AI as Thinking Partners: Supporting Expert Cognitive Workflow with AI

"AI is advancing rapidly, and while the current progress in areas like large language models is amazing, many are already looking toward the future." — David Jensen, DARPA

When ChatGPT appeared, it astonished the world with its ability to write, summarize, and converse in ways that felt almost human. For the first time, people could experience an AI that seemed to "understand" language. Yet alongside the excitement came anxiety: could such systems eventually replace human workers? Many companies have already begun adopting AI for tasks once performed by entry-level employees, leading to hiring freezes in those roles.

Behind the hype lies a deeper truth. Systems like ChatGPT learn from vast amounts of online text, but they still struggle to understand why humans think and act as we do. They can produce fluent language, but they cannot truly reason, weigh alternatives, or grasp the intentions behind communication. In real workplaces, experts often find such systems useful for surface-level tasks—but not for the deeper reasoning, planning, and creativity that define human expertise.

My research begins where today's AI falls short. Traditional Natural Language Processing (NLP) has focused on building models that understand and generate text, but these systems still fail to engage with the processes of human thought. My work bridges **language and cognition to develop novel algorithms and interaction frameworks that assist experts at real-world workplaces in their complex cognitive workflows**. Rather than replacing human reasoning, I design AI that augments humans —systems that learn from how people plan, reason, and create, and that can anticipate cognitive bottlenecks, scaffold difficult tasks, and adapt dynamically to expert strategies. The goal is not automation, but **human-centered amplification**: building AI that understands cognitive intent, supports deep analytical thinking, and collaborates as a true partner in expert judgment.

(1) Collecting Expert Workflows and Developing Algorithms to Cognitively Think. Language mirrors how we *think*. When we write, read, or exchange ideas, we reveal our mental processes: how we plan arguments, revise thoughts, and engage with others' perspectives. By studying these patterns, we can design AI systems that don't just mimic text, but also reflect the cognitive processes that shape human understanding.

My group studies how people think **through writing and reading**. We examine how writers plan, draft, and refine their work, and how readers focus, infer meaning, and form

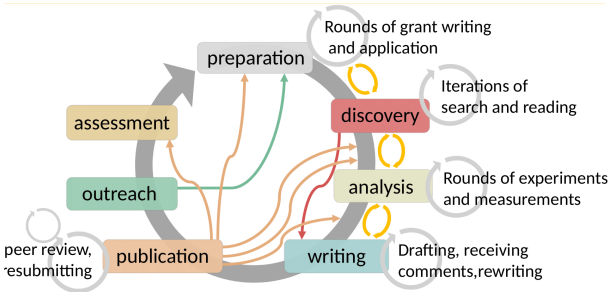

Scientific Research Cycle

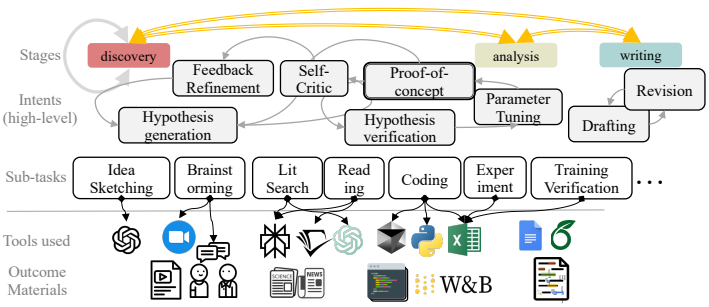

Scientific Workflow (Our collected data)

judgments as they interpret complex material. To capture these processes, we collect and analyze **workflow data** — real-world records of domain experts performing demanding, long-term intellectual tasks. In collaboration with UMN EduPsy and support from NSF, our recent study shows how scientists plan, draft, conduct experiments, and refine manuscripts over multiple months of research cycle [3], tracing how they integrate AI tools into their workflows and reflect on their use (See the collected workflow data above). We also capture fine-grained cognitive signals, such as eye-tracking, conversation logs, and writing keystrokes data, when experts read and write for specific goals, enabling us to infer underlying intentions and attention patterns.

Through collaborations with scientists, lawyers, and journalists, our lab is among the first in the world to collect such high-fidelity cognitive workflow data, offering crucial insights into complex, long-horizon tasks where current AI struggle. These data form the foundation for advancing what we call **cognitively aligned AI**—novel algorithms designed to learn human cognitive states and generate contextually aligned suggestions. Such systems can anticipate experts' next actions, estimate cognitive load across tasks, or track attention through smart interfaces to support focused reading. Moreover, working with cognitive scientists, we develop assessment frameworks of frontier AI models that assess their cognitive capabilities, like memory, planning, coherence, and adaptability, to reveal where machines still fall short of truly human-like understanding.. These efforts position our lab at the unique frontier of human-centered AI research in the world.

(2) AI as Partners in Expert Thinking.: My research extends beyond the lab to real professional environments, **developing human-AI interaction designs and building real-world assistants** that collaborate with experts to enhance reasoning, reduce cognitive burdens, and amplify creativity.

