iCon: A Modular AI Architecture for Open-Ended, Self-Supervised Learning

(Extended Abstract)

Timur Ryspekov and Yevgeniy Ussatov

Chariot Technologies Lab, Inc., USA

Keywords: open-ended learning, self-supervised learning, modular AI, intrinsic motivation

1 Introduction

As the limitations of large language models (LLMs) become increasingly apparent, there is a growing demand for novel approaches to advance current artificial intelligence (AI) systems toward Artificial General Intelligence (AGI) [1] capable of continual, self-supervised learning. In contrast to the self-learning [2] and self-adapting [3] techniques known to be implemented in LLMs, we introduce iCon, a modular "interpretable container" architecture (based on a proprietary framework), and a corresponding self-evolving, modular AI system capable of autonomously adapting to novel tasks by refining the knowledge of existing modules and dynamically expanding its functional capacity through integration of new expert components. In our research, this AI system has proven to be capable of open-ended, self-supervised learning.

2 Methodology

The iCon-based AI system (Figure 1) is composed of a hybrid ensemble of domain expert modules (generative AI models, machine learning (ML) components, and rule-based functionalities), each paired with a dedicated verification expert to ensure output accuracy. A Conductor LLM decomposes user input query into sub-tasks, routes each to the appropriate domain expert, and compiles responses into a system output. If a sub-task fails verification or no suitable expert exists, the Conductor invokes the Architect, which either finetunes the underperforming module or instantiates a new expert to address the capability gap. Additionally, an internal Oracle module generates input prompts designed to expose knowledge gaps, facilitating self-supervised performance enhancement and functional development. See References [6–11] for open-source models utilized in this research.

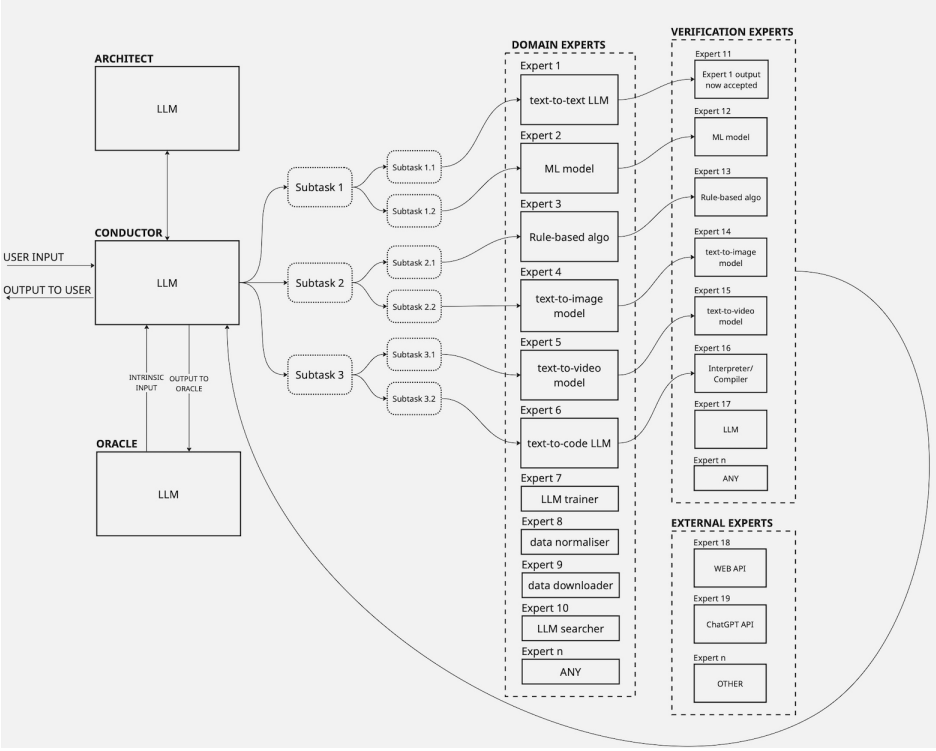

Figure 1: Example structure of iCon-based self-evolving AI system

3 Results

To analyze the system's self-learning capability, we created a baseline iCon-based system (Table 1) and asked it to teach itself school cirricula [4, 5]. Upon execution, the system autonomously expanded to a new state over the course of its runtime (Table 2).

