified and diversified them in a multi-step, human-led dataset creation process. Due to the labor-intensive process of creating this dataset, we did not have OR experts overlap in creating the same data instances. Instead, we prioritized developing a larger benchmark and relied on our training session and a rigorous verification stage (as described below) to ensure consistency in difficulty and quality.

The selection, creation, and verification process was comprised of three steps and is detailed in Figure 3. In step 1, problems were filtered based on a set of criteria, which was

defined prior to the selection process. The criteria for selection include the problem complexities, diversity in the resulting models types, and practicality for real-world use cases. To regulate the complexity of the dataset, OR experts were instructed to select problems where the number of components in their mathematical model is within pre-defined limits (e.g., the number of decision variables should be between 2 and 7). As a result, the average and standard deviation of the model components are as follows: sets 1.97/0.89, parameters 4.08/2.19, variables 3.14/2.29, objectives 1/0.0, and constraints 4.60/3.21.

To regulate problem diversity, each OR expert was given a set of application domains and instructed to select at least three problems from each domain. Given that most optimization models lack concrete natural language descriptions, our OR experts crafted handwritten problem descriptions and, where applicable, modified the problem context to ensure coverage across diverse application domains. Once the OR expert was satisfied with the problem description, they would verify that the natural language description exactly mapped to the original mathematical model.

In step 2 (Figure 3), two OR experts were responsible for creating and annotating the question, target options, and target answer of each instance of data. Questions were designed to promote multi-step reasoning and they were created along with the options and target answer by referencing both the mathematical model and problem description. For cases where multiple modeling approaches are possible, OR experts ensured that incorrect options were truly incorrect considering all different models.

Finally, in step 3 (Figure 3), each instance was verified by another cohort of two OR experts. These experts verified that each data instance was complete and correct, devoid of ambiguities in its question and options, required multi-step reasoning, and free of sensitive information (e.g., real-world people and organizations names).

Experiment Setup for Evaluation

To provide an initial assessment of the difficulty of ORQA and gain insights, we ran a series of experiments to benchmark different LLM models using various prompting strategies.

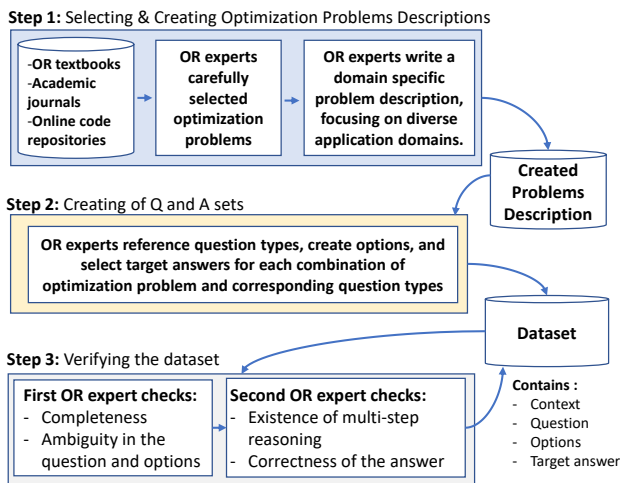


Figure 3: Selection, creation, and verification process for the ORQA dataset.

Baseline models. Specifically, we run inference on open-source LLMs such as the 11B parameter FLAN-T5 XXL (Chung et al. 2024), Falcon-7B-Instruct (Almazrouei et al. 2023), Mistral model series (Jiang et al., 2023), Mixtral series (Jiang et al. 2024), as well as Llama2, 3 and 3.1 models of different sizes (Touvron et al. 2023; Dubey et al. 2024). We made a deliberate decision not to use closed-source LLMs in the spirit of scientific reproducibility. Prompting an LLM through APIs has unquantifiable uncertainties in how inputs and outputs are processed. Moreover, the LLMs may change or be deprecated, potentially invalidating our results. Model endpoints are commonly deprecated as newer models are made available ¹.

We evaluated each model on the 1468 instances of data from the test set for its standard and CoT prompting capabilities in both zero-shot and few-shot settings, as described in Section 3. Table 2 outlines the accuracy performance of the LLMs evaluated using this benchmark. Furthermore, we report the average F1 scores in Table 6 in Appendix of Mostajabdaveh et al. (2024a). As a preliminary human baseline, a single expert with a related Ph.D. achieved 93% accuracy on a random set of 100 instances without any in-context learning examples.