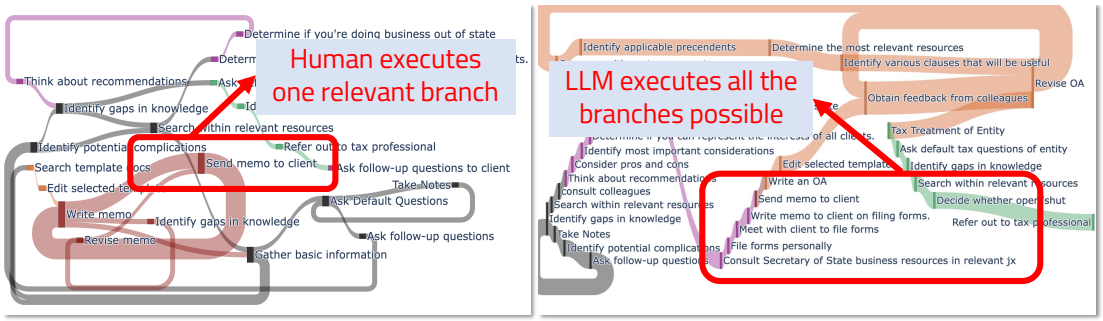

Human Legal Workflow | AI Legal Workflow

In collaboration with UMN Law School and Thomson Reuters and support by Open Philanthropy, our recent work [1] reveals **striking differences between how human and AI lawyers** in complex legal reasoning: Humans think intuitively, revise plans as cases evolve, and improvise creative solutions, while AI reason rigidly along predefined structures (See the legal workflow above). Based on these insights, we are currently building **next-gen legal assistants** that respect this balance, making up human lawyers' mistakes and identifying inconsistencies, while preserving human intuition and judgment.

In partnership with Grammarly and the Allen Institute for AI (AI2), we develop **AI-augmented reading and writing interfaces** for scientists (see figure below). Our **augmented PDF reader** [2] supports researchers navigating technical papers by analyzing how they read and what cognitive burdens they have (e.g., understanding heavy math

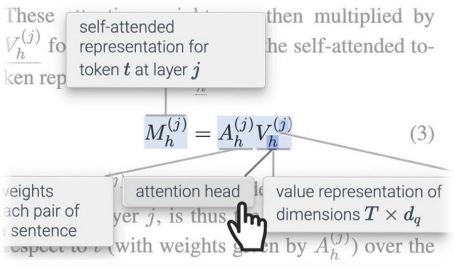

Augmented Reading Interface

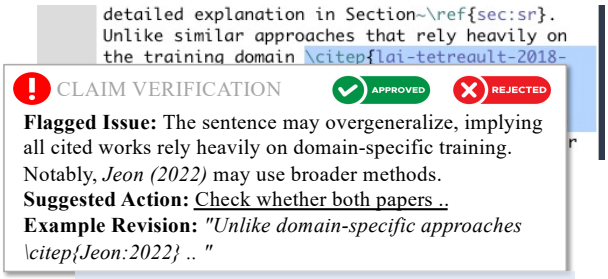

Proactive, In-Situ Writing Assistance

symbols). Our prior work [3] found human researchers' workflow, though **inefficient, embody autonomy and intellectual ownership**—insights that guide the **design of writing assistants offering cognitive scaffolding**, such as providing in-situ suggestions of missing literature, proactively detecting logical gaps, or improving readability without disrupting creative flow.

Beyond law and science, we are actively expanding expert-AI alignment to domains such as journalism, education, and medicine. Through initiatives like In2Writing (Intelligent and Interactive Writing Assistants) and AIFoW (AI for the Future of Work) workshops, we unite researchers and practitioners to reimagine how AI can empower human reasoning across fields. With AP News, Hubbard School of Journalism, and Minnesota local newsrooms, we build journalist-led AI to assist with research and editorial verification, ensuring their data and intellectual governance. In K–12 classrooms, we design AI systems that ease teachers' administrative load and student tracking, while promoting emotional and social interactions with students. Through these efforts, the Minnesota NLP group has become **an interdisciplinary hub for human-centered AI research**, creating systems that reflect human values, support cognitive growth, and strengthen our capacity to work and think together.

Future of Work. The goal of my research is to redefine what it means to work and think in the age of AI—not through replacement, but through **augmenting humans with AI**. I aim to build AI that complements human intelligence, values, and empathy, empowering people to collaborate effectively with these systems. I envision scientists exploring new ideas with AI, lawyers ensuring fairness and clarity, and educators deepening reflection and understanding. All through AI as a co-thinker that expands, not diminishes, human potential. The next chapter of AI will not be driven by bigger and better models, but by **deeper alignment with the human mind**, creating technology that is not only intelligent, but profoundly human, growing in harmony with the people it serves.

References

[1] Debarati Das, Khanh Chi Le, Ritik Sachin Parkar, Karin De Langis, Brendan Madson, Chad M. Berryman, Robin M Willis, Daniel H Moses, Brett McDonnell, Daniel Schwarcz, and Dongyeop Kang. Lawflow : Collecting and simulating lawyers' thought processes. In *Conference on Language Modeling (COLM)*, 2025.

[2] Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. Augmenting scientific papers with just-in-time, position-sensitive definitions of terms and symbols. In Proceedings of the *SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2021.

[3] Khanh Chi Le, Linghe Wang, Minhwa Lee, Ross Volkov, Luan Tuyen Chau, and Dongyeop Kang. Scholawrite: A dataset of end-to-end scholarly writing process, 2025. *under review*.