

Ontology-Guided Multi-Agent Labeling of Reactor-System Schematics for GENESYS™

Seungyeon Lee^a, Doyeong Lim^b, In Cheol Bang^{a*}

^a Department of Nuclear Engineering, Ulsan National Institute of Science and Technology (UNIST)
50 UNIST-gil, Ulju-gun, Ulsan, 44949, Republic of Korea

^b Department of Nuclear Engineering, Chosun University, Chosundae 1-gil 10, Dong-gu, Jeonnam-Gwangju, 61452, Republic of Korea

*Corresponding author: icbang@unist.ac.kr

***Keywords :** Nuclear system design, Artificial intelligence, Multi-agent system, Large language model, Ontology

1. Introduction

Artificial intelligence (AI) is reshaping the energy sector both as a growing source of electricity demand and as an emerging engineering tool. At the same time, generative AI is enabling the automated production and verification of technical artifacts, suggesting a broader role in advanced nuclear-system development.

Prominent applications of AI to nuclear design have focused on problems that can be expressed through well-defined numerical variables and performance objectives. Sobes *et al.* applied AI to optimize reactor-core geometry and reported improvements in thermal-performance metrics [1]. At the system level, Liu *et al.* represented reactor systems as heterogeneous graphs, with components as nodes and physical interactions as edges [2]. However, these approaches rely on predefined system topologies rather than deriving structured representations directly from engineering schematics.

System schematics compactly represent reactor-system anatomy, including major components and functional connections. A schematic is useful for generative system design only when it preserves the underlying system architecture. Current general-purpose text-to-image models can generate visually convincing images, but visual plausibility does not ensure compositional or engineering correctness [3]. In reactor-system schematics, such limitations may manifest as missing components and structurally inconsistent relations. Therefore, unconstrained image generation alone is insufficient for producing engineering-consistent reactor schematics.

To address this limitation, we are developing GENESYS™ (GEnenerative NuclEar SYstem design framework), in which domain knowledge constrains the representation and generation of reactor-system anatomy. This study focuses on the ontology-based data-construction stage of GENESYS™. An ontology is commonly defined as an “explicit specification of a conceptualization” [4]. In this framework, the reactor anatomy ontology serves as a machine-readable rulebook for representing reactor-system concepts and structural constraints.

Since an ontology does not provide image-level training examples, the dataset must identify the components, their image locations, and their structural relations in each schematic. Constructing such annotations entirely by hand is difficult to scale and requires nuclear-domain expertise, whereas unreviewed automatic labeling may

introduce significant errors. We therefore propose an ontology-guided, human-in-the-loop multi-agent labeling workflow that combines AI-assisted annotation, expert review, and deterministic validation. The workflow is demonstrated on an OPR1000 schematic to construct traceable visual annotations and component-relation graphs while identifying gaps in the underlying ontology.

2. Methodology

2.1 Reactor Anatomy Ontology

The reactor anatomy ontology provides the semantic schema for labeling and validating reactor schematics. It defines canonical component and flow-segment classes, named ports, media, expected directions, relation types, functional branch roles, and reactor-type-specific required or forbidden constraints. The three relation types are flow (FLOW), containment (CONTAINS), and mechanical (MECHANICAL). Image-specific evidence, such as masks, flow polylines, and coordinates, is stored in the annotation rather than in the ontology. By binding these observations to ontology-defined concepts, the annotation can be converted deterministically into a typed component-relation graph and checked for semantic and topological consistency.




Fig. 1. Overview of the multi-agent labeling workflow

2.2 Multi-Agent Labeling System

Fig. 1. summarizes the LangGraph-based multi-agent labeling workflow, in which specialized agents sequentially propose components, visible masks, and semantic relations. ChatGPT-family models perform orchestration and proposal generation, SAM 2.1 generates mask candidates, and Claude-family models serve as independent critics. Reconciled outputs are

reviewed at three human-in-the-loop gates: H1 confirms component instances and ontology gaps, H2 accepts or edits visible masks, and H3 verifies flow segments, named port bindings, and flow, containment, and mechanical relations. The approved annotation is retained as the image-level visual ground truth, while the observed graph and validation report are generated deterministically from the annotation and ontology.

