A note on the development of Artificially Intelligent Operational Researchers: Ideation and the way ahead

Version 1, August 2025
Santanu Banerjee, The SenseiAlgo Team

Introduction

Operations Research (OR) is an interesting branch of Computer Science and Mathematics which involves mathematically modelling problems to construct formulations for optimization. The process of developing formulations involves understanding the problem at hand deeply and scrutinizing relevant details that could affect the decision making. Mathematical Modelling of real-world problem scenarios is the most challenging part of the process. Thus advanced degrees are offered for Operations Research specialities. Due to the immense value OR personnel provide within organizations, corporates are always on the lookout to integrate highly skilled OR scientists within their teams. An interesting proposal in this regard is the development of AI-Operations Researchers for the development of mathematical models from problem statements, as well as designing Heuristics to provide fast solutions to these problems.

Multiple organizations are looking to develop/utilize AI Researchers for complex formulation/heuristic development. The development of AI Scientists have already been ideated in (Lu, et al., 2024) and (Yamada, et al., 2025). Named entity recognition for developing mathematical formulations from natural problem statements have been discussed in (Kadıoğlu, et al., 2024) and (Singirikonda, Kadioglu, & Uppuluri, 2025). However, the problem statements demonstrated in these studies are not complicated enough to replace the OR Engineers; the problems are actually quite naïve. Developing the necessary workflow and frameworks to achieve the development of elaborate formulations to complex problems could be done with multi-level AI Agents crews; investigation in this regard is necessary as breakthroughs would essentially be priceless. Since the entry barrier to ML/AI is generally lower than the entry-barrier to hardcore optimization modelling, the advent of an AI Operations Researcher would pave the way for:

- easy identification of problems which are in the scope of optimization
- problem statement development (in natural/conversational language)
- problem feature extraction from documents
- identifying constrains from the problem statement
- initializing variables for the problem (and ensuring enough variables to generate linear problems which are easier to solve)
- identification of the objective function from the problem statement

Once these are done and a formulation is ready, other AI Agents would be able to code the mathematical formulation in the specific syntax of the solver being used (current LLMs like ChatGPT, Claude, etc. are able to achieve this when provided with the formulation). Solvers could be made to run as tools to check the feasibility of the formulation (on synthetic data developed by AI Agents). Ultimately a closed feedback loop needs to be created to allow upgrading the mathematical formulation in case infeasibilities arise (dealing with infeasibilities is a common challenge in the field of optimization). Solvers like Gurobi/CPLEX have special tools like the Irreducible Inconsistent Subset (IIS) to identify the specific constraints which are infeasible and this would also help in providing valuable feedback to the initial Formulation Development AI Agent.

