ICT with Industry 2025

20 – 25 January 2025, Lorentz Center@Oort

Under the umbrella of ICT-research Platform Netherlands (IPN) the workshop ICT with Industry 2024 was organized by NWO in collaboration with the research schools ASCI, IPA and SIKS.

Scientific

The workshop strives for direct and rapid interaction between ICT researchers and industrial partners with the following objectives:

➔ To stimulate contact between ICT research and industrial R&D.

➔ To obtain creative solutions for industrial problems and to find new approaches that could lead to such solutions.

➔ To give insight into the wide range of possibilities ICT research offers and thereby enable accelerated innovation.

➔ To enrich the PhD students' and postdocs' experience in collaborating with industry.

The ICT with Industry 2025 workshop took place between 20 and 25 January 2025. Throughout the workshop week, case studies are worked on in teams of 4-8 researchers. Specially appointed senior researchers serve as academic team leaders who liaise with case owners and guide discussions during the event.

All groups have at least partly achieved the wanted results. This was evident during the presentation on the final day and the scientific reports by the groups. The teams were invited to present their case at the pitch podium during ICT.OPEN2025.

Name organizer Kurt Driessens (UM, NL)
Name organizer Nicole van der Meulen (SURF, NL)
Name organizer Katja van Erkel (NWO, NL)
Name organizer Suzanne van der Heijden (NWO, NL)

Summary of the cases

Alliander: Data Driven Approach towards More Efficient Newton-Raphson Power Flow Calculation for Distribution Grids

Power flow (PF) calculations are fundamental to power system analysis to ensure stable and reliable grid operation. The Newton-Raphson (NR) method is commonly used for PF analysis due to its rapid convergence when initialized properly. However, as power grids operate closer to their capacity limits, ill-conditioned cases and convergence issues pose significant challenges. This work, therefore, addresses these challenges by proposing strategies to improve NR initialization, hence minimizing iterations and avoiding divergence. We explore three approaches: (i) an analytical method that estimates the basin of attraction using mathematical bounds on voltages, (ii) a data-driven model leveraging supervised learning and physics- informed neural networks (PINNs) to predict optimal initial guesses, and (iii) a reinforcement learning (RL) approach that incrementally adjusts voltages to accelerate convergence. These methods are tested on benchmark systems. This research is particularly relevant for modern power systems, where high penetration of renewables and decentralized generation require robust and scalable PF solutions. Our findings provide a pathway for more efficient real-time grid operations, which, in turn, support the transition toward smarter and more resilient electricity networks.

Index Terms—State estimation, power flow calculation, estimating basin of attraction, graph neural networks, physics-informed neural networks, reinforcement learning.

Contractuo: How to make an LLM change its mind?

Ensuring the consistency of large language model (LLM) outputs remains a critical challenge, particularly in domains requiring high reliability, such as legal compliance analysis. This study investigates methods to enhance response consistency by deploying the LLaMA model on a dedicated server and evaluating its performance under different prompting strategies. Specifically, we assess the model's ability to determine whether legal documents

comply with relevant legislation. Two approaches are compared: a baseline direct-questioning method and a structured “chain of guidance” strategy. We explored Auto CoT (Automatic Chain-of-Thought) for generating step-by-step reasoning in complex compliance tasks, but it struggled to produce coherent reasoning, likely due to insufficient legal context. Providing relevant examples through few-shot learning may enhance its performance. Additionally, we explored domain adaptation by developing an adapter fine-tuned on legal regulations, designed to improve the model’s performance in compliance analysis. While server limitations prevented full-scale deployment, the preparation of the dataset and model setup was successfully completed. To quantify response consistency, we analyze the BERT embeddings of model-generated answers using cosine similarity. The results indicate little difference in response consistency between the two prompting strategies. However, domain-adaptive fine-tuning remains a promising approach for enhancing AI-driven legal reasoning. This study contributes to efforts to improve LLM reliability, suggesting that structured prompting and domain-specific adaptation can mitigate response variability in high-stakes applications.

ING: Legacy software

Legacy systems employed in large companies are critical to maintaining business continuity but present unique challenges. In the case of ING, the NRC application, with approximately 2,000,000 lines of code written in COBOL, PL1, and JCL, suffers from incorrect, incomplete, or even missing comprehensive documentation, making it difficult for new engineers to understand its architecture. The absence of systematic documentation and the highly interconnected layers complicate both maintenance and knowledge transfer. This report investigates the utility of generative artificial intelligence (genAI), more specifically, large language models, to address these challenges by enhancing code understanding, documentation, and onboarding processes. We present COBOT: a prototype of a code understanding chatbot tailored to the NRC application. We describe the incremental design process employed to create and validate COBOT’s architecture. Finally, we derive several practical implications found when designing this chatbot, and recommendations for future work on this tool.

ING: Software refactoring

Vortex is a legacy system from ING that comprises approximately 2.5 million lines of PL/SQL code. The system lacks consistent documentation and automated tests, which poses challenges for refactoring and feature implementation. In this study, we investigated the feasibility of using large language models (LLMs) to support the translation of Vortex into a modern Java-based system. We used 10 pairs of PL/SQL-to-Java code snippets and 15 Java classes to perform code translation tests with several LLMs.

After multiple stages of refinement, our proposed pipeline uses the chain of guidance prompting technique, along with cosine similarity for identifying optimal code sample combinations and TF-IDF vectorization for feature extraction. Our results suggest that LLMs can provide syntactically correct translations and, when guided with rich domain-specific context and task-specific examples, achieve functional correctness. However, limitations such as non-determinism, scalability issues, and small dataset size constrained the findings. Future work includes clustering PL/SQL code for better sample selection, and improving similarity measures. These advancements lay the groundwork for scalable, automated solutions in modernizing large legacy systems.