3. Results and Discussion

The ontology and annotations were co-developed using one full-system OPR1000 schematic and five auxiliary PWR schematics before the formal H1–H3 approval process. During preliminary labeling, components and connections that could not be represented by the existing ontology were reviewed and used to revise its vocabulary and rules. Consequently, all 67 component instances identified in the OPR1000 schematic had canonical ontology representations before H1.

One representative refinement event illustrates this process. In an auxiliary PWR schematic, six component instances were initially unresolved. Human review determined that these instances corresponded to five missing component classes. Adding the classes and their associated ports expanded the ontology from 10 to 15 component classes and from 16 to 36 named ports. Relabeling the same schematic reduced the unresolved component-class gaps from six to zero. Repetition of this process across the PWR cases produced the ontology used for the reported OPR1000 labeling run, containing 57 component classes and 128 canonical named ports.

Table I. Summary of OPR1000 labeling outcomes

Stage	Outcome
H1 Component	67 approved, 0 ontology gaps
H2 Mask	66 accepted, 1 edited
H3 Graph	124 relations approved
Final Validation	7 errors, 3 review items

KSNP (Korean Standard Nuclear Power Plant)




Fig. 2. H3-approved relation overlay of OPR1000 schematic

After ontology alignment, the human-in-the-loop workflow was performed as summarized in Table I. The H3-approved candidate graph contained 11 flow-segment instances, 153 port occurrences, and 124 relations,

comprising 100 FLOW, 20 CONTAINS, and 4 MECHANICAL relations. The corresponding H3-approved relation overlay is shown in Fig. 2.

Although the Graph Critic reported no blocking topology issue at H3, deterministic final validation identified seven errors and three review items. The errors consisted of three exclusive-mask overlaps, one unsupported branch-role assignment, and three direct-connection-rule conflicts. The direct-connection conflicts arose because the detailed annotation preserved intermediate components, whereas earlier ontology rules assumed simplified direct connections. Mask errors were returned to relabeling, while semantic conflicts were reviewed to determine whether the annotation or the ontology and validator should be revised. The pilot therefore demonstrates an auditable process in which labeling both uses and tests the ontology, rather than providing an independent measurement of labeling accuracy or generalizability.

4. Conclusion and Future Work

This study presented an ontology-guided, human-in-the-loop labeling process as the data foundation of GENESYS™. Labeling was used not only to construct annotations but also to expose representation gaps, revise the ontology, and relabel affected schematics. Future work will broaden the versioned ontology and dataset beyond the present PWR cases to advanced-reactor systems, including heat-pipe-cooled molten-salt reactor concepts, and use the resulting typed graphs to develop system-level generative models. Ultimately, GENESYS™ could accelerate the exploration of engineering-consistent reactor concepts, support the development of original reactor technologies and independent intellectual-property portfolios, and broaden the low-carbon nuclear options available for future energy systems.

ACKNOWLEDGMENTS

This work was supported by Korea Institute of Energy Technology Evaluation and Planning (KETEP) grant funded by the Korea government (MOTIE) (RS-2024-00403194, Next-Generation Nuclear Technology Creation IP-R&D Talent (Human Resources) Development Project)

REFERENCES

- [1] V. Sobes, B. Hiscox, E. Popov, R. Archibald, C. Hauck, B. Betzler, and K. Terrani, "AI-Based Design of a Nuclear Reactor Core," *Scientific Reports*, Vol. 11, 19646, 2021
- [2] Y. Liu, F. Alsafadi, T. Mui, D. O'Grady, and R. Hu, "Development of Whole System Digital Twins for Advanced Reactors: Leveraging Graph Neural Networks and SAM Simulations," *Nuclear Technology*, Vol. 211, No. 9, pp. 2206–2223, 2025
- [3] T. Zhang, H. Lin, Z. Liu, C. Chen, and W. Zhang, "SciFlow-Bench: Evaluating Structure-Aware Scientific Diagram Generation via Inverse Parsing," *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 17747–17765, 2026
- [4] T. R. Gruber, "A Translation Approach to Portable Ontology Specifications," *Knowledge Acquisition*, Vol. 5, pp. 199–220, 1993