Expert	Model	Fine-tuning	Type
Conductor Lvl2	Qwen 32B R1 (Distilled)	LoRA #1	LLM
Conductor Lvl1	Qwen 32B R1 (Distilled)	LoRA #2	LLM
Architect	Qwen 32B R1 (Distilled)	LoRA #3	LLM
Verification	Qwen 32B R1 (Distilled)	LoRA #4	LLM
Normalisation	LLaMA 3.1 8B	Fine-tuned	LLM
Memory Lvl1	LLaMA 3.2 3B	Fine-tuned	LLM
Memory Lvl2	LLaMA 4 Scout	Default	LLM
Coding	Codestral	Default	LLM
Text Formatter/Analysis	Qwen 13B R1 (Distilled)	LoRA	LLM
52 Rule-based	Custom Logic	Rule Set (52)	Rule-based

Table 1: Zero-state of iCon-based self-evolving system used in this research

Expert	Model	Fine-tuning	Type
Math Expert	Qwen 13B R1 (Distilled)	LoRA – Math	LLM
Physics Expert	Qwen 13B R1 (Distilled)	LoRA – Physics	LLM
Chemistry Expert	Qwen 13B R1 (Distilled)	LoRA – Chemistry	LLM
Biology Expert	Qwen 13B R1 (Distilled)	LoRA – Biology	LLM
History Expert	Qwen 13B R1 (Distilled)	LoRA – History	LLM
Geography Expert	Qwen 13B R1 (Distilled)	LoRA – Geography	LLM
Literature Expert	Qwen 13B R1 (Distilled)	LoRA – Literature	LLM
Computer Science Expert	Qwen 13B R1 (Distilled)	LoRA – Computer Science	LLM
Art Expert	Qwen 13B R1 (Distilled)	LoRA – Art	LLM
Music Expert	Qwen 13B R1 (Distilled)	LoRA – Music	LLM
Philosophy Expert	Qwen 13B R1 (Distilled)	LoRA – Philosophy	LLM
1024 Rule-based	Custom Logic	Rule Set (1024)	Rule-based

Table 2: Autonomously made changes from zero to end state after self-learning

4 Conclusion

The iCon-based modular AI system demonstrates self-evolving capabilities through expert refinement and autonomous functional expansion, driven by both intrinsic and extrinsic stimuli. This architecture enables dynamic adaptation, self-correction, and enhancement of internal competencies, supporting open-ended, self-supervised learning. As these systems further mature, they offer a scalable approach to developing next-generation AI.

References

[1] T. Feng, C. Jin, J. Liu, K. Zhu, H. Tu, Z. Cheng, G. Lin, and J. You. *How Far Are We From AGI: Are LLMs All We Need? Transactions on Machine Learning Research (TMLR)*, 2024. URL: https://openreview.net/forum?id=H2ZKqfNd0U.

[2] T. Ferdinan, J. Kocoń, and P. Kazienko. *Into the Unknown: Self-Learning Large Language Models*. In *Proceedings of the 2024 IEEE International Conference on Data Mining Workshops (ICDMW)*, IEEE, 2024. URL: https://sentic.net/sentire2024ferdinan.pdf.

[3] A. Zweiger, J. Pari, H. Guo, E. Akyürek, Y. Kim, and P. Agrawal, "Self-Adapting Language Models," *arXiv preprint* arXiv:2506.10943, 2025. [Online]. Available: https://arxiv.org/abs/2506.10943

[4] Duke University Libraries. *Open Education in K–12: Freely Available K–12 Teaching Resources*. Available at: https://guides.library.duke.edu/k12-OER. Accessed 12 July 2025.

[5] Institute for the Study of Knowledge Management in Education (ISKME). *OER Commons*. Available at: https://www.oercommons.org/. Accessed 12 July 2025.

Models

[6] DeepSeek AI. *DeepSeek-R1-Distill-Qwen-32B* (model card). Hugging Face. Available at: https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-32B. Accessed 12 July 2025.

[7] DeepSeek AI Unsloth AI. *DeepSeek-R1-Distill-Qwen-14B* (model card)—distilled 14 B variant that serves as the closest public release to the 13 B checkpoint. Hugging Face. Available at: https://huggingface.co/unsloth/DeepSeek-R1-Distill-Qwen-14B. Accessed 12 July 2025.

[8] Meta AI. *Llama-3.1-8B* (model card). Hugging Face. Available at: https://huggingface.co/meta-llama/Llama-3.1-8B. Accessed 12 July 2025.

[9] Meta AI. *Llama-3.2-3B* (model card). Hugging Face. Available at: https://huggingface.co/meta-llama/Llama-3.2-3B. Accessed 12 July 2025.

[10] Meta AI. *Llama-4 Scout (17B-16E)* (model card). Hugging Face. Available at: https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E. Accessed 12 July 2025.

[11] Mistral AI. *Codestral: an open-weight code-generation model*. Product announcement (blog). Available at: https://mistral.ai/news/codestral. Accessed 12 July 2025.