The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)

KnowPilot: Your Knowledge-Driven Copilot for Domain Tasks

Zekun Xi[1], Ziyan Jiang[1,3], Yujie Bao[1], Zhenqian Xu[1], Shumin Deng[2]*

[1]Zhejiang University
[2]National University of Singapore
[3]Dalian University of Technology
{xizekun2023, 231sm}@zju.edu.cn, shumin@nus.edu.sg

Abstract

Despite the rapid advancement of agents, their deployment in real world industry scenarios often encounters challenges due to a lack of domain-specific knowledge. To address this gap, we present KnowPilot: a Domain-Specific Knowledge Augmented Agent System. KnowPilot is an open-source framework that integrates task-specific priors, explicit knowledge, and experiential knowledge to enhance agent performance in specialized applications. It combines knowledge retrieval from structured repositories with a memory system capable of capturing expert experience through human–AI interaction.

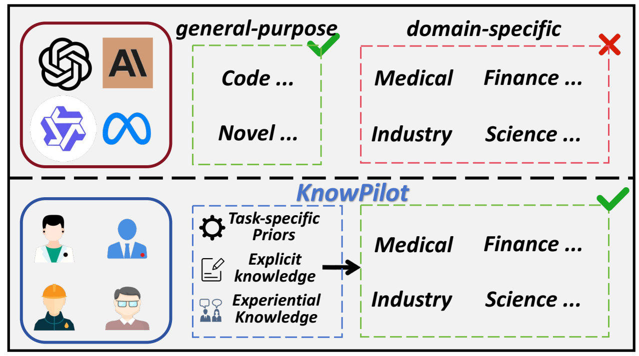

Figure 1: Current agents often perform well on generic tasks but struggle to address domain specific problems due to a lack of specialized knowledge. KnowPilot integrates task-specific priors, explicit knowledge, and experiential knowledge to support specific applications.

Introduction

Agents have demonstrated tremendous potential and value in the field of artificial intelligence (Yao et al. 2023; Nakano et al. 2022; Yue 2025; Li et al. 2025b; Gao et al. 2024; Li et al. 2025a). From personal chatbots to automated workflows (Qiao et al. 2025; Borgeaud et al. 2022; Yang et al. 2025a), the widespread application of agents is profoundly reshaping how we process information and carry out tasks.

Despite these advances, most existing systems remain confined to general purpose applications. According to relevant studies, the accuracy of general agent systems on specialized tasks, such as medical question answering (Wang et al. 2025) or legal document analysis(Barron et al. 2025; Cui et al. 2024), remains substantially lower than their performance on general applications like writing or coding, highlighting a gap in domain specific expertise (Guha et al. 2023; Yang et al. 2025b; Xi et al. 2025; Kim et al. 2024).

The root cause of this gap lies in the absence of domain specific knowledge (Song et al. 2025). Injecting domain specific knowledge into agents faces a unique dual challenge. (Abu-Salih 2021; Guha et al. 2023; Yang et al. 2025b; Xi et al. 2025). Firstly, a vast amount of critical domain knowledge exists in private, non-public databases, such as internal corporate documents, reports, and databases. Secondly, and more critically, much of the most valuable domain knowledge does not exist in a structured text format but is embodied as **the tacit knowledge of experts** (Yue 2025). This experience is deeply embedded in the daily work and decision making of experts and can only be captured and refined through repeated, in-depth interaction and dialogue.

To address these challenges, we propose KnowPilot[1], a framework designed to systematically endow large language models with domain specific knowledge by combining textual knowledge injection with interactive experience learning, thereby building more professional generative domain specific agents. Our core idea is that the model **should not only inject static, text-based knowledge but also learn dynamic, tacit, practical experiential knowledge by simulating dialogues with experts, transforming these valuable interactive processes into reusable knowledge assets.**

KnowPilot enables the agent to go beyond mastering textbook knowledge, allowing it to "think" and "interact" as a domain expert would, thus providing powerful and dependable support for tackling intricate, high-stakes challenges in real-world scenarios. We have open-sourced the project code and will provide long-term maintenance.

Design and Implementation

As shown in Figure 2(a), the fundamental concept of KnowPilot[2] is the fusion of three heterogeneous knowledge.

