


AI as Thinking Partners: Supporting Experts' Cognitive Workflows

Dongyeop Kang

I aim to build human-centric language technologies that cognitively align human and machine thinking, advancing AI as thinking partners for experts. Today's language models still struggle with long-horizon, iterative, self-reflective reasoning, which is essential for high-fidelity decision making in real-world professional work. My work bridges language and cognition to develop novel algorithms, benchmarks, and interaction frameworks that support experts in complex, real-world cognitive workflows. What differentiates my agenda is its training signal: models learn not only from the final text people produce, but from the cognitive processes behind it—keystroke-level revisions, gaze, and expert workflow traces. I design AI that augments humans: models that learn from how people plan, reason, and create; anticipate cognitive bottlenecks; scaffold difficult tasks; and adapt dynamically to expert strategies.



Human-centric AI: learning from human cognition data toward expert-level thinking partners for people and society

The central premise of my work is that AI should extend and enhance human cognition: supporting expert reasoning, reducing cognitive burdens, and accommodating diverse individual goals. I organize my research into three tightly interconnected pillars: (1) *Cognitive Scaffolding*—collecting human cognition data (§1), building cognitively-aligned AI models (§2), and benchmarking human and machine cognition (§3); (2) *Thinking Assistants for Experts* (§4); and (3) *Societal Alignment* through pluralistic AI (§5).

§1 Collecting Human Cognition Data: Human cognition involves complex, multi-layered thinking processes. For example, writing is a non-linear and iterative cognitive process. By analyzing human writing-process data—such as our *ScholaWrite* corpus of keystroke-level scholarly writing [72] and iterative revision datasets [23, 63, 80]—we can gain insights into these thought processes, which can then be used to improve the planning and reasoning capabilities of AI models [91, 22]. Similarly, human perception can be studied through reading behaviors, such as eye-tracking data or explicit perception annotations, offering clues about how people perceive and process information through reading [31, 16, 35]. We are now scaling this capture to full research workflows (*SciFlow* [81]) and containerized, scientist-verified experimental tasks (*SciVeri* [2]), enabling us to enhance the cognitive capabilities of LLMs and develop AI assistants that better support human thinking.

§2 Cognitively-aligned AI Models: Cognitive alignment seeks to extend LLM capabilities to complex human behaviors, including neurosymbolic reasoning [53], task compositionality [91, 30], hierarchical planning [44, 49], abstraction, and social cognition [17]. We develop novel learning paradigms that either mimic human cognitive processes directly or incorporate cognitive data (e.g., writing, reading, expert reasoning, interaction) into training objectives. This includes self-supervised planning, structural and multi-attribute alignment, self-evolving reward models (e.g., Meta Policy Optimization [61]), cognitively-efficient modeling strategies, and self-evolving agents mid-trained on curated workflow trajectories toward human-AI co-discovery [75].

§3 Cognitive Benchmarking: Current LLM evaluation largely relies on shallow, monotonic tasks. We address this through three directions: (i) assessing both the potential [79, 101, 57, 58, 89, 56, 5] and risks [68, 13, 63, 80] of AI-generated data and AI-based evaluation, identifying issues such as cognitive bias, stylistic artifacts, and shallow synthesis; (ii) building *CogBench*, a large-scale cognitive assessment framework [20, 19] that integrates rigorous cognitive science methodologies with scalable evaluation to contrast human and machine cognition, in open collaboration with interdisciplinary researchers; and (iii) developing long-horizon behavioral evaluation for agentic systems [62, 83].

§4 Thinking Assistants for Experts: In domains such as science and law, there is a significant cognitive gap

between human experts and current LLMs. We design interactive, domain-specific AI systems that facilitate productive collaboration, adapt to expert workflows, and provide cognitively-aligned assistance. For scientists, we have developed reading tools like *Semantic Reader* [35, 78] and *SciTalk* [88], as well as writing assistants that detect overloaded symbols and discourse-level inconsistencies and promote iterative refinement [22, 72]. For legal professionals, we built *LawFlow* [14], a dataset capturing lawyers’ long-horizon planning and reasoning, and are building legal assistants on top of it.


§5 Societal Alignment via Pluralism: AI technologies must reflect diverse human perspectives to achieve societal alignment. We develop data-centric methods to detect, characterize, and augment underrepresented viewpoints [79, 96, 58, 55], and model-centric methods to encode pluralism at individual, group, and societal levels using distributional alignment and societal value modeling [56, 32, 17]. This mitigates risks of monolithic outputs and supports socially inclusive AI.

Together, these agendas form my long-term vision: *human-centric thinking partners* that perform complex, multi-step reasoning, reduce cognitive load, and reflect the diversity and values of human society. My research integrates linguistics, social sciences, and cognitive sciences, and is supported by industry and government partners such as Grammarly, Naver, Google, NSF, Cisco, Sony, Accenture, Thomson Reuters, and Open Philanthropy, with active research collaborations with Amazon, AI2, Naver, and Google. I hold affiliations with cross-college labs at UMN and collaborate widely with faculty and students in computer science, law, psychology, education, journalism, design, and medicine. I also foster cross-community collaboration through workshop organization: I co-organized the first “CtrlGen: Controllable Generative Modeling” workshop at NeurIPS 2021, and co-founded the “In2Writing: Intelligent and Interactive Writing Assistants” workshop series (co-organizer at ACL 2022 and CHI 2023; advisory board for CHI 2024), connecting the ML, HCI, NLP, and professional writing communities. I also co-founded the “Pluralistic Alignment” workshop at NeurIPS 2024 (second edition at ICML 2026) to emphasize diversity in AI.