Prompting strategies. We evaluate the LLM’s ability to reason with different prompting strategies. First, all-at-once (standard) prompting evaluates the LLM’s robustness to perform these reasoning steps without explicitly prompting it to reason. Conversely, CoT prompting is implemented as a two-step approach following similar works on multi-choice QA (Wei et al. 2022; Yoran et al. 2023; Kojima et al. 2022). Specifically, the first prompt elicits the LLM to reason “step-by-step”. Then, the generated reasoning is added to the prompt and given to the LLM to generate the final answer. Both standard prompting and CoT are evaluated for zero-shot and few-shot capabilities, where few-shot prompts are

¹<https://platform.openai.com/docs/deprecations>

Model	Standard (Acc)			CoT (Acc)	
	0-shot	1-shot	3-shot	0-shot	1-shot
Llama3.1-8B-I	0.588	0.615	0.618	0.563	0.324
Llama3.1-70B-I	0.702	0.721	0.735	0.689	0.292
Llama3.1-405B-I	0.723	0.753	0.772	0.695	0.360
Llama3-8B-I	0.535	0.573	0.592	0.530	0.364
Llama3-70B-I	0.676	0.716	0.710	0.671	0.448
Llama2-7B-Chat	0.368	0.375	0.403	0.368	0.282
Llama2-13B-Chat	0.409	0.437	0.454	0.432	0.313
Llama2-70B-Chat	0.526	0.552	0.589	0.518	0.372
FLAN-T5-XXL-11B	0.503	-	-	0.457	-
Falcon-7B-I	0.245	0.246	0.245	0.242	0.243
DeepSeek-M-7B-I	0.478	0.552	0.559	0.379	0.514
NuminaMath-7B	0.250	0.484	0.525	-	0.290
Mistral-7B-I-v0.1	0.467	0.475	0.483	0.460	0.407
Mistral-7B-I-v0.3	0.523	0.555	0.555	0.539	0.543
Mixtral-8x7B-I-v0.1	0.588	0.606	0.612	0.565	0.565

Table 2: Accuracy for each model across different prompting strategies. In model names, *I* stands for Instruct. Empty entries for FLAN-T5 are due to engineering limitations as a result of prompts exceeding the LLM’s input token limit. Due to the slow generation speed of the NuminaMath model, which took around 10 days to generate reasoning steps, we skipped the 0-shot CoT experiment.

created by randomly sampling instances with the same question type from the validation split. Note that for ICL examples, ground-truth reasonings are excluded in standard prompting but included in CoT.

Evaluation protocol. Our custom evaluation framework takes inspiration from Robinson and Wingate (2023) and binds each option to a symbol (i.e., A, B, C, D). We further discovered that appending the prompt “Therefore, among A through D, the answer is (” to the end reliably guided all tested LLMs to output the required format. We bind each answer to a symbol to avoid punishing models that display the proper reasoning ability to reach the correct solution, but hallucinate or cannot output one of the options exactly. We calculate the accuracy by comparing the exact match between the generated answer and the corresponding ground truth.

An example of the different components of the prompt is shown in Figure 4. Standard and CoT prompting strategies are comprised of the same components shown in the figure. The only difference is that zero-shot prompts omit the explicit REASONING step, and CoT prompts would use two different INPUT TAGS between the trigger prompt and the answer-eliciting prompt. As mentioned, CoT is performed through a two-step approach. The first step extracts the REASONING component (Figure 4). The second step appends the REASONING component to the trigger prompt to generate the final answer.

To create the reason extraction prompt, we take the same format as the few-shot prompts, but append the TRIGGER PROMPT (e.g., “Let’s think step by step”) after the list of OPTIONS.

