Poster Abstract: P&ID-to-Graph: LLM-Assisted Digitalization of HVAC Diagrams

Lu, C.J.; Walker, Shalika; Struck, Christian; Itard, L.C.M.; Saelens, Dirk

DOI
10.1145/3736425.3772104

Publication date
2025

Document Version
Final published version

Citation (APA)
Lu, C. J., Walker, S., Struck, C., Itard, L. C. M., & Saelens, D. (2025). *Poster Abstract: P&ID-to-Graph: LLM-Assisted Digitalization of HVAC Diagrams*. 296-297. Abstract from Proceedings of the 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, Golden, Colorado, United States. https://doi.org/10.1145/3736425.3772104

Poster Abstract: P&ID-to-Graph: Exploring LLM-assisted Approaches for HVAC Diagrams Digitalization

Chujie Lu
KU Leuven/TU Delft
Belgium
c.j.lu@tudelft.nl

Shalika Walker
Kropman B.V.
The Netherlands
shalika.walker@kropman.nl

Christian Struck
Saxion University of Applied Science
The Netherlands
c.struck@saxion.nl

Laure Itard
TU Delft
The Netherlands
l.c.m.itard@tudelft.nl

Dirk Saelens
KU Leuven
Belgium
dirk.saelens@kuleuven.be

Abstract

Digitalization of HVAC piping and instrumentation diagrams (P&IDs) is essential for advancing the intelligent transformation of building systems and the building services industry. This work explores Large Language Models (LLMs) for zero-shot P&ID digitization, focusing on symbol detection. Three LLM-assisted approaches are evaluated. The results show that directly applying LLMs to P&ID digitization is highly challenging. By segmenting P&IDs into local crops and pairing them with the full diagram annotated with bounding boxes to provide global context, the LLM achieves improved symbol recognition. The findings highlight both the promise of LLMs and the need for further refinement to enable reliable HVAC P&ID digitization

CCS Concepts

• **Computing methodologies** → *Knowledge representation and reasoning*; • **Applied computing** → *Engineering*.

Keywords

Piping and Instrumentation Diagram, HVAC systems, LLM, Digitalization

ACM Reference Format:
Chujie Lu, Shalika Walker, Christian Struck, Laure Itard, and Dirk Saelens. 2025. Poster Abstract: P&ID-to-Graph: Exploring LLM-assisted Approaches for HVAC Diagrams Digitalization . In *The 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BUILDSYS '25), November 19–21, 2025, Golden, CO, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3736425.3772104

1 Introduction

Although Building Information Modeling (BIM) is rapidly evolving, Heating, Ventilation, and Air Conditioning (HVAC) systems in buildings are still commonly documented through piping and instrumentation diagrams (P&IDs). However, compared to other engineering domains, the building services industry remains at a relatively low level of digitalization, where these HVAC P&IDs in most existing buildings are typically available only as scanned images or static PDF format. Converting HVAC P&IDs into machine-readable formats is therefore a critical step toward enabling building performance analysis, model-based control, and fault detection [4, 5].

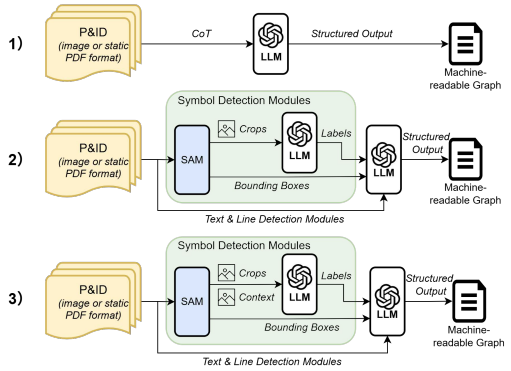

Figure 1: LLM-assisted approaches for HVAC diagrams digitalization

At present, the digitalization of HVAC P&IDs is usually carried out manually by domain experts. The process demands significant time, carries a high risk of errors, and its reliability varies with the proficiency of the expert. To address this, previous studies have explored computer vision and deep learning approaches, including Convolutional Neural Networks and Transformers, which have shown promising results [2, 6]. However, such methods depend on large annotated datasets that are expensive to create, while also suffering from imbalanced label distributions and densely connected graphical elements. In addition, HVAC P&IDs differ significantly from other process industry diagrams in their frequent use of non-standard symbols.

Recently, Large Language Models (LLMs) have demonstrated remarkable success across diverse domains and applications from natural language processing tasks to multimodal tasks, and their use is also emerging in the field of building energy systems [1]. This work explores the capabilities of LLMs for HVAC P&ID digitization, with a particular focus on symbol detection under zero-shot settings. To this end, we design and evaluate three LLM-assisted

BUILDSYS '25, Golden, CO, USA

ACM ISBN 979-8-4007-1945-5/2025/11
https://doi.org/10.1145/3736425.3772104

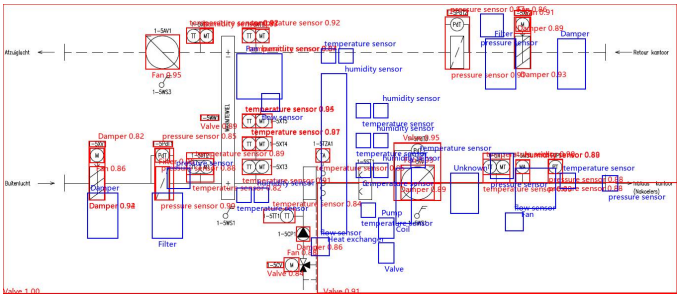

Figure 2: Visualization for zero-shot LLM-assisted digitalization. Blue boxes indicate DLD, while red boxes indicate CA-SALD

approaches that progressively incorporate segmentation and contextual information. The comparison of these approaches provides initial evidence of the opportunities and limitations of LLMs in digitalizing HVAC P&ID, and offers insights into how to enhance performance in P&ID digitization.