1 Collecting Human Cognition Data

Language serves as a window into human cognition. Writing and reading are not linear acts but complex, iterative processes that reveal how people plan, reason, and perceive information. For example, a well-written manuscript reflects a writer’s long-term, multi-stage thinking, far beyond simple next-token prediction. Similarly, reading behaviors captured through eye-tracking or annotated perceptions provide insights into how individuals interpret, learn from, and engage with text. This first research direction—the data foundation of my *Cognitive Scaffolding* pillar—focuses on systematically collecting and analyzing human cognition data from both *production* (writing) and *perception* (reading) to better understand these processes and inform cognitively-aligned AI systems. We call this **cognitive scaffolding**: the systematic use of temporally grounded human cognitive traces—rather than isolated final outputs—as training signals.

1.1 Human Writing as Process Data




Left: human–AI collaborative iterative text revision [23, 60]; right: writing intents in scholarly writing [72]

Human writing involves richer and more deliberate processes than next-token generation: drafting, revising, planning, proofreading, and synthesizing knowledge—often collaboratively. Domain-specific writing, such as drafting scientific manuscripts, adds further challenges, requiring sustained practice, peer interaction, and integration of prior literature. We first examine the *iterative nature of human text revision*. We have collected revision data across domains (*IteraTeR* [23, 60]) and built systems that emulate iterative refinement until quality stabilizes (*R3* [22], *CoEdit* [91]). Human–AI collaborative editing consistently outperforms either humans or AI working alone [22, 91, 74]. Our studies show that explicitly modeling the revision process—whether through AI alone or human–AI collaboration—consistently yields higher text quality, demonstrating the effectiveness of test-time scaling. Guiding models with human revision intents (e.g., clarity, coherence) further produces revision traces that better mimic the human writing process, making these models better suited to supporting human writers. We also study how varying levels of LLM access reshape writers’ behavior [10].

Instead of optimizing isolated skills, we model the coordinated dynamics of full expert processes. In *ScholaWrite* [69, 72], we captured end-to-end scientific writing trajectories—nearly 62K keystroke-level text changes with intention annotations, collected over four months of real manuscript writing—and showed that training on human process data improves alignment and writing quality. This dataset provides a strong empirical basis for cognitively-aligned writing assistants that model the full lifecycle of text production.

Building on this, *SciFlow* [81] extends ScholaWrite’s capture to end-to-end research workflows: multi-month trajectories from ideation through revision, recording keystrokes, screen, workspace artifacts, and AI prompt-response traces. A two-project pilot is already recording, annotated under our process ontology and codebook that structures raw logs into events, spans, and episodes carrying intent labels, dependency edges, and failure-recovery outcomes; we are scaling toward 50 expert-led longitudinal projects across 10+ domains, with agent-generated and synthetic arms extending the corpus to several thousand trajectories, a third of which carry the failure and divergence signals that finished artifacts never contain. Relatedly, we find that structured constraints on sensemaking yield more novel research output [82].




(left) SciFlow [81] extends ScholaWrite’s writing workflow to the whole research cycle; (middle) both corpora are annotated under one process codebook (events, spans, episodes); (right) SciVeri [2] turns published studies into scientist-verified, step-scored task cards served through Sci-MCP.

Its companion, *SciVeri* [2], covers the breadth side: it atomizes published studies into scientist-verified, step-scored task cards and wraps domain simulators (e.g., AlphaFold, OpenFOAM) as typed tools behind our *Sci-MCP* server, so one agent can reproducibly run any community task. The platform is live, drawing on our campus-wide collection effort with 50+ faculty across 15 UMN departments, and is scaling toward 450 tasks. Together, the two corpora form an **ontology of the scientific process**—training signal for process learning (§2) and a benchmarking substrate for agents (§3)—toward a longitudinal repository of cognitive workflows (science, law, education) that allows models to learn not isolated tasks but the temporal dynamics of human reasoning.

1.2 Long-horizon Workflow Data from Domain Experts

My long-term goal is cognitively inspired AI that performs complex, multi-step tasks involving causality, abstraction, and memory while reducing human cognitive load. Such systems require multiple cognitive functions working in concert across an entire workflow (e.g., scientific processes, legal reasoning), not excellence in any single capability. I therefore collect longitudinal workflows from both human experts and LLMs, and use them to improve models’ cognitive capabilities on complex, high-fidelity tasks in knowledge-intensive domains.




Left: scholarly writing intents over time [72]; right: legal planning & execution [14]

Comparing Human and LLM Workflows. The first step is to investigate how human workflows differ from AI workflows and to design methods to align them. Our prior work on collecting long-term human workflow data (in scholarly writing [72] and legal processing [14]) has revealed clear distinctions between human and AI processes. When end-to-end writing data is used to train models to emulate human writing trajectories, the resulting systems exhibit more human-aligned dynamics and improved writing quality [72, 69]. Similarly, in legal planning and execution, we find that humans often act more spontaneously, sometimes repeating inefficient steps [14], whereas AI tends to rigidly follow pre-defined plans. We observe these contrasting behavioral patterns in our study of human- and AI-based legal workflows with practicing lawyers [93], providing key insights for designing cognitively-aligned thinking assistants that better bridge workflows between lawyers and AI agents.

