Project Title:

Self-Evolving AI Agents for Industrial Process Optimization

Supervisor:

Principal Supervisor: Professor Jianxin (Kevin) Li, PhD

Summary

The proposed PhD research aims to develop trustworthy self-evolving AI agents for industrial process optimization in dynamic, data-rich and uncertain environments. The project will integrate large language models, multi-agent collaboration, industrial knowledge graphs, process mining, simulation-based optimization and human-in-the-loop feedback to enable agents to analyse process data, identify bottlenecks, diagnose causes, test alternative strategies and recommend optimization actions. Its significance lies in moving beyond static industrial AI systems toward adaptive, explainable and continuously improving decision-support. The research has strong potential impact in advanced manufacturing, energy, logistics and Industry 5.0 by improving efficiency, safety, sustainability and industrial innovation.

Project Details

Industrial process optimization is a critical challenge across manufacturing, logistics, renewable energy, food production, construction, healthcare service delivery and business operations. Effective optimization requires more than prediction; it requires process understanding, causal diagnosis, scenario evaluation, risk assessment, trustworthy adaptation and continuous decision support. For example, a renewable energy process may need to optimise solar generation, battery storage, building demand, carbon reduction and financial return under changing weather, technology and market conditions. Similarly, advanced manufacturing processes need to manage machine status, material flows, energy use, workforce allocation, product quality and safety requirements.

Traditional methods such as mathematical programming, rule-based scheduling, statistical process control and process mining are useful, but they often assume stable process structures and predefined optimization objectives. In practice, industrial environments are uncertain and continuously evolving. Process requirements change, new disturbances emerge, tools and data sources are updated, and previous optimization strategies may become outdated. Recent advances in self-evolving AI agents provide a new opportunity to move beyond static industrial analytics. These agents can interact with environments, use tools, collaborate with other agents, accumulate experience, refine workflows and improve their internal components over time.

However, current agentic AI systems are not yet sufficiently reliable for industrial optimization. They often lack structured industrial knowledge, quantitative optimization capability, lifecycle management, version control, safe update mechanisms, reusable domain-specific skills, long-term memory, self-correction and evolution-aware evaluation. In high-stakes industrial contexts, uncontrolled self-modification may also introduce safety, compliance and operational risks. Therefore, this project will focus on developing a trustworthy and auditable self-evolving AI agent framework that can adapt to changing industrial conditions while preserving safety,

performance and reliability.

This research will integrate heterogeneous industrial data, including event logs, production records, sensor readings, maintenance notes, quality reports, energy usage data, workforce schedules, supply-chain information, simulation outputs and business constraints. A dynamic industrial process knowledge graph will be developed to represent activities, resources, machines, workers, materials, products, risks, constraints, performance indicators, tools, agent skills and optimization actions. This knowledge graph will act as a shared reasoning and memory layer for multiple AI agents.

Key Research Questions (RQs):

- **RQ1:** How can industrial process knowledge be represented to support self-evolving AI agent reasoning and optimization?
- **RQ2:** How can multiple AI agents collaborate to identify bottlenecks, diagnose causes and recommend optimization strategies?
- **RQ3:** How can AI agents continuously evolve their knowledge, memory, tools, subagents and optimization policies in response to changing industrial conditions?
- **RQ4:** How can simulation-based optimization be integrated with agentic AI reasoning for trustworthy industrial decision support?

Expected Outcomes:

- A trustworthy agentic AI framework for industrial process optimization.
- An industry-ready decision-support prototype.
- High-impact journal publications in AI, 5.0 Industry for advanced manufacturing.
- Establish new industry collaboration and engagement.

Selected publications for reading

1. Fang, J., Peng, Y., Zhang, X., et al. (2025). A Comprehensive Survey of Self-Evolving AI Agents: A New Paradigm Bridging Foundation Models and Lifelong Agentic Systems. [Open access via arXiv]
2. Zhang, W., Zhao, Z., Wen, H., et al. (2026). Autogenesis: A Self-Evolving Agent Protocol. [Open access via arXiv]
3. Zhang, Z., Lu, S., Qian, H., He, D., & Liu, Z. (2026). AgentFactory: A Self-Evolving Framework Through Executable Subagent Accumulation and Reuse. [Open access via arXiv]
4. Fan, Z., Jin, W., Zhang, F., et al. (2026). Evolving-RL: End-to-End Optimization of Experience-Driven Self-Evolving Capability within Agents. [Open access via arXiv]
5. Qiu, J., Qi, X., Wang, H., et al. (2025). ALITA-G: Self-Evolving Generative Agent for Agent Generation. [Open access via arXiv]
6. Guo, D., Zhou, T., Liu, D., et al. (2026). Towards Self-Evolving Agent Benchmarks: Validatable Agent Trajectory via Test-Time Exploration. [Open access via arXiv]
7. Sampath, S., & Baskaran, A. (2026). Adaptive Orchestration: Scalable Self-Evolving Multi-Agent Systems. [Open access via arXiv]