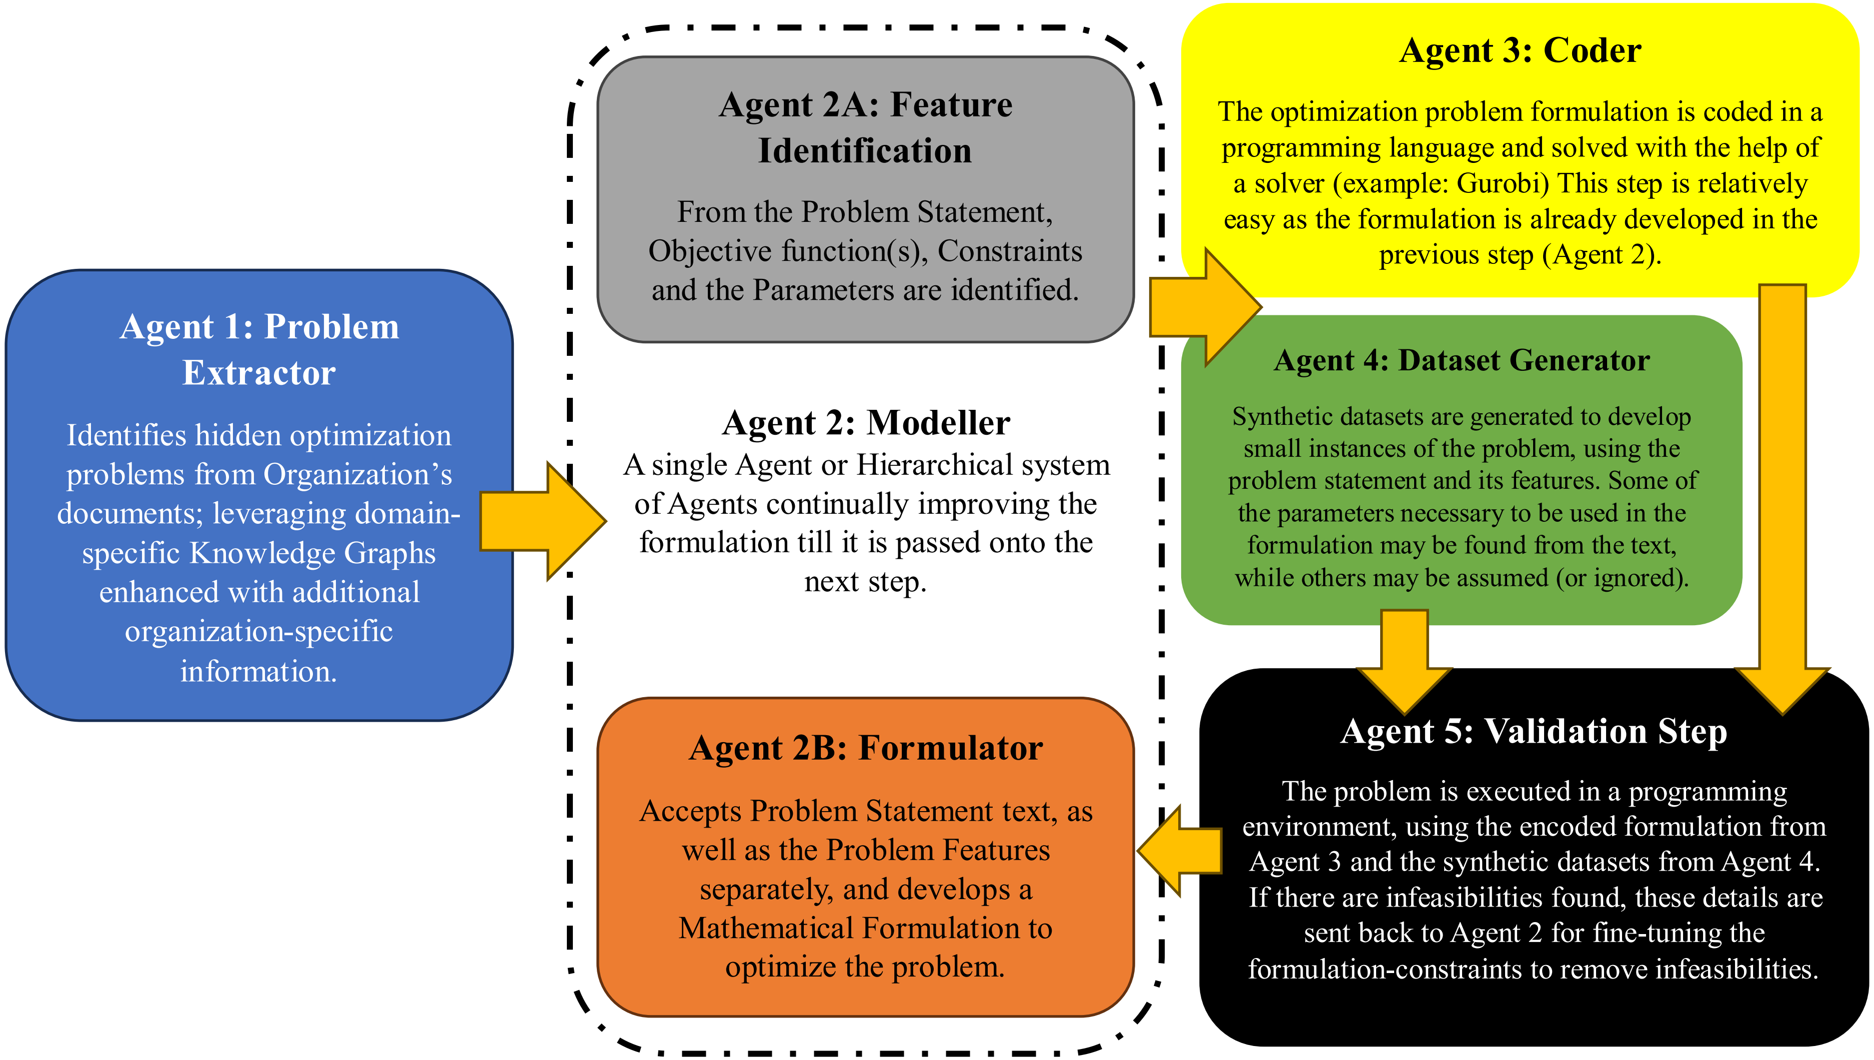

Figure 1: Flow-diagram of an AI Agentic Operations Research Engineer

Agent Usage

In this regard, the AI Agentic workflow that could be used is illustrated in Figure 1. Here, a closed-loop multi-agent framework is proposed for the development and testing of optimization models developed from industrial process documents, plant schemas, operation manuals, and other relevant knowledge sources. We discuss an Agent 1 starting out to identify an (hidden) optimization problem that could help boost sales, sustainability or operational time (as well as other KPIs). This Agent passes the problem statement to the Agent system 2; the above flow may also be warm-started by directly providing the problem statement to the Agent 2, and continuing with the subsequent flow. Agent 2A extracts additional information from the text, including objective function, constraints and parameters (NER4OPT can also be leveraged here (Kadıoğlu, et al., 2024), and so can knowledge graphs be put to use). Based on these features and keeping the problem in the context, the Agent 2B formulates the problem (mathematical programming formulation). Agent 3 develops code to convert the mathematical problem into an optimization solver-specific syntax; Agent 4 develops synthetic datasets for input to the formulation. Finally Agent 5 runs the formulation with the developed dataset(s) and tries to identify any pending issues; in case infeasibilities arise, or the optimization output is not as intended, this Agent 5 transfers control back to the Agent 2 with the additional details of the infeasibilities or other modelling concerns which further help in tweaking the optimization model's constraints.