User Cognitive Behaviors from Multimodal Context. To scale workflow acquisition beyond keystroke logging and annotation, we are developing multimodal workflow capture and cognitive state modeling. We built systems that extract fine-grained workflow dynamics from raw egocentric YouTube videos (e.g., programming, cooking, tutorials, market research) and infer latent intents using theory-grounded psychometrics. Our platform induces


probabilistic finite user state machines (FUSMs)—compact graphs of user actions and the latent intents driving transitions between them—from workflow video via iterative, bottom-up action and intent induction, and derives cognitive metrics from multimodal signals such as gaze, physiology, voice, and video from XR devices. We plan to train lightweight cognitive reasoning models to infer hidden states online and adapt assistance in real time; our *SERUM* framework takes a first step by extracting and refining user states from interaction data [90], and seeds the *Scientist Model* of expert judgment described in §2.



Left: SERUM’s [90] multi-pass pipeline revisits prior context across frames, building a refined user model of actions and intents for better informed proactive suggestions. Right: psychometrics from user workflow

1.3 Reading Data

Whereas writing data reveals the generative processes of human thought, reading data captures the perceptual and evaluative side of cognition—how people interpret, prioritize, and react to information. We have collected explicit annotations of stylistic understanding [31] and used them to train LLMs that align with readers’ stylistic judgments [33]. We have also gathered eye-tracking data to study how readers engage with stylistic cues, narrative content, or preference-ranking annotation tasks [16, 86, 21], showing that gaze patterns provide richer cognitive signals than explicit annotations or model-based interpretations. As shown in the accompanying table, these fine-grained gaze patterns also provide interpretable signals of user behaviors. By collecting data that closely mirrors human perception and evaluative behavior, we aim to train AI systems that more faithfully reproduce human-like decision making and perception processes [56]. These results show that perceptual signals capture evaluative cognition that is systematically absent from text-only supervision; a scaffolding framework for recovering such lost context is developed in [11].




Eye-tracking captures readers’ preference decisions and stylistic perception (left); fine-grained gaze metrics provide interpretable signals of reading behavior (right)

2 Cognitively-aligned AI Models

The goal of *cognitive alignment* is to enhance LLMs’ ability to emulate complex human behaviors—such as task compositionality, planning, abstraction, reasoning, and memory—by either (i) mimicking cognitive processes directly or (ii) incorporating human cognition data into training objectives. In line with Halliday’s Systemic Functional Linguistics (SFL)—a theory that language simultaneously encodes content, social relations, and structure [42]—these efforts address a single question: how can models internalize human-like structure across reasoning, planning, social judgment, and efficiency?


2.1 Reasoning and Compositionality

Human reasoning often bridges incomplete information through external knowledge and iterative refinement. We integrate neural and symbolic systems [53, 24, 52] to fill knowledge gaps while enhancing interpretability [46]. To test and improve compositionality, we created benchmarks and models for composite and chained tasks—including *CoEdit*, our publicly released instruction-tuned text-editing system [91], and *Chain-of-Instructions*



Language encodes content, social relations, and structure (SFL)

[30]—using compositional data augmentation and task-space densification, i.e., synthesizing training instances for unseen combinations of atomic tasks [8]. Human reasoning is often iterative, refining ideas or solutions step-by-step. We mimic this by developing models that emulate iterative reasoning patterns, showing that step-by-step improvements lead to higher-quality, more diverse outputs [23, 60, 22, 37].




Task compositionality: composite text-editing instructions [91] and chained instructions [30]

2.2 Planning

Humans achieve coherence by organizing content into a unified structure through decisions about topics, ordering, abstraction, and communication strategy. We therefore model planning as hierarchical optimization of text generation, coordinating multiple text units to guide surface realization.

Discourse-Guided Planning. Coherence emerges from plans that align sentence sequences with discourse goals, grounded in rhetorical structure theory and script theory. We develop supervised planners that incorporate discourse relations [46, 47], topical keywords [49], and social or persuasive goals [29]. Explicit plans based on relations, keywords, and goals enable more controlled, interpretable, and coherent generation through theory-driven discourse guidance [65, 77, 29].




Discourse-guided planning with structural and perception-aligned rewards

Self-Supervised Planning and Structural Alignment. Large language models rely on next-token prediction, which creates a gap in complex planning and long-horizon writing. We address this gap using alignment methods such as reinforcement learning to optimize high-level policies with human or model-based feedback. Our self-supervised role-playing framework [44] applies reinforcement learning to simulate dialogue interactions and optimize explicit communication goals. This paradigm has evolved into reinforcement learning with AI feedback, where one agent generates text and another provides structured feedback, supporting long-term policy learning. Recent work further aligns text structure with human writing through structural rewards [66, 63, 94, 103] and integrates human perception signals via dense reward alignment, assigning rewards at the span level rather than to whole responses [70].

2.3 Social Cognition

Language encodes interpersonal and societal dimensions of social cognition. Text style emerges from interacting factors such as formality, emotion, and metaphor, reflecting an author’s personality and serving specific communicative goals. Understanding these cross-style relationships is key to capturing the nuances of human communication. We built datasets for cross-style analysis, including PASTEL [45] and xSLUE [50], showing that multi-style learning outperforms single-style approaches. Certain style pairs, such as impoliteness–offense, are strongly correlated, while contradictory combinations yield less appropriate outcomes. This underexplored area remains vital for modeling complex stylistic patterns. More recently, we expanded to high-level affective states such as nostalgia [59] and skepticism [84], collecting annotated data for social cognition research. To align models with multiple social aspects, we developed methods for multi-attribute control via data balancing [15], multi-task fine-tuning [50], policy learning with dynamic reward re-weighting [17], or distillation [9].




