

Leveraging LLMs for Operations Research Problems

Can LLMs reason about mathematical modelling from natural language descriptions?

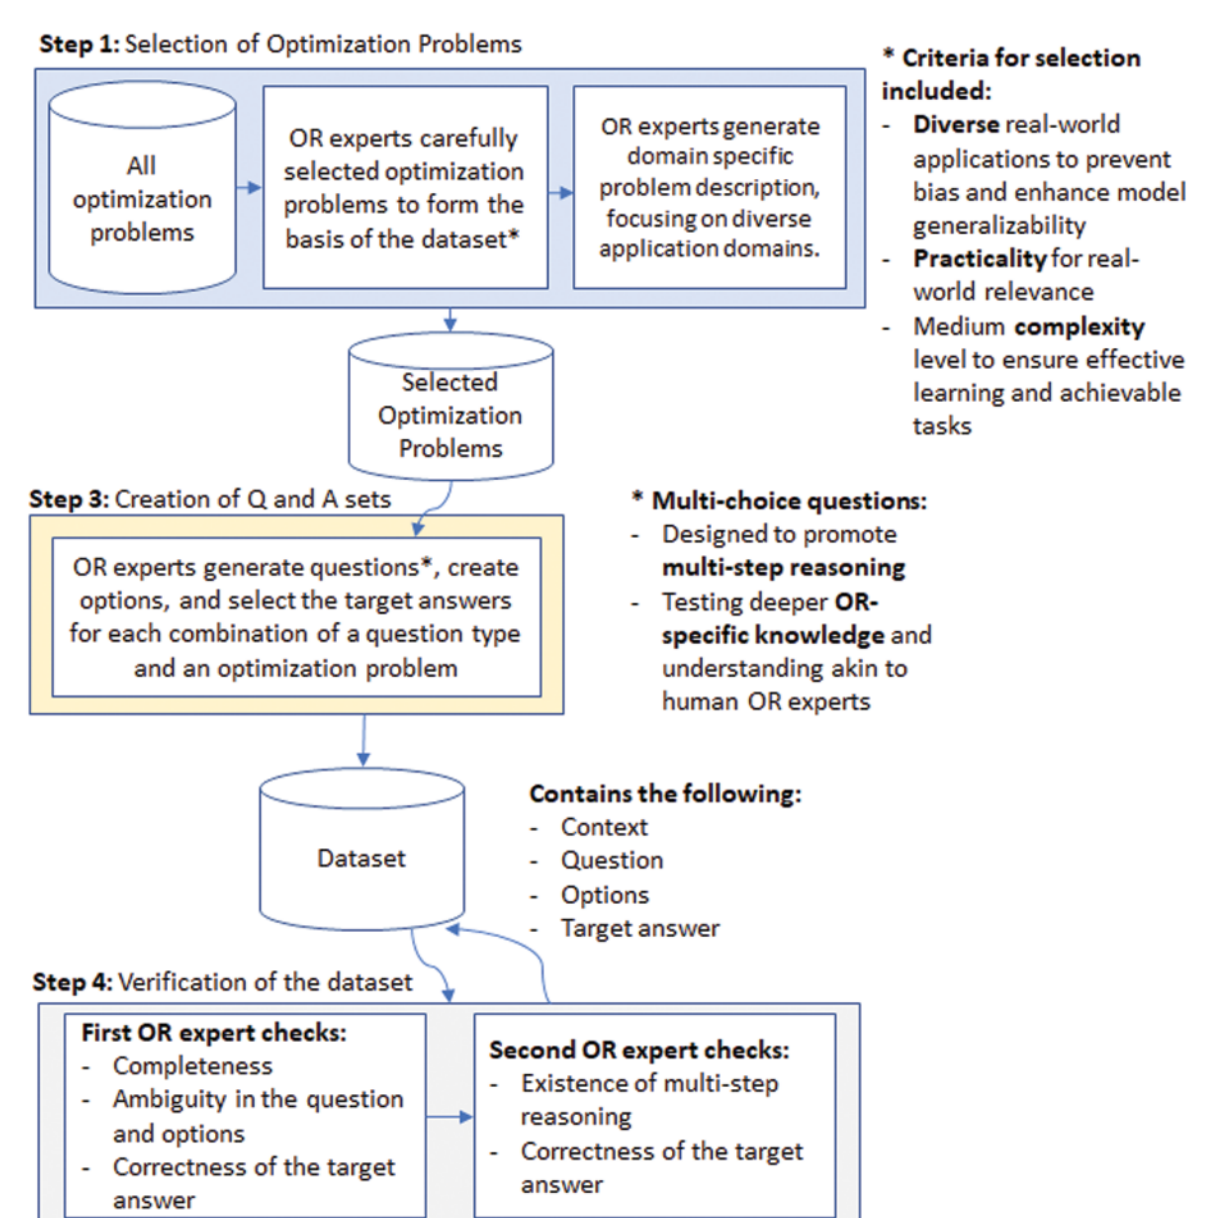
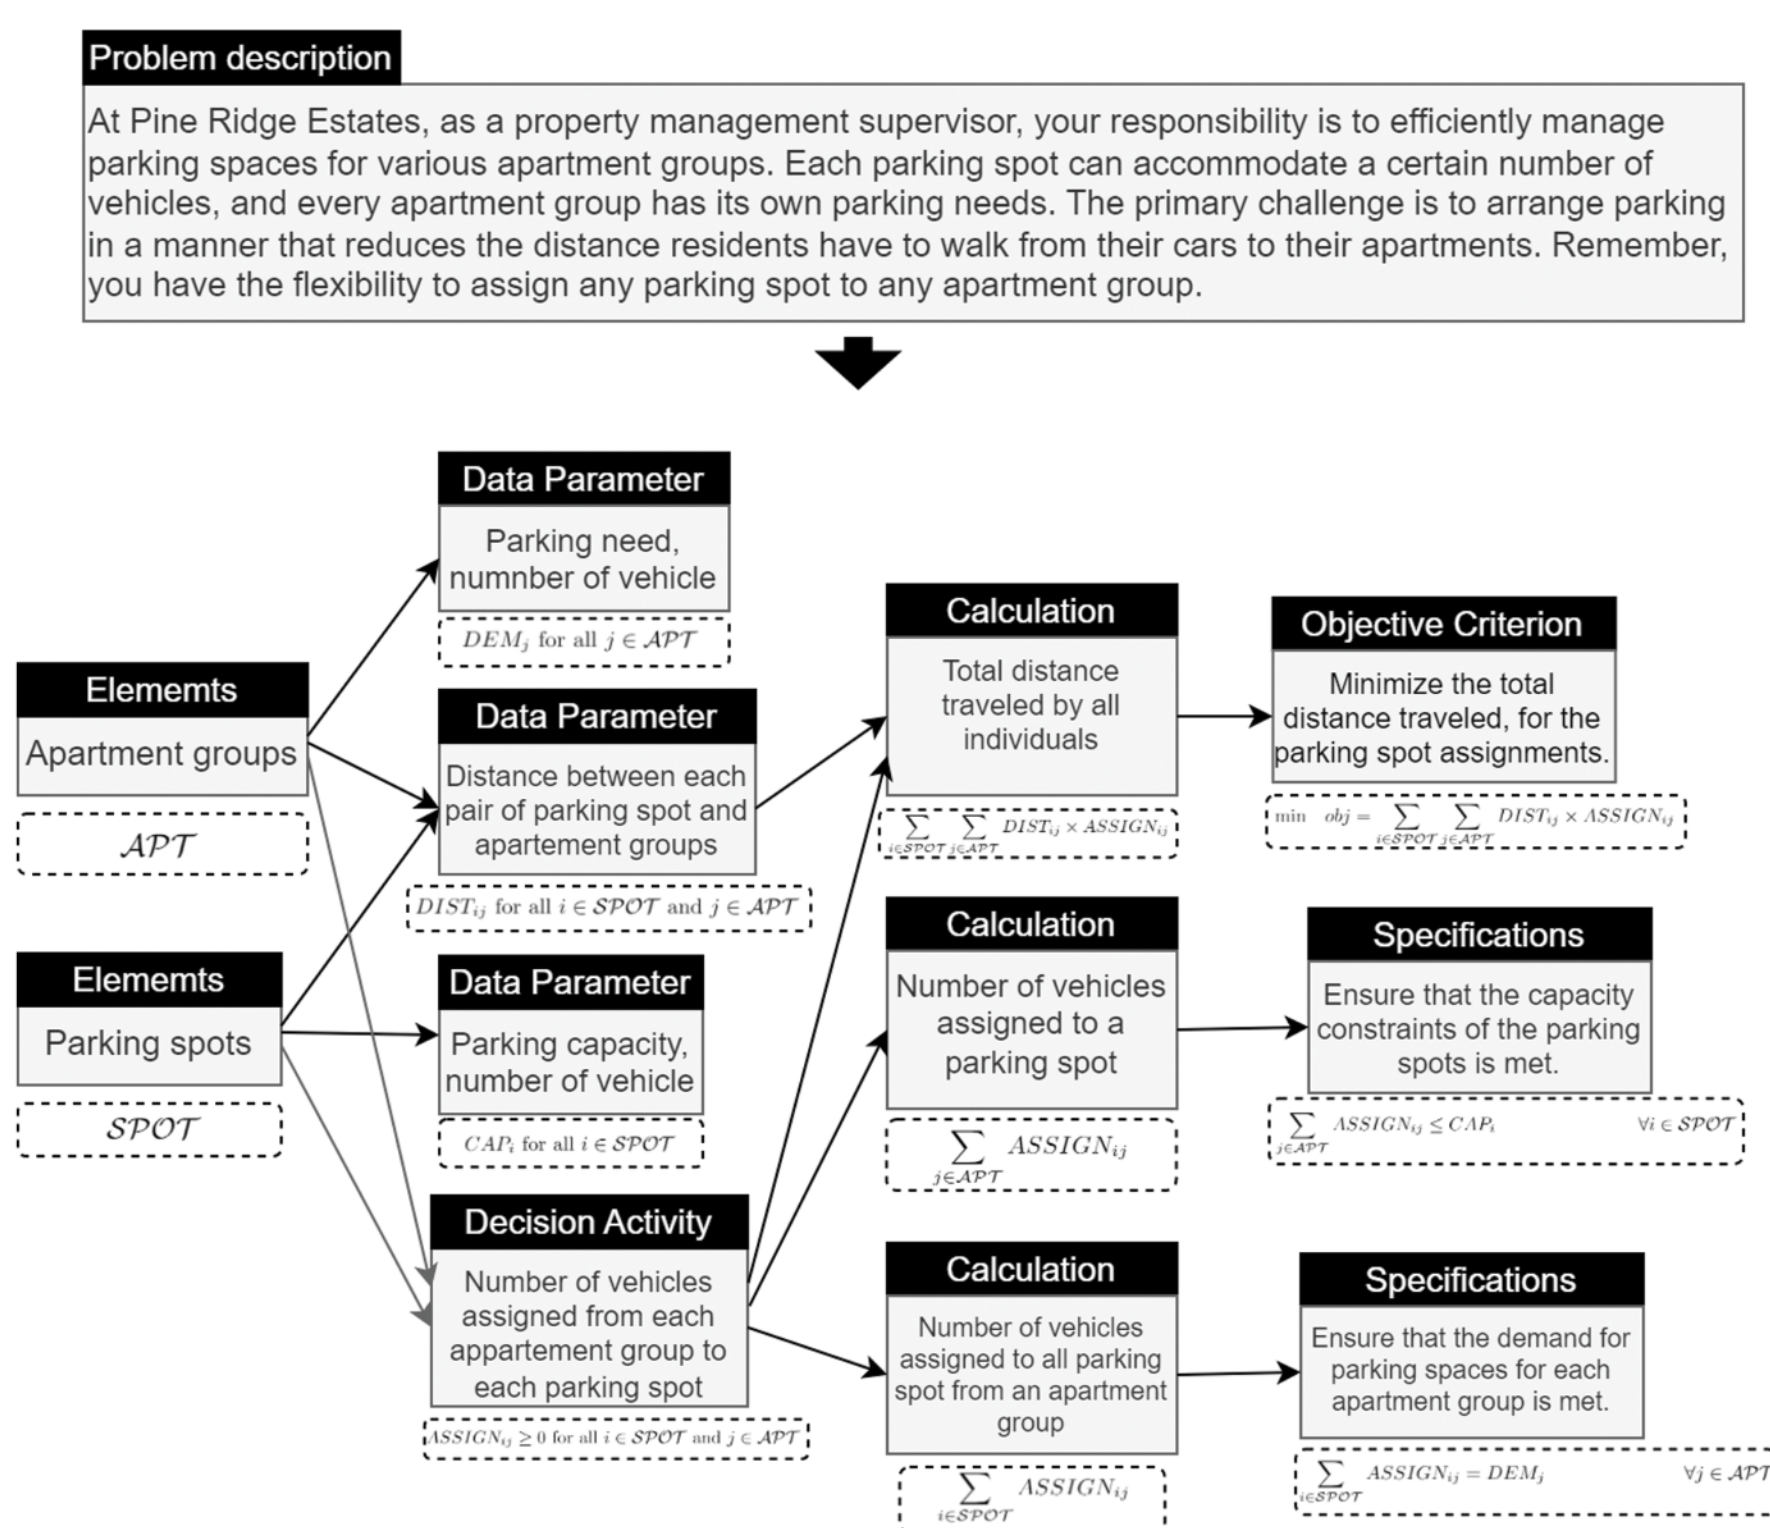
Samarendra Chandan Bindu Dash