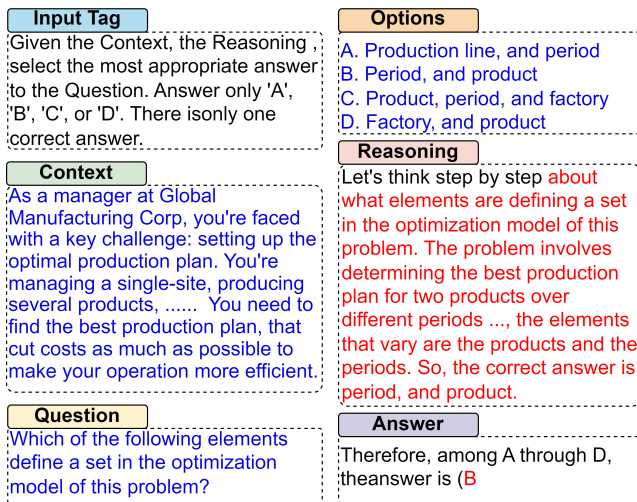


Figure 4: The different components of a prompt. The pre-defined text is in black; we provide an example (in blue), and the LLM output (in red).

Results and Discussion

We show the potential of ORQA as a useful benchmark for LLMs across multiple dimensions of interest. Here are some key findings:

- Model size contributes to reasoning performance.

Not surprisingly, an LLM's reasoning ability is correlated to the size of an LLM. As shown in Table 2, this is true when comparing models within the same family (e.g., Llama 3.1). However, some models such as Mistral-7B and Flan-T5 perform better than Llama2-13B despite being smaller. This supports the findings from (Jiang et al. 2023) that Mistral 7B outperforms Llama2-13B across their evaluation benchmarks potentially due to Mistral's application of sliding window attention that allows Mistral to better handle the long natural language model descriptions (Jiang et al. 2023).

- CoT generally drops the performance. ICL examples benefit standard but not CoT prompting. Answering ORQA questions requires multi-step reasoning, and CoT prompting has been shown to elicit reasoning in LLMs (Wei et al. 2023). However, our experiments surprisingly indicate decreased performance on ORQA when using CoT (Table 2 and Figure 6). By investigating the generated reasoning, we found that the models often ignored instructions. For example, although the prompt specifies that only one option is correct, the models would attempt to generate reasoning that selects none or multiple options. We also report hallucinations where models would create their own options and select those. Finally, the reasonings were often incorrect, as highlighted in Figures 9 and 10 in Appendix of Mostajab-daveh et al. (2024a). A promising direction is to explore more advanced CoT prompting techniques, including possible extensions of the faithful CoT reasoning presented by Lyu et al. (2023).

- Trigger prompts impact performance. Prompt ensembling improves CoT performance. Inspired by Liévin et al. (2024), we experimented with different trigger prompts

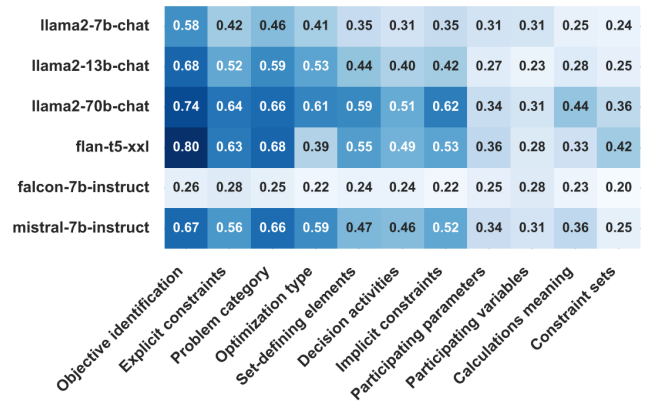


Figure 5: Heatmap of LLM performance on different question types. Performance is the average accuracy over the five prompting strategies.




Figure 6: Performance comparison of Standard and CoT prompting. Questions of Category B & C require more OR and/or model building knowledge.

for CoT. These experiments were conducted using the following settings: 0-shot with Llama-3.1-70B-Instruct, temperature set to 0.7, and each trigger prompt was run five times. The results are recorded in Table 3. Different trigger prompts can change the performance, with scores ranging from 0.648 to 0.689.