*Corresponding Author.

[1]Video: https://zjunlp.github.io/project/KnowPilot/video.
[2]Project: https://github.com/XeeKee/KnowPilot.

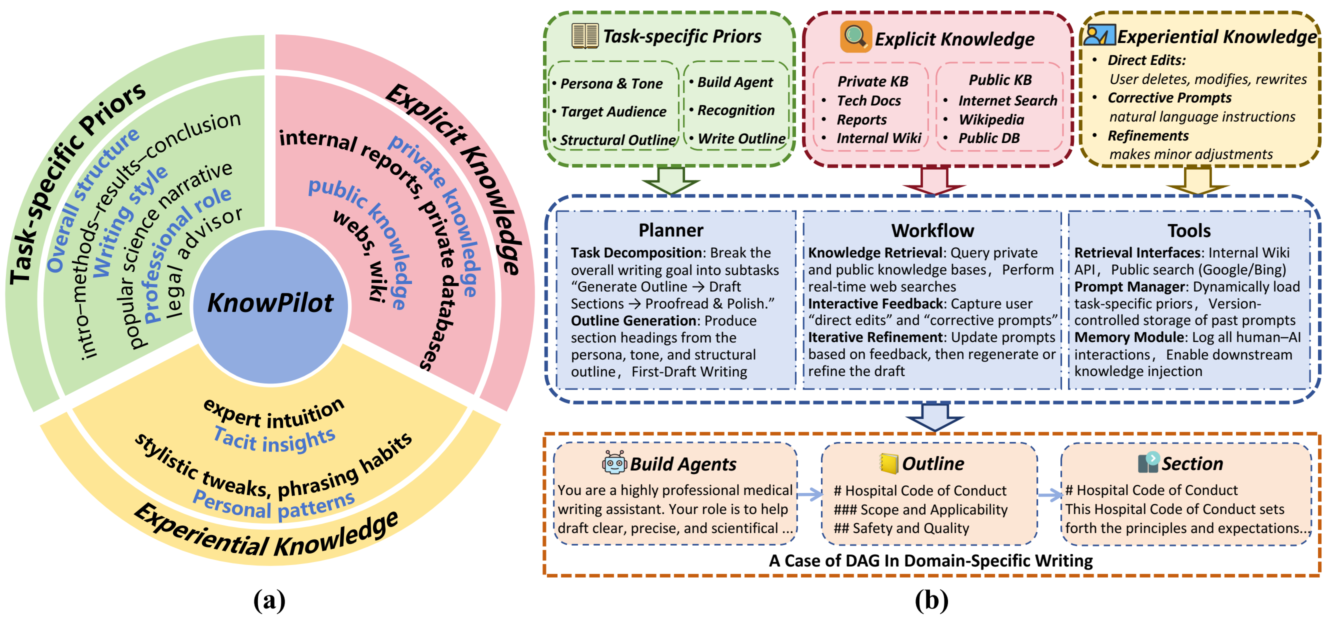

Figure 2: The overall architecture of KnowPilot. This framework integrates three types of heterogeneous knowledge sources: task-specific priors, experiential knowledge, and explicit knowledge.

Task-specific Priors. The first step of KnowPilot is to incorporate task-specific priors, defining the tone and structure to align the output with user objectives. The system generates a configuration file based on user requirements, specifying persona and style, and then produces a structural outline.

Explicit Knowledge. After incorporating task-specific priors, the second step of KnowPilot is to prepare explicit knowledge (Nonaka 1994; Li et al. 2025b; Gao et al. 2024; Li et al. 2025a). Explicit knowledge refers to information that can be clearly documented and efficiently accessed to ensure accuracy and professionalism.

Specifically, KnowPilot draws on two main sources: private knowledge bases (e.g., enterprise documents, project reports, internal wikis) and open-domain resources (e.g., Wikipedia or real-time web search) (Reimers and Gurevych 2019; Jiang et al. 2024).

Experiential Knowledge. The value of an expert lies not only in the possession of facts, but more importantly in how those facts are applied, organized, and articulated. In human–AI interaction, users often revise outputs due to issues of logic, style, or failure to align with domain best practices, and these very revisions constitute the system's most valuable experiential data.