2 Methods

As shown in Figure 1, three LLM-assisted approaches are proposed. Their details are described as follows: (1) **Direct LLM Digitalization (DLD):** DLD directly processes P&ID diagrams in image format (e.g., PNG, JPG) as input, where the LLM applies chain-of-thought (*CoT*) reasoning to generate machine-readable graph outputs (e.g., JSON, XML, RDF). (2) **Segmentation-Assisted LLM Digitalization (SALD):** SALD first applies the Segment Anything Model (SAM) [3] to segment P&ID diagrams into cropped symbol regions (*Crops*). Each crop is then classified by the LLM into symbol categories (*Labels*). The symbol labels with their bounding boxes, and outputs from text and line detection modules, are subsequently integrated by the LLM to generate machine-readable graph outputs. (3) **Context-Aware Segmentation-Assisted LLM Digitalization (CA-SALD):** Based on SALD, CA-SALD enriches each cropped symbol region with the full P&ID diagram annotated with bounding boxes (*Context*) to enhance symbol recognition.

3 Preliminary Results

The work presents initial efforts exclusively focusing on evaluating the symbol detection module, the key component of the approaches. A P&ID of an air handling unit provided by a Dutch building services company was used to evaluate the three proposed approaches. The diagram was manually labeled using Label Studio [7] to provide ground-truth data. For DLD, GPT-5 was applied, while for SALD and CA-SALD, GPT-4o was used (with temperature set to 0).

Figure 2 shows the results of zero-shot LLM-assisted P&ID digitization. CA-SALD achieves relatively accurate bounding boxes for diagram symbols by leveraging SAM, whereas direct application of LLMs, even state-of-the-art models (e.g., GPT-5), fails to localize symbols precisely. Furthermore, DLD depends on lengthy and complex CoT prompts to handle symbol, text, and line detection as well as structured output generation. This one-step process lacks interpretability and controllability, making DLD difficult to integrate into practical workflows where human-in-the-loop oversight is essential.

A common performance evaluation metric for symbol detection is the Average Precision (AP) [6]. SALD achieves an mAP50 of 16.99%, while CA-SALD reaches 31.05%. It indicates that adding annotated full diagrams as global context can enhance recognition performance, although there remains considerable room for further improvement. More specifically, CA-SALD successfully detects all the *Fan* and *Damper* symbols, and improves the detection precision of *Filter* and sensor (*temperature sensor*, *humidity sensor* and *pressure sensor*) symbols. However, *Coil*, *Heat exchanger*, *Humidifier*, *Pump*, and *Valve* symbols remain unrecognized (AP=0). This may result from variations in symbol representation, for instance, a *Heat exchanger* may appear as either a recovery wheel or a panel, or from densely populated regions in the P&ID, where SAM struggles to produce accurate segmentation.

4 Conclusion

This exploratory study shows that directly applying LLM to digitalize HVAC P&ID is highly challenging. Integrating segmentation models improves symbol recognition, but performance remains limited without global context. By combining cropped symbols with the annotated full diagram, CA-SALD achieves substantial gains in mAP50, demonstrating that contextual information is essential for reliable P&ID digitization. Nevertheless, several symbols remain unrecognized, indicating the need for further refinement (i.e., few-shot learning, fine-tuning, and LLM agents, etc.) and larger-scale evaluation.

Acknowledgments

This work is supported by the Junior Postdoctoral Fellowship (1255626N) of Research Foundation Flanders (FWO) and the Brains4Buildings Project (MOOI32004).

References

[1] Yang Deng, Donghua Xie, Rui Liang, Jingyun Zeng, Samson Tai, and Dan Wang. 2024. BuildProg: Program Generation for Testing ML-based Building Load Forecasting models via LLM and Prompt Engineering. In *Proceedings of the 11th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation* (Hangzhou, China) *(BuildSys '24)*. Association for Computing Machinery, New York, NY, USA, 248–249. doi:10.1145/3671127.3698698

[2] Mohit Gupta, Chialing Wei, and Thomas Czerniawski. 2024. Semi-supervised symbol detection for piping and instrumentation drawings. *Automation in Construction* 159 (2024), 105260. doi:10.1016/j.autcon.2023.105260

[3] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. 2023. Segment Anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 4015–4026.

[4] Chujie Lu and Laure Itard. 2025. Leveraging LLM for P&ID-based Automated Code Generation in HVAC Fault Detection and Diagnosis. In *Proceedings of the 15th REHVA HVAC World Congress - CLIMA 2025* (Milan, Italy). Springer Nature.

[5] Chujie Lu, Ziao Wang, Martín Mosteiro-Romero, and Laure Itard. 2025. Diagnostic Bayesian network in building energy systems: Current insights, practical challenges, and future trends. *Energ Build* 341 (2025), 115845. doi:10.1016/j.enbuild.2025.115845

[6] Maximilian F. Theisen, Kenji Nishizaki Flores, Lukas Schulze Balhorn, and Artur M. Schweidtmann. 2023. Digitization of chemical process flow diagrams using deep convolutional neural networks. *Digital Chemical Engineering* 6 (2023), 100072. doi:10.1016/j.dche.2022.100072

[7] Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. 2020-2025. Label Studio: Data labeling software. https://github.com/HumanSignal/label-studio Open source software available from https://github.com/HumanSignal/label-studio.