- Fair performance on reading comprehension. Poor performance if a question requires OR knowledge and model building knowledge. As shown in Figure 5, the questions requiring only reading comprehension (left-side of the heatmap) resulted in fair performance. However, the questions that require OR and/or model building knowledge (right-side of the heatmap) were too difficult and many LLMs performed poorly. These trends can also be seen in Figure 6 where LLMs performed worse on the more difficult questions. To enable LLMs to deal with such complex questions is a critical venue for future work. One possibility is providing domain-specific knowledge bases or enabling API calls during reasoning, which has been shown to enhance an LLM's ability to perform knowledge-intensive

CoT Prompt (0-shot with Llama-3.1-70B-Instruct)	Average	STD	Best
Let’s think step by step	0.688	0.001	0.689
Let’s work by elimination	0.648	0.000	0.649
Let’s reflect on each answer option like an operations research expert	0.689	0.001	0.689
Let’s use step by step inductive reasoning, given the mathematical nature of the question	0.674	0.003	0.676
Let’s think step by step like an operations research expert	0.685	0.000	0.685
Prompt ensembling (majority vote)	0.696	0.008	0.702

Table 3: Impact of preceding trigger prompts on Llama-3.1-70B-Instruct using 0-shot CoT. Results are averaged over five independent runs per prompt. The bottom row shows the accuracy of ensembling all prompts.

reasoning tasks (Yao et al. 2023). Another option would be to explore architecture alternatives to transformers such as StripedHyena (Poli et al. 2023), that could better deal with the concerning findings of Dziri et al. (2024) which suggest transformer-based LLMs reduce compositional reasoning into linearized subgraph matching. In essence, these findings seem to indicate that transformer-based LLMs are inherently limited in their ability to perform complex multi-step reasoning. Along these lines, ORQA could represent a useful benchmark for testing LLMs based on those novel architectures.

- In-context learning can reduce the number of reasoning steps and reasoning errors. We analyzed reasoning steps generated by Llama2-13b-chat model in 0-shot and 1-shot CoT settings on our validation set. OR experts manually evaluated a total of 90 instances (45 per setting) and classified the reasoning into four categories: correct reasoning, incorrect logic, insufficient knowledge, and incorrect reading comprehension. Representative examples are provided in Appendix of Mostajabdaveh et al. (2024a) (Figure 10). As shown in Table 4, the number of reasoning steps in the 0-shot setting is significantly higher than in the 1-shot setting. This increases the likelihood of errors in intermediate steps, explaining the lower accuracy in 0-shot reasoning. Interestingly, in 20% of 0-shot tests, the model arrived at the correct answer despite some reasoning errors. Lastly, correct reasoning strongly correlates with finding the correct answer. Please refer to Appendix of Mostajabdaveh et al. (2024a) for our analysis on the relations between reasoning error types and question categories.

Metric	0-shot	1-shot
Instances with correct answer	35.6%	33.3%
Instances where all are reasoning steps are correct	15.6%	31.1%
Incorrect reasoning, correct answer	20.0%	6.7%
Incorrect answer, correct reasoning	0.0%	4.4%
Avg. number of steps per instance	7.93	4.53
Avg. accuracy of steps per instance	0.682	0.611

Table 4: Statistics on generated reasoning steps and answers using Llama2-13b-Chat on the validation set with CoT.

- Length of ICL examples has a greater influence on accuracy than the similarity of question types. We ex-

plore the impact of ICL example selection on ORQA by running four experiments on the Llama2-13B-Chat with 1-shot CoT: (1) random ICL selection, (2) same question type, (3) similar prompt length, (4) same question type and similar prompt length. Table 5 presents the results. We found that while both the length of ICL examples and the similarity of question types improved the performance, an example of similar length had more impact.

Approach	Accuracy
Random selection	0.300
Same question type	0.313
Similar length	0.353
Similar length & same question type	0.362

Table 5: Comparison of different ICL example selection approaches.

Conclusion and Future Works

LLMs have received recognition and popularity through interfaces like ChatGPT. However, as they are being used more in high-stakes fields such as medicine and law, we must understand the reasoning behind these LLM responses. Thus, we developed the Operations Research QA (ORQA) benchmark and utilized it to evaluate the generalization and reasoning abilities of some popular pretrained LLMs within a novel and technically complex domain. We explored different prompting strategies to evaluate perspectives of reasoning using ORQA. Our results show that there is still considerable room for improvement across different LLMs with Llama3.1-405B-Instruct achieving the highest accuracy of 0.772 and human expert accuracy (on a subset) of 0.93.