Gerald Penn

ACADEMIC SUPERVISOR

Rindranirina Ramamonjison

INDUSTRY SUPERVISOR



PROJECT SUMMARY

This research aims to explore the capabilities of pretrained large language models to reason about mathematical modeling in the domain of operations research. We worked on multiple research questions and projects under this.

One of our major contributions to this field is creating a novel MCQ dataset with more than 1400 questions. To develop the dataset, we started with 130+ problems from various domains along with their optimization model. Following that, we created MCQ questions, which could be roughly classified into 11 distinct categories based on whether they ask about the model parameters, variables, or objective. The questions are designed in such a way that if someone has the corresponding mathematical model, they can easily answer the question. Hence, if only given the NL description and question, if the LLM can answer the question, then it suggests that the LLM can reason about the mathematical model without explicitly providing it. We ran inference with various SOTA models such as LLAMA2, Falcon, and Mistral, comparing the performance of different prompting techniques such as few-shot, chain-of-thought, etc. We will be publishing a paper on the findings soon.

Currently, we are working on developing a solution to help planners make production-scheduling plans, answering questions related to plans, and dynamically changing plans based on updated requirements. We will be exploring different LLMs and language-modeling techniques to achieve our goal. Our target is to reduce the OR planning and request-handling overhead by 60%. We aim to submit a patent.