To capture and leverage such data, KnowPilot implements a meticulous recording mechanism to track every meaningful user intervention. When users interact with the drafts generated by the model, the system records the following types of behaviors:

Direct Edits: When users delete or revise model-generated sentences or paragraphs, KnowPilot records the changes to capture wording preferences, enabling more accurate future outputs.

Corrective Prompts: When users give instructions such as "make this section more objective" or "add a counter argument," KnowPilot logs both the prompts and the before-and-after text, helping correct model misunderstandings and improve reliability (Schick et al. 2023; Nakano et al. 2022).

Refinements: After accepting a draft, users may request small tweaks to words or phrases to enhance professionalism. KnowPilot records both versions and learns from the contrasts to better match the user's style.

Knowledge Fusion. As shown in Figure 2(b), KnowPilot guides the generation of high quality content by fusing three types of knowledge.

The system loads the user's configuration file, incorporates past experiential knowledge to build context, and generates an editable outline. It then retrieves relevant information from private and open domain knowledge bases and progressively composes the content. Users can adjust or issue new instructions at any time, and all interactions are recorded as experiential knowledge to improve the efficiency and alignment of future tasks.

Conclusion

In this work, we presented KnowPilot, a human AI collaborative writing system. KnowPilot supports local deployment of large language models via vLLM and Docker, and can operate entirely on private knowledge bases.

Acknowledgments

This work was supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2023) (Grant A-8001996-00-00) and the Academic Research Fund Tier 1 (FY2025) (Grant T1 251RES2507).

References

Abu-Salih, B. 2021. Domain-specific knowledge graphs: A survey. *Journal of Network and Computer Applications*, 185: 103076.

Barron, R. C.; Eren, M. E.; Serafimova, O. M.; Matuszek, C.; and Alexandrov, B. S. 2025. Bridging Legal Knowledge and AI: Retrieval-Augmented Generation with Vector Stores, Knowledge Graphs, and Hierarchical Non-negative Matrix Factorization. arXiv:2502.20364.

Borgeaud, S.; Mensch, A.; Hoffmann, J.; Cai, T.; Rutherford, E.; Millican, K.; van den Driessche, G.; Lespiau, J.-B.; Damoc, B.; Clark, A.; de Las Casas, D.; Guy, A.; Menick, J.; Ring, R.; Hennigan, T.; Huang, S.; Maggiore, L.; Jones, C.; Cassirer, A.; Brock, A.; Paganini, M.; Irving, G.; Vinyals, O.; Osindero, S.; Simonyan, K.; Rae, J. W.; Elsen, E.; and Sifre, L. 2022. Improving language models by retrieving from trillions of tokens. arXiv:2112.04426.

Cui, J.; Ning, M.; Li, Z.; Chen, B.; Yan, Y.; Li, H.; Ling, B.; Tian, Y.; and Yuan, L. 2024. Chatlaw: A Multi-Agent Collaborative Legal Assistant with Knowledge Graph Enhanced Mixture-of-Experts Large Language Model. arXiv:2306.16092.

Gao, Y.; Xiong, Y.; Wang, M.; and Wang, H. 2024. Modular RAG: Transforming RAG Systems into LEGO-like Reconfigurable Frameworks. arXiv:2407.21059.

Guha, N.; Nyarko, J.; Ho, D. E.; Ré, C.; Chilton, A.; Narayana, A.; Chohlas-Wood, A.; Peters, A.; Waldon, B.; Rockmore, D. N.; Zambrano, D.; Talisman, D.; Hoque, E.; Surani, F.; Fagan, F.; Sarfaty, G.; Dickinson, G. M.; Porat, H.; Hegland, J.; Wu, J.; Nudell, J.; Niklaus, J.; Nay, J.; Choi, J. H.; Tobia, K.; Hagan, M.; Ma, M.; Livermore, M.; Rasumov-Rahe, N.; Holzenberger, N.; Kolt, N.; Henderson, P.; Rehaag, S.; Goel, S.; Gao, S.; Williams, S.; Gandhi, S.; Zur, T.; Iyer, V.; and Li, Z. 2023. LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. arXiv:2308.11462.