We acknowledge that structured reasoning metrics, such as those in ROSCOE (Golovneva et al. 2023), reasoning graph accuracy and similarity in STREET (Ribeiro et al. 2023), and entailment tree comparisons (Dalvi et al. 2021), are promising for automating the evaluation of reasoning steps in ORQA. By making ORQA publicly available, we look forward to seeing it used as a benchmark for LLMs, especially for models trained specifically on the task of QA or OR and related fields. We also encourage experts to expand this dataset by introducing not only more OR multi-choice QA problems, but also from other technical fields similarly requiring expert-level knowledge and reasoning skills with limited publicly available text corpora.

References

- AhmadiTeshnizi, A.; Gao, W.; and Udell, M. 2024. OptiMUS: Scalable Optimization Modeling with (MI) LP Solvers and Large Language Models. In *Forty-first International Conference on Machine Learning*.
- Almazrouei, E.; Alobeidli, H.; Alshamsi, A.; Cappelli, A.; Cojocar, R.; Debbah, M.; Goffinet, E.; Heslow, D.; Launay, J.; Malaric, Q.; Noun, B.; Pannier, B.; and Penedo, G. 2023. The Falcon Series of Open Language Models. *arXiv:2311.16867*.
- Alzahrani, N.; Alyahya, H. A.; Alnumay, Y.; Alrashed, S.; Alsubaie, S.; Almushaykeh, Y.; Mirza, F.; Alotaibi, N.; Altwairesh, N.; Alowisheq, A.; et al. 2024. When benchmarks are targets: Revealing the sensitivity of large language model leaderboards. *arXiv preprint arXiv:2402.01781*.
- Arkoudas, K. 2023. GPT-4 can't reason. *arXiv preprint arXiv:2308.03762*.
- Baktash, J. A.; and Dawodi, M. 2023. Gpt-4: A review on advancements and opportunities in natural language processing. *arXiv preprint arXiv:2305.03195*.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Clark, J. H.; Choi, E.; Collins, M.; Garrette, D.; Kwiatkowski, T.; Nikolaev, V.; and Palomaki, J. 2020. TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages. *Transactions of the Association for Computational Linguistics*, 8: 454–470.
- Dalvi, B.; Jansen, P.; Tafjord, O.; Xie, Z.; Smith, H.; Pipatanangkura, L.; and Clark, P. 2021. Explaining Answers with Entailment Trees. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 7358–7370. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.
- Dziri, N.; Lu, X.; Sclar, M.; Li, X. L.; Jiang, L.; Lin, B. Y.; Welleck, S.; West, P.; Bhagavatula, C.; Le Bras, R.; et al. 2024. Faith and fate: Limits of transformers on compositionality. *Advances in Neural Information Processing Systems*, 36.
- Fan, Z.; Ghaddar, B.; Wang, X.; Xing, L.; Zhang, Y.; and Zhou, Z. 2024. Artificial Intelligence for Operations Research: Revolutionizing the Operations Research Process. *arXiv preprint arXiv:2401.03244*.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; and Wang, H. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Golovneva, O.; Chen, M.; Poff, S.; Corredor, M.; Zettlemoyer, L.; Fazel-Zarandi, M.; and Celikyilmaz, A. 2023. ROSCOE: A Suite of Metrics for Scoring Step-by-Step Reasoning. *arXiv:2212.07919*.