Cross-style correlations [50]

2.4 Cognitive Efficiency

Abstraction enables humans to process knowledge efficiently, bypassing detailed reasoning when shortcuts suffice. Inspired by this, we design algorithms that optimize data usage, computation, and memory in LLM training, reducing hallucinations and improving reliability. Our work includes data-efficient methods that maintain performance with fewer training samples [58, 89]. We continue to develop cognitively inspired techniques that optimize LLMs’ *cognitive utility function*, balancing performance with resource efficiency [97].

2.5 Self-Evolving Agents with Adaptive Cognitive Scaffolding


Beyond modeling human workflows, I build agents whose own thinking mechanisms evolve over time. The core idea is to treat long-horizon reasoning as an iterative control problem: agents should revise, critique, and improve their plans and judgments using explicit meta-level feedback or self-reflection on their thinking trajectories, rather than relying on static prompts or fixed reward models.



Self-evolving mechanisms: meta policy optimization [61], adaptive judge updates [41], and thought-template evolution [37]


Our earlier work showed that explicitly modeling revision loops in writing reliably improves quality for both standalone agents and human–AI collaboration, highlighting structured test-time scaling as a principled mechanism for better outcomes [23, 60, 22]. More recently, we focus on methods that update the agent’s internal policies and judges as it gains experience. Meta Policy Optimization (MPO) [61] reframes alignment as a two-level adaptive process by dynamically updating the reward model through meta-reward learning, reducing reward hacking by allowing evaluation criteria to change as training reveals new failure modes; *Metaⁿ* generalizes this recursion to multiple meta-levels of self-improvement [64]. Complementary inference-time approaches evolve meta-prompts during deployment, selectively updating them based on self-inconsistency signals to enable continual self-improvement without full retraining [41].

A key technical thread is converting test-time experience into reusable, evolvable knowledge. Instead of compressing long rollouts into a single prompt, we store per-instance reasoning experience as reusable thought-templates, then iteratively refine them using language feedback as text gradients, enabling long-context models to retrieve and apply these templates RAG-style across new tasks [37]. When tool-use experience introduces utility constraints such as budget, latency, and compute, we also model these signals explicitly and learn a meta-utility function at test time via bandit-style optimization, improving test-time decision making under constraints (ongoing work).



Our current emphasis is scalable synthetic collection and selection of complex workflow trajectories. Evolution Fine-Tuning (EFT) [75] turns 156,731 agent evolution-search trajectories (371 tasks, 10 domains) into mid-training supervision, lifting open models by +10.22% on 22 held-out tasks, transferring discovery strategies across domains (math, science, coding), and surpassing the best human on autocorrelation tasks. Across evolving agent frameworks (OpenEvolve, Terminus 2, and OpenHands), we find that trajectory quality and diversity dominate raw scale: a small, carefully selected set of trajectories—filtered for branching behavior, refinement strategies, and informative failures—can outperform conventional large-scale Best-of-N and generic test-time training approaches. Curated trajectory data functions as mid-training scaffolding that reduces train–test mismatch, and the resulting *Finch* model family seeds the process-learning agenda below.

We explore meta-scaling frameworks (e.g., AsyncThink) that learn how to allocate branching and refinement budgets dynamically across stages of a workflow [25].



In domains with unverifiable judges and multi-step pipelines, such as data analysis and visual insight generation, fixed branching factors waste compute and miss opportunities. Our approach learns stage-specific scaling decisions, for example, increasing parallel exploration when outputs lack diversity or reallocating compute when judge scores are uninformative, yielding more strategic use of fixed compute budgets.

These threads converge in my agenda on **process learning for co-discovery**: process mid-training writes the discovery process into an open policy model (*Finch-Sci*)—imitating verified transitions, preferring recovery over abandoned branches, and reinforcing against a process reward—while a separately trained *Scientist Model* learns the expert’s judgment, forecasting the next correction (accept, repair, or redirect) and supplying an implicit process reward [66, 61]. The two co-optimize at inference in a dual-loop harness, to be released openly on our *OpenGalapagos* agent-scaffold platform [1], and expert interventions logged during deployment feed back into training (§3). Long term, I aim to unify human workflow modeling with self-evolving agent training into efficient agentic learning regimes: systems that learn from expert trajectories, accumulate test-time experience as structured memory, and continuously improve their reasoning policies and evaluation criteria.

2.6 Extension to Other Modalities


While my work primarily focuses on language, multidisciplinary collaborations have extended it to images and video, showing that multimodal models benefit from complementary cross-modal effects. Effective multimodal learning, however, requires careful design of fusion strategies and cognitive objectives such as reasoning and planning. Key results include the following: (i) vision–language models improve understanding of global connectivity and graph motif analysis in graph images [12], (ii) explicit temporal memory in video LLMs yields more coherent frames [6], (iii) aligning vision–language models with LLM feedback enhances reasoning in video and image tasks [54, 5, 4, 85], (iv) integrating multimodal representations into a unified latent space (a “Platonic” representation, where different modalities converge to a shared underlying structure) accelerates cross-modal understanding [36], and (v) LLMs can plan robotic manipulation tasks from language-based instructions [98]. Beyond multimodality, I also explore (i) autoregressive–diffusion hybrids for text generation [104], (ii) model uncertainty and calibrated abstention [39, 58, 67, 102] and generalization [3, 40], and (iii) theoretical analyses of learning frameworks [76].

3 Cognitive Benchmarking and Evaluation

Evaluation is undergoing a structural shift as LLMs move from standalone models to persistent, tool-using collaborators embedded in real workflows. In this setting, leaderboard-style testing on fixed benchmarks is insufficient: evaluation becomes a measurement engineering problem, where the goal is to produce metrics that remain valid, interpretable, and safety-relevant as models, data, and deployment contexts drift; our recent benchmarks make evaluation principles and test difficulty explicit [34, 7]. My work develops evaluation methods and infrastructure across three directions: (i) quantifying the benefits and risks of AI-generated data and LLM-based evaluation, (ii) benchmarking human and machine cognition with theory-grounded assessments, and (iii) replacing static benchmarking practices with long-horizon behavioral measurement for agentic systems [62, 83].