This Agentic approach for the development of OR Engineers needs extensive testing and demonstration for large optimization problem use-cases. Additional Agents will need to be integrated subsequently which can develop Heuristics for the problem to help obtain good solutions swiftly. As of now, OR Engineers still seem to be safe from AI taking over their jobs, but their time could run out soon!

References

Kadıoğlu, S., Dakle, P. P., Uppuluri, K., Politi, R., Raghavan, P., Rallabandi, S., & Srinivasamurthy, R. (2024, November 26). Ner4Opt: named entity recognition for optimization modelling from natural language. *Constraints, 29*, 261-229. doi:https://doi.org/10.1007/s10601-024-09376-5

Lu, C., Lu, C., Lange, R. T., Foerster, J., Clune, J., & Ha, D. (2024). The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery. Retrieved from https://arxiv.org/abs/2408.06292

Singirikonda, A., Kadioglu, S., & Uppuluri, K. (2025). Text2Zinc: A Cross-Domain Dataset for Modeling Optimization and Satisfaction Problems in MiniZinc. Retrieved from https://arxiv.org/abs/2503.10642

Yamada, Y., Lange, R. T., Lu, C., Hu, S., Lu, C., Foerster, J., . . . Ha, D. (2025). The AI Scientist-v2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search. Retrieved from https://arxiv.org/abs/2504.08066