- Hillier, F. S.; and Lieberman, G. J. 2015. *Introduction to operations research*. McGraw-Hill.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. I.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; Mensch, A.; Savary, B.; Bamford, C.; Chaplot, D. S.; Casas, D. d. I.; Hanna, E. B.; Bressand, F.; et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Joshi, C. K.; Cappart, Q.; Rousseau, L.-M.; and Laurent, T. 2022. Learning the travelling salesperson problem requires rethinking generalization. *Constraints*, 27(1): 70–98.
- Kandpal, N.; Deng, H.; Roberts, A.; Wallace, E.; and Raffel, C. 2023. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, 15696–15707. PMLR.
- Khatun, A.; and Brown, D. G. 2024a. A Study on Large Language Models' Limitations in Multiple-Choice Question Answering. *arXiv preprint arXiv:2401.07955*.
- Khatun, A.; and Brown, D. G. 2024b. A Study on Large Language Models' Limitations in Multiple-Choice Question Answering. *arXiv:2401.07955*.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213.
- Kweon, S.; Kim, J.; Kwak, H.; Cha, D.; Yoon, H.; Kim, K.; Yang, J.; Won, S.; and Choi, E. 2024. EHRNoteQA: An LLM Benchmark for Real-World Clinical Practice Using Discharge Summaries. *arXiv:2402.16040*.
- Laban, P.; Kryściński, W.; Agarwal, D.; Fabbri, A. R.; Xiong, C.; Joty, S.; and Wu, C.-S. 2023. LLMs as factual reasoners: Insights from existing benchmarks and beyond. *arXiv preprint arXiv:2305.14540*.
- Lanham, T.; Chen, A.; Radhakrishnan, A.; Steiner, B.; Denison, C.; Hernandez, D.; Li, D.; Durmus, E.; Hubinger, E.; Kernion, J.; et al. 2023. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*.
- Liévin, V.; Hother, C. E.; Motzfeldt, A. G.; and Winther, O. 2024. Can large language models reason about medical questions? *Patterns*, 5(3).
- Lyu, Q.; Havaladar, S.; Stein, A.; Zhang, L.; Rao, D.; Wong, E.; Apidianaki, M.; and Callison-Burch, C. 2023. Faithful chain-of-thought reasoning. In Park, J. C.; Arase, Y.; Hu, B.; Lu, W.; Wijaya, D.; Purwarianti, A.; and Krisnadhi, A. A., eds., *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, 305–329. Nusa Dua, Bali: Association for Computational Linguistics.
- Mallen, A.; Asai, A.; Zhong, V.; Das, R.; Khashabi, D.; and Hajishirzi, H. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 9802–9822. Toronto, Canada: Association for Computational Linguistics.
- Mihaylov, T.; Clark, P.; Khot, T.; and Sabharwal, A. 2018. Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering. *arXiv:1809.02789*.
- Mostajabadeh, M.; Salman, F. S.; and Tahmasbi, N. 2022. Two dimensional guillotine cutting stock and scheduling problem in printing industry. *Computers & Operations Research*, 148: 106014.
- Mostajabadeh, M.; Yu, T. T.; Chandan Bindu Dash, S.; Ramamonjison, R.; Byusa, J. S.; Carenini, G.; Zhou, Z.; and Zhang, Y. 2024a. Evaluating LLM Reasoning in the Operations Research Domain with ORQA. *arXiv:2412.17874*.
- Mostajabadeh, M.; Yu, T. T.; Ramamonjison, R.; Carenini, G.; Zhou, Z.; and Zhang, Y. 2024b. Optimization modeling and verification from problem specifications using a multi-agent multi-stage LLM framework. *INFOR: Information Systems and Operational Research*, 62(4): 599–617.