3.1 Potentials and Risks of AI-generated Data

As LLMs began replacing or supplementing human annotation, I studied both the acceleration they enable and the new failure modes they introduce, including opaque error provenance and artifact reinforcement. On the opportunity side, I developed cost-reducing supervision and data-expansion techniques via symbolic rules [24, 53], annotation imputation [79], and information-theoretic measures [57, 58, 89, 67], as well as disagreement-aware sampling [96]. LLMs can generate annotations, prompts, simulated dialogues, and evaluation data, accelerating the research and development process. We have leveraged LLM-based evaluations [56, 5] and multi-agent simulations [44, 8, 99] to enrich training data.




Risks of LLM-generated data and LLM as evaluators: cognitive biases [68], discourse shifts [63], shallow synthesis [80], and artifact amplification [13]

In parallel, I characterized how overreliance can create an artificial data ecosystem (*Artifactuality* [13]) that distorts measurement and behavior through cognitive biases (*CoBBLEr* [68, 51]), discourse shifts [63], shallow synthesis

[80], artifact amplification [13], and behavioral inconsistency across tasks and contexts [83, 95]. These results motivate evaluation protocols that explicitly model feedback loops between generators, judges, and downstream learners, rather than treating synthetic data and LLM judges as neutral measurement instruments.

3.2 Cognitive Assessment of LLMs


Traditional cognitive science experiments provide detailed insights into human cognition but are difficult to scale, creating a bottleneck for theory-driven data. To address this, we are developing *CogBench*¹, a benchmarking framework that connects cognitive psychology with artificial intelligence. CogBench integrates controlled human experiments with systematic evaluation of LLMs across executive control, working memory, temporal reasoning, and narrative comprehension. Although LLMs often perform well superficially, our findings reveal key cognitive limitations compared to humans. Models show a gap between capacity and control, with high working memory capacity but weak executive control and cognitive flexibility, as seen in tasks such as the Wisconsin Card Sorting Test [19]. They also rely heavily on prototypical patterns, producing inconsistent temporal and causal judgments in narratives [20, 92]. Moreover, while internal representations can detect incoherence, model outputs do not reliably distinguish coherent from incoherent narratives [18].



Limited executive functioning in LLMs [19]

By grounding AI evaluation in cognitive science and scaling it through benchmarking, CogBench characterizes the cognitive limits of current AI systems and provides the foundation for expert-facing systems where cognitive alignment can be assessed end to end.

Toward Long-horizon Agentic Evaluation. Finally, as LLMs become agents embedded in extended workflows, we move beyond static benchmarks toward behavioral measurement over long, multi-step episodes. Our recent studies show that current LLM agents are strong tool users but weak navigators of long-horizon tasks [62], and that their behavior is often inconsistent across tasks and contexts [83], motivating persistent, workflow-embedded evaluation of agentic systems. *SciVeri* [2] operationalizes this: agents are benchmarked on scientist-verified tasks under increasing scaffolding levels and scored per step against expert-defined target metrics rather than final outcomes alone. The end state is *deployed evaluation*: agents running on partner labs’ live science through a scientist-facing console, scored on a pre-registered taxonomy spanning data quality, capability, collaboration, and discovery—where a claimed discovery must clear a ladder of corpus distance, blind-expert novelty judgment, and community certification [82]—and every logged expert intervention becomes fresh training data, closing the human–AI discovery loop.



Human–AI co-discovery

4 Thinking Assistants for Experts

In knowledge-intensive domains such as scientific research and legal practice, significant cognitive gaps remain between human experts and current LLM capabilities. These gaps arise from limitations in handling complex multi-step tasks, incorporating dynamic feedback, and adapting to evolving contexts. Full automation often oversimplifies the nuanced reasoning these tasks require, leaving AI systems misaligned with expert workflows.

My objectives are twofold: (i) benchmark AI cognition on end-to-end professional tasks to identify cognitive gaps, and (ii) create collaborative platforms where experts and AI agents interact, exchange feedback, and adapt over time. We focus first on building expert support systems in two flagship domains—*science* and *law*—and extend to media domains such as advertising and journalism (§4.3).


4.1 Assistants for Scientists

Reading Assistance. In collaboration with AI2 and UC Berkeley, I contributed to the development of *Semantic Reader* [35, 78], a tool publicly deployed on Semantic Scholar that provides in-situ definitions and explanations of terms as scientists read papers [48, 38], thereby reducing cognitive load. More recently, we launched *SciTalk*, which transforms research papers into short-form videos, making knowledge more accessible to wider audiences [88].

Writing Assistance. Scientific writing is a paradigmatic case of the iterative cognitive processes described in §1. I envision the future of scientific writing as a collaborative effort between scientists and AI agents, where both iteratively co-develop high-quality research outputs.

Writing assistance must go beyond supporting the immediate writing context. It should understand the full

¹<https://minnesotanlp.github.io/CogBench-web/>



In-situ reading [35, 78] and writing assistance for scientists

workflow, capturing the writer's intent and providing in-situ support while minimizing distraction. I also focus on building interaction-centric AI systems that evolve through user feedback to better support scientific writing.

