

Research Interest. Large Language Model (LLM) powered agents have demonstrated immense promise in automating many complex workflows, such as coding, literature review, and data analysis. However, their applicability to real-world tasks remains limited, with a core bottleneck being their lack of adaptability to new environments, diverse data formats, and extensive toolsets. Broadly, my research interest is to address this bottleneck by **designing agentic frameworks that generalize across tasks**, focusing on **efficient self-improvement** to minimize manual engineering and **robust task-scale scaling** to handle challenging, large-scale data reasoning tasks.

Efficient Self-Improvement. During my internship on ServiceNow’s CoreLLM research team, I observed that the development of state-of-the-art agentic systems relied heavily on manual, labor-intensive engineering efforts. Inspired by this observation, I wondered if agentic system architectures could be autonomously optimized. In **Agentic Reasoning Module** [2], I proposed a framework where an evolutionary tree search-guided meta-agent iteratively proposes and reviews novel architectures. We demonstrate that these automatically designed agentic systems are powerful and generalize effectively, achieving state-of-the-art performance across math, science, and commonsense reasoning benchmarks, and across three different backbone LLMs. Notably, these architectures generalized without requiring domain or model specific re-optimization. This work resulted in a first-author publication which I presented at the **NeurIPS 2025 Math-AI Workshop** and is currently under review at **ICLR 2026**.

With the promise of automatic agentic system design demonstrated on textual reasoning tasks, my next target was to extend this paradigm to automate the design of heterogeneous agentic systems capable of reasoning over multiple data sources of various modalities, including text, charts, tables, long documents, and web data. In this problem setting, the agentic system designer must learn to design architectures that are capable of dynamic interactions with a vast array of tools, such as multi-modal models, OCR tools, web search, etc. I am currently working on this problem as a part time research scientist at ServiceNow.

Robust Task-Scale Scaling. With the capabilities of agents rapidly increasing, I am passionate about developing agentic systems that effectively utilize compute and data to address problems of larger scale. While a visiting researcher at ServiceNow, I found that current agentic frameworks for math reasoning failed to scale effectively to competition-level problems. I identified that the key limitation was the reliance on a single tool invocation per reasoning step, which frequently led to tool invocation mistakes and erroneous tool selection. To address this limitation, I proposed **Multi-TAG** [1], an agentic framework that guides the backbone LLM to invoke multiple tools at each reasoning step and aggregate their outputs to produce accurate, cross-tool verified reasoning paths. Multi-TAG significantly outperformed state-of-the-art frameworks with parallel compute scaling and achieved the largest improvements on the most difficult subset of questions. These results highlight the importance of effectively scaling agent compute for solving complex reasoning tasks. This work resulted in a first-author publication at **Findings of EMNLP 2025** and the **NeurIPS 2025 Math-AI Workshop**.

Motivated by the success of this project, I moved on to tackle a project of even greater scale: automatically documenting linguistic feature differences between non-standard English dialects and Standard American English [3]. Collaborating with Prof. Noah Smith’s *Noah’s ARK* lab and Prof. Yulia Tsvetkov’s *Tsvetshop* lab, I designed an agentic system that analyzes a massive corpus of dialectal texts in parallel to propose and verify linguistic feature differences. We found that our approach significantly outperforms non-agentic baselines, such as simple prompting of frontier models, on LLM-as-a-judge evaluations, demonstrating the potential for agentic systems to conduct novel

scientific research through large-scale data analysis. We are currently working on conducting the final human evaluation experiments and preparing the work for publication in January 2026.

Future Research Directions. In the future, as agents need to query more diverse formats of data to perform larger-scale tasks, improving multi-modal data understanding becomes critical. A key challenge that arises is how to efficiently equip agents with modality-specific data processing capabilities, such as understanding the data format and learning effective modality-specific tool invocation patterns to retrieve and process data. One direction I plan to pursue to tackle this challenge is to build systems to automatically design modular heterogeneous agentic systems, where separate agentic modules are optimized for each data modality independently and in parallel. Then, a central orchestration system can flexibly invoke these optimized components as required when performing tasks. Furthermore, as task-scale increases, agents' long-horizon planning and dynamic plan editing capabilities will need to be improved. Inspired by the recent success of reinforcement learning for training long-CoT reasoning models, I plan to explore the potential of reinforcement learning for optimizing agentic orchestration systems that learn to effectively produce and update plans and generalize across problems of different domains. By combining task-agnostic orchestrators with modality-optimized modules, we can produce powerful yet flexible agentic systems that efficiently learn to manipulate vast amounts of data. Based on the skills I have gained from my past research experiences, I am confident in my ability to effectively lead these projects, working together with the [School] research community.

Career Goals. During my PhD, I intend to continue unlocking the potential of agentic systems to solve challenging, real world problems. I believe that a PhD offers the ideal environment for me to develop my research skills, collaborate with peers with a similar vision, and contribute my ideas to the academic community. After my PhD, I plan to continue my career in academia by becoming a professor, allowing me to continue leading and advising research on building autonomous agents that are capable of tackling the world's biggest problems.

Why [School].

References

- [1] **B. Yao**, and V. Yadav, “A Toolbox, Not a Hammer – Multi-TAG: Scaling Math Reasoning with Multi-Tool Aggregation,” Findings of Conference on Empirical Methods in Natural Language Processing (EMNLP), 2025.
- [2] **B. Yao**, S. K. R. Malay, and V. Yadav, “ARM: Discovering Agentic Reasoning Modules for Generalizable Multi-Agent Systems,” The 5th Workshop on Mathematical Reasoning and AI (NeurIPS), 2025.
- [3] **B. Yao**, O. Ahia, K. Ahuja, N. Smith, Y. Tsvetkov, “Documenting Dialectal Differences Using Language Models,” In preparation.