- Petropoulos, F.; Laporte, G.; Aktas, E.; Alumur, S. A.; Archetti, C.; Ayhan, H.; Battarra, M.; Bennell, J. A.; Bourjolly, J.-M.; Boylan, J. E.; et al. 2024. Operational Research: methods and applications. *Journal of the Operational Research Society*, 75(3): 423–617.
- Poli, M.; Wang, J.; Massaroli, S.; Quesnelle, J.; Nguyen, E.; and Thomas, A. 2023. StripedHyena: Moving Beyond Transformers with Hybrid Signal Processing Models. <https://github.com/togethercomputer/stripedhyena>. Online; accessed August 15, 2024.
- Ramamonjison, R.; Li, H.; Yu, T.; He, S.; Rengan, V.; Banitalebi-Dehkordi, A.; Zhou, Z.; and Zhang, Y. 2022a. Augmenting Operations Research with Auto-Formulation of Optimization Models From Problem Descriptions. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, 29–62.
- Ramamonjison, R.; Yu, T.; Li, R.; Li, H.; Carenini, G.; Ghaddar, B.; He, S.; Mostajabdeh, M.; Banitalebi-Dehkordi, A.; Zhou, Z.; et al. 2022b. Nl4opt competition: Formulating optimization problems based on their natural language descriptions. In Ciccone, M.; Stolovitzky, G.; and Albrecht, J., eds., *Proceedings of the NeurIPS 2022 Competitions Track*, volume 220 of *Proceedings of Machine Learning Research*, 189–203. PMLR.
- Ribeiro, D.; Wang, S.; Ma, X.; Zhu, H.; Dong, R.; Kong, D.; Burger, J.; Ramos, A.; Wang, W.; Huang, Z.; Karypis, G.; Xiang, B.; and Roth, D. 2023. STREET: A Multi-Task Structured Reasoning and Explanation Benchmark. arXiv:2302.06729.
- Robinson, J.; and Wingate, D. 2023. Leveraging large language models for multiple choice question answering. In *The Eleventh International Conference on Learning Representations, ICLR*.
- Sánchez, J. M. G.; et al. 2021. Modelling in Mathematical Programming. *International Series in Operations Research and Management Science, Springer*, (978-3): 030–57250.
- Sawada, T.; Paleka, D.; Havrilla, A.; Tadepalli, P.; Vidas, P.; Kranias, A.; Nay, J. J.; Gupta, K.; and Komatsuzaki, A. 2023. ARB: Advanced Reasoning Benchmark for Large Language Models. arXiv:2307.13692.
- Shen, K.; and Kejriwal, M. 2023. An experimental study measuring the generalization of fine-tuned language representation models across commonsense reasoning benchmarks. *Expert Systems*, 40(5): e13243.
- Talmor, A.; Herzig, J.; Lourie, N.; and Berant, J. 2019. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. arXiv:1811.00937.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Truong, T. H.; Baldwin, T.; Verspoor, K.; and Cohn, T. 2023. Language models are not naysayers: an analysis of language models on negation benchmarks. In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*, 101–114.
- Valmeekam, K.; Marquez, M.; Olmo, A.; Sreedharan, S.; and Kambhampati, S. 2023. PlanBench: An Extensible Benchmark for Evaluating Large Language Models on Planning and Reasoning about Change. arXiv:2206.10498.
- Valmeekam, K.; Olmo, A.; Sreedharan, S.; and Kambhampati, S. 2022. Large language models still can't plan (a benchmark for LLMs on planning and reasoning about change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Vidal, T.; Laporte, G.; and Matl, P. 2020. A concise guide to existing and emerging vehicle routing problem variants. *European Journal of Operational Research*, 286(2): 401–416.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q.; and Zhou, D. 2023. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv:2201.11903.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Xiao, Z.; Zhang, D.; Wu, Y.; Xu, L.; Wang, Y. J.; Han, X.; Fu, X.; Zhong, T.; Zeng, J.; Song, M.; et al. 2023. Chain-of-Experts: When LLMs Meet Complex Operations Research Problems. In *The Twelfth International Conference on Learning Representations*.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K. R.; and Cao, Y. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. In *The Eleventh International Conference on Learning Representations*.
- Yoran, O.; Wolfson, T.; Bogin, B.; Katz, U.; Deutch, D.; and Berant, J. 2023. Answering Questions by Meta-Reasoning over Multiple Chains of Thought. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 5942–5966.
- Zhao, R.; Li, X.; Joty, S.; Qin, C.; and Bing, L. 2023. Verify-and-Edit: A Knowledge-Enhanced Chain-of-Thought Framework. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5823–5840.
- Zhou, H.; Gu, B.; Zou, X.; Li, Y.; Chen, S. S.; Zhou, P.; Liu, J.; Hua, Y.; Mao, C.; Wu, X.; et al. 2023. A survey of large language models in medicine: Progress, application, and challenge. *arXiv preprint arXiv:2311.05112*.
- Zhou, P.; Pujara, J.; Ren, X.; Chen, X.; Cheng, H.-T.; Le, Q. V.; Chi, E. H.; Zhou, D.; Mishra, S.; and Zheng, H. S. 2024. Self-discover: Large language models self-compose reasoning structures. *arXiv preprint arXiv:2402.03620*.