Our design principles include:

- *In-Situ Aids for Writing*: Document-level tools that detect overloaded symbols, redundant terminology, and logical inconsistencies using discourse-level modeling for real-time writing support.
- *Workflow-aware Interfaces*: Features that flag inconsistent scientific claims or overstated text and suggest revisions. For instance, a scientist drafting a related work section in Overleaf may trigger backend reasoning agents to automatically retrieve cited literature, detect errors or overstatements, and suggest improved citations with inferred rationales.
- *Human-Centric Design and Co-Learning*: Avoiding the "illusion of clarity" [87] by carefully designing interfaces that support scientists' thinking processes while reinforcing human-AI collaboration.
- *Learning from Human Workflow*: Collecting human revision and writing workflow data to build cognitively-aligned systems for naturally integrated interactions [22, 72, 74, 99].
- *Learning from Interaction*: Adaptive systems that improve via user feedback, as demonstrated in collaborative editing [22] and iterative taxonomy building [73].
- *Trustworthy Models*: With NSF support, we develop foundation models integrating reliable, cross-modal knowledge for robust writing assistance.

Other Scientific Tasks. Beyond reading and writing, scientific workflows involve cognitively demanding stages such as peer review, collaboration, experimentation, and verification. Our lab is developing specialized agents and interfaces to support different stages of scientific tasks. For instance, we created *PeerRead* [43] (NAACL 2018), the first large-scale dataset of papers and reviews from computer science conferences, enabling research on acceptance prediction and aspect-specific review generation. In collaboration with visualization researchers, we are also building systems that automatically generate effective visualizations and insights (A2P-Vis [26]) and that scaffold the design of multi-agent analysis workflows (*FlowForge* [28]), adapting interactively to scientists' needs. Through the SciFlow and SciVeri community platforms [81, 2], we are extending this expert-in-the-loop approach into deployed co-discovery with partner labs across 15 UMN departments, where scientists steer agents on their own live research with per-step accept, repair, or redirect feedback.

4.2 Assistants for Lawyers


With Open Philanthropy and UMN Law School, we developed *LawFlow* [14], a benchmark for complex legal processing tasks requiring long-horizon reasoning and planning. The dataset captures detailed workflow data from both human lawyers and AI agents, including brainstorming, planning, research, and client communication for tasks such as advising a startup.

Our findings reveal that human legal reasoning is recursive and exploratory, whereas AI workflows are typically linear and exhaustive. From these human-AI comparisons, we derive design principles for legal interfaces, including collaborative planning, task division, and workflow flagging, and we are actively developing interfaces guided by these insights. Reflections from practicing lawyers [93] further highlight that generative AI already improves junior lawyer performance, yet adoption lags because evaluating AI outputs is cognitively demanding, especially for less experienced practitioners.

These challenges create several risks, including overtrust or overreliance on AI, shortcut-driven reasoning [71], and the potential intensification of expert polarization. Junior lawyers in particular risk being sidelined without proper scaffolding. To mitigate these risks, we emphasize the need for tools that upskill junior lawyers, train them to critically validate AI outputs, and provide workflow-aware assistance that supports lawyers as thinking and co-learning partners, ultimately lowering cognitive burden rather than replacing expertise.

4.3 Advertising and Journalism

Advertising and journalism are expert domains where nuanced affective responses, not just factual accuracy, drive outcomes. In collaboration with Naver and MCAL, I contribute to developing computational metrics and



Design principles for legal human-AI workflows from LawFlow [14]


algorithms for advertising in both academia and industry. For example, in the advertising market, we model nuanced affective responses such as skepticism [84] and nostalgia [59], moving beyond sentiment-based metrics. Recently, we have also studied how advertising is shifting from keyword-based search advertising toward generative ad platforms (e.g., Perplexity, Google AI Overviews), conducting comparative research on search patterns and consumer engagement to better understand the impact of AI on the advertising ecosystem [27]. Finally, I have worked on transforming the workflows of advertising services by integrating AI technologies. For instance, we are developing ad systems that support off-policy evaluation (OPE)—estimating the value of new ad policies from logged data without live deployment—in deterministic environments such as ad auctions [100].

By capturing detailed expert reasoning, we aim to build cognitively inspired models and expert-aware interfaces that adapt to the dynamic cognitive states of professionals. This approach will extend to other domains such as education, medicine, and programming, forming a broad *Expert Benchmarking* initiative to support specialized AI systems tailored to the cognitive demands of diverse expert roles.

5 Societal Alignment via Pluralism

Our society thrives on diverse perspectives shaped by different backgrounds, identities, and cultures, yet rapid AI advances often collapse them into a generalized “average.” Because cognition itself is pluralistic, alignment is not an add-on but a core requirement of cognitively-aligned AI: we must build people-centric technologies that capture and encourage diversity, ensuring inclusivity and societal alignment.

My goal is to build socially aware AI models that represent the full spectrum of human opinions through two complementary approaches: **data-centric** (§5.1) and **model-centric** (§5.2).



Pluralistic alignment across objectives, data, and models

5.1 Data-centric Alignment

Diversity is difficult to define and capture: it manifests at individual, group, and community levels, and varies over time and context. Our “pluralistic representation” approach encodes disagreement and perspective variation directly into datasets. We develop methods to detect, characterize, and augment marginal viewpoints by (i) predicting missing annotations with collaborative filtering over annotations [79], (ii) analyzing annotation properties via learning dynamics [58], (iii) quantifying the disagreement level using demographic factors [96], and (iv) collecting perspectives from language learners to diversify model predictions [101, 55].