Jiang, Y.; Shao, Y.; Ma, D.; Semnani, S. J.; and Lam, M. S. 2024. Into the Unknown Unknowns: Engaged Human Learning through Participation in Language Model Agent Conversations. arXiv:2408.15232.

Kim, S.; Suk, J.; Longpre, S.; Lin, B. Y.; Shin, J.; Welleck, S.; Neubig, G.; Lee, M.; Lee, K.; and Seo, M. 2024. Prometheus 2: An Open Source Language Model Specialized in Evaluating Other Language Models. arXiv:2405.01535.

Li, Y.; Li, Y.; Wang, X.; Jiang, Y.; Zhang, Z.; Zheng, X.; Wang, H.; Zheng, H.; Huang, F.; Zhou, J.; and Yu, P. S. 2025a. Benchmarking Multimodal Retrieval Augmented Generation with Dynamic VQA Dataset and Self-adaptive Planning Agent. In *ICLR*. OpenReview.net.

Li, Y.; Shen, X.; Yao, X.; Ding, X.; Miao, Y.; Krishnan, R.; and Padman, R. 2025b. Beyond Single-Turn: A Survey on Multi-Turn Interactions with Large Language Models. arXiv:2504.04717.

Nakano, R.; Hilton, J.; Balaji, S.; Wu, J.; Ouyang, L.; Kim, C.; Hesse, C.; Jain, S.; Kosaraju, V.; Saunders, W.; Jiang, X.; Cobbe, K.; Eloundou, T.; Krueger, G.; Button, K.; Knight, M.; Chess, B.; and Schulman, J. 2022. WebGPT: Browser-assisted question-answering with human feedback. arXiv:2112.09332.

Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. *Organization science*, 5(1): 14–37.

Qiao, S.; Fang, R.; Qiu, Z.; Wang, X.; Zhang, N.; Jiang, Y.; Xie, P.; Huang, F.; and Chen, H. 2025. Benchmarking Agentic Workflow Generation. arXiv:2410.07869.

Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *EMNLP/IJCNLP (1)*, 3980–3990. Association for Computational Linguistics.

Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Zettlemoyer, L.; Cancedda, N.; and Scialom, T. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. arXiv:2302.04761.

Song, Z.; Yan, B.; Liu, Y.; Fang, M.; Li, M.; Yan, R.; and Chen, X. 2025. Injecting Domain-Specific Knowledge into Large Language Models: A Comprehensive Survey. arXiv:2502.10708.

Wang, W.; Ma, Z.; Wang, Z.; Wu, C.; Ji, J.; Chen, W.; Li, X.; and Yuan, Y. 2025. A Survey of LLM-based Agents in Medicine: How far are we from Baymax? arXiv:2502.11211.

Xi, Z.; Yin, W.; Fang, J.; Wu, J.; Fang, R.; Zhang, N.; Yong, J.; Xie, P.; Huang, F.; and Chen, H. 2025. OmniThink: Expanding Knowledge Boundaries in Machine Writing through Thinking. arXiv:2501.09751.

Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; Zheng, C.; Liu, D.; Zhou, F.; Huang, F.; Hu, F.; Ge, H.; Wei, H.; Lin, H.; Tang, J.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Zhou, J.; Lin, J.; Dang, K.; Bao, K.; Yang, K.; Yu, L.; Deng, L.; Li, M.; Xue, M.; Li, M.; Zhang, P.; Wang, P.; Zhu, Q.; Men, R.; Gao, R.; Liu, S.; Luo, S.; Li, T.; Tang, T.; Yin, W.; Ren, X.; Wang, X.; Zhang, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Wang, Z.; Cui, Z.; Zhang, Z.; Zhou, Z.; and Qiu, Z. 2025a. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*.

Yang, H.; Chen, H.; Guo, H.; Chen, Y.; Lin, C.-S.; Hu, S.; Hu, J.; Wu, X.; and Wang, X. 2025b. LLM-MedQA: Enhancing Medical Question Answering through Case Studies in Large Language Models. arXiv:2501.05464.

Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. arXiv:2210.03629.

Yue, M. 2025. A Survey of Large Language Model Agents for Question Answering. arXiv:2503.19213.