5.2 Model-centric Alignment

As NLP tasks grow more subjective, labels often shift from discrete to continuous values. Models that ignore this fluidity risk excluding certain groups. We design pluralistic alignment methods at multiple granularity levels: (i) *Individual preferences*: Model disagreement as a proxy for diversity using pairwise preference learning [56], (ii) *Group distributions*: Extract diverse opinions from LLMs [32] and align them with survey-based population distributions, and (iii) *Societal values*: Dynamically combine and resolve value conflicts using crowdsourced resolution scenarios [17].

As AI becomes embedded in education, industry, and politics, systems that ignore plurality risk amplifying monolithic perspectives and deepening polarization; pluralistic AI instead models coexistence among differing values, beliefs, and cultural backgrounds. To build a research community around these questions, I co-founded the Pluralistic Alignment workshop at NeurIPS 2024 (second edition at ICML 2026), bringing together ML, HCI, and social science researchers. My future work will operationalize pluralistic alignment through quantitative diversity metrics, disagreement-aware training objectives, and cross-cultural evaluation datasets.

References

- [1] Galapagos: Build your own scaffold. <https://open-galapagos.com/>, 2026.
- [2] SciVeri: A benchmark for agents that do science. <https://sciveri-bench.com/>, 2026.
- [3] Dyah Adila and Dongyeop Kang. Understanding out-of-distribution: A perspective of data dynamics. In *ICBINB@NeurIPS*, 2021.
- [4] Daechul Ahn, Yura Choi, San Kim, Youngjae Yu, Dongyeop Kang, and Jonghyun Choi. ISR-DPO: Aligning large multimodal models for videos by iterative self-retrospective DPO. In *Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, 2025.
- [5] Daechul Ahn, Yura Choi, Youngjae Yu, Dongyeop Kang, and Jonghyun Choi. Tuning large multimodal models for videos using reinforcement learning from AI feedback. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.

- [6] Daechul Ahn, Daneul Kim, Gwangmo Song, Seung Hwan Kim, Honglak Lee, Dongyeop Kang, and Jonghyun Choi. Story visualization by online text augmentation with context memory. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, October 2023.
- [7] Shady Ali and Dongyeop Kang. Who you test matters: Difficulty alignment in LLM benchmark analysis. In *Workshop on Evaluating Evaluations (EvalEval) at ACL*, 2026.
- [8] Hyungjoo Chae, Yongho Song, Kai Tzu iunn Ong, Taeyoon Kwon, Minjin Kim, Youngjae Yu, Dongha Lee, Dongyeop Kang, and Jinyoung Yeo. Dialogue chain-of-thought distillation for commonsense-aware conversational agents. In *The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- [9] Seyyed Saeid Cheshmi, Hahnemann Ortiz, James Mooney, and Dongyeop Kang. Reasoning beyond literal: Cross-style multimodal reasoning for figurative language understanding. In *Findings of the Association for Computational Linguistics (EACL Findings)*, 2026.
- [10] Julia Christenson, Karin de Langis, Shirley Anugrah Hayati, and Dongyeop Kang. Effects of varying LLM access on essay writing behavior. In *Workshop on Innovative Use of NLP for Building Educational Applications (BEA) at ACL*, 2026. Oral.
- [11] Debarati Das. *LMs as Sociotechnical Actors: A Scaffolding Framework for Recovering Context Lost in Text-Only Settings*. PhD dissertation, University of Minnesota, Minneapolis, MN, May 2026. Committee: Jaideep Srivastava, Dongyeop Kang, Stevie Chancellor, Jisu Huh.
- [12] Debarati Das, Ishaan Gupta, Jaideep Srivastava, and Dongyeop Kang. Which modality should I use - text, motif, or image? : Understanding graphs with large language models. In Kevin Duh, Helena Gomez, and Steven Bethard, editors, *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 503–519, Mexico City, Mexico, June 2024. Association for Computational Linguistics.
- [13] Debarati Das, Karin De Langis, Anna Martin-Boyle, Jaehyung Kim, Minhwa Lee, Zae Myung Kim, Shirley Anugrah Hayati, Risako Owan, Bin Hu, Ritik Parkar, Ryan Koo, Jonginn Park, Aahan Tyagi, Libby Ferland, Sanjali Roy, Vincent Liu, and Dongyeop Kang. Under the surface: Tracking the artifactuality of LLM-generated data, 2024.
- [14] 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.
- [15] Debarati Das, David Ma, and Dongyeop Kang. Balancing effect of training dataset distribution of multiple styles for multi-style text transfer. In *Findings of the Annual Meeting of the Association for Computational Linguistics (ACL) Findings*, 2023.
- [16] Karin de Langis and Dongyeop Kang. A comparative study on textual saliency of styles from eye tracking, annotations, and language models. In Jing Jiang, David Reitter, and Shumin Deng, editors, *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 108–121, Singapore, December 2023. Association for Computational Linguistics.
- [17] Karin de Langis, Ryan Koo, and Dongyeop Kang. Dynamic multi-reward weighting for multi-style controllable generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [18] Karin de Langis, Püren Öncel, Ryan Peters, Andrew Elfenbein, Laura Kristen Allen, Andreas Schramm, and Dongyeop Kang. Mary, the cheeseburger-eating vegetarian: Do LLMs recognize incoherence in narratives? In *European Chapter of the Association for Computational Linguistics (EACL)*, 2026. Oral.
- [19] Karin de Langis, Jong Inn Park, Bin Hu, Khanh Chi Le, Andreas Schramm, Michael C. Mensink, Andrew Elfenbein, and Dongyeop Kang. Strong memory, weak control: An empirical study of executive functioning in LLMs. In *European Chapter of the Association for Computational Linguistics (EACL)*, 2026. Oral.
- [20] Karin de Langis, Jong Inn Park, Andreas Schramm, Bin Hu, Khanh Chi Le, and Dongyeop Kang. How LLMs comprehend temporal meaning in narratives: A case study in cognitive evaluation of LLMs. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 29174–29191, 2025.
- [21] Karin de Langis, William Walker, Khanh Chi Le, and Dongyeop Kang. Tracing how annotators think: Augmenting preference judgments with reading processes. In *International Conference on Language Resources and Evaluation (LREC)*, 2026.
- [22] Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*, pages 96–108, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [23] Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, and Dongyeop Kang. Understanding iterative revision from human-written text. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3573–3590, Dublin, Ireland, May 2022. Association for Computational Linguistics (ACL).
- [24] Steven Y. Feng, Varun Gangal, Dongyeop Kang, Teruko Mitamura, and Eduard Hovy. GenAug: Data augmentation for finetuning text generators. In *Deep Learning Inside Out (DeeLIO) Workshop at EMNLP 2020*, Online, November 2020.
- [25] Shuyu Gan, James Mooney, Pan Hao, Renxiang Wang, Mingyi Hong, Qianwen Wang, and Dongyeop Kang. Scaling unverifiable rewards: A case study on visual insights. In *Findings of the Association for Computational Linguistics (ACL Findings)*, 2026.
- [26] Shuyu Gan, Renxiang Wang, James Mooney, and Dongyeop Kang. A2P-Vis: an analyzer-to-presenter agentic pipeline for visual insights generation and reporting. In *Workshop on Visualization for Generative AI (VISxGenAI) at IEEE VIS*, 2025. Honorable Mention.
- [27] Gabriel Garlough-Shah, Jong Inn Park, Shirley Anugrah Hayati, Dongyeop Kang, and Jisu Huh. Consumer engagement with AI-powered search engines: Implications for the future of search advertising research and practice. *Journal of Advertising Research*, 2026.
- [28] Pan Hao, Dongyeop Kang, Nicholas Hinds, and Qianwen Wang. FlowForge: Guiding the creation of multi-agent workflows with design space visualization as a thinking scaffold. In *IEEE Transactions on Visualization and Computer Graphics (IEEE VIS)*, 2025.
- [29] Shirley A. Hayati, Dongyeop Kang, Qingxiaoyang Zhu, Weiyan Shi, and Zhou Yu. INSPIRED: Toward sociable recommendation dialog systems. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, November 2020.

- [30] Shirley Anugrah Hayati, Taehee Jung, Tristan Bodding-Long, Sudipta Kar, Abhinav Sethy, Joo-Kyung Kim, and Dongyeop Kang. Chain-of-instructions: Compositional instruction tuning on large language models. In *Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, 2025.
- [31] Shirley Anugrah Hayati, Dongyeop Kang, and Lyle Ungar. Does BERT learn as humans perceive? understanding linguistic styles through lexica. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [32] Shirley Anugrah Hayati, Minhwa Lee, Dheeraj Rajagopal, and Dongyeop Kang. How far can we extract diverse perspectives from large language models? In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [33] Shirley Anugrah Hayati, Kyumin Park, Dheeraj Rajagopal, Lyle Ungar, and Dongyeop Kang. StyLEx: Explaining styles with lexicon-based human perception. In *European Chapter of Association for Computational Linguistics (EACL)*, 2023.
- [34] Shirley Anugrah Hayati, Ruizi Wang, and Dongyeop Kang. From rubrics to recipe: Principle-centric benchmark for evaluating large language models. In *Workshop on Evaluating Evaluations (EvalEval) at ACL*, 2026.
- [35] 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.
- [36] Minoh Jeong, Min Namgung, Zae Myung Kim, Dongyeop Kang, Yao-Yi Chiang, and Alfred Hero. Anchors aweigh! sail for optimal unified multi-modal representations, 2024. under review.
- [37] Soyeong Jeong, Taehee Jung, Sung Ju Hwang, Joo-Kyung Kim, and Dongyeop Kang. When thoughts meet facts: Reusable reasoning for long-context LMs. In *Findings of the Association for Computational Linguistics (ACL Findings)*, 2026.
- [38] Hwiyeol Jo, Dongyeop Kang, Andrew Head, and Marti A. Hearst. Modeling mathematical notation semantics in academic papers. In *Findings on Empirical Methods in Natural Language Processing (EMNLP Findings)*, 2021.
- [39] Taehee Jung, Dongyeop Kang, Hua Cheng, Lucas Mentch, and Thomas Schaaf. Posterior calibrated training on sentence classification tasks. In *2020 Annual Conference of the Association for Computational Linguistics (ACL)*, 2020.
- [40] Taehee Jung, Joo-Kyung Kim, Sungjin Lee, and Dongyeop Kang. Cluster-guided label generation in extreme multi-label classification. In *European Chapter of Association for Computational Linguistics (EACL)*, 2023.
- [41] Seungyeon Jwa, Daechul Ahn, Reokyoung Kim, Dongyeop Kang, and Jonghyun Choi. Becoming experienced judges: Selective test-time learning for evaluators. In *European Chapter of the Association for Computational Linguistics (EACL)*, 2026. Oral.
- [42] Dongyeop Kang. *Linguistically Informed Language Generation: A Multifaceted Approach*. PhD dissertation, Carnegie Mellon University, Pittsburgh, PA, May 2020. Available at https://kilthub.cmu.edu/articles/thesis/Linguistically_Informed_Language_Generation_A_Multi-faceted_Approach/24990648.
- [43] Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. A dataset of peer reviews (PeerRead): Collection, insights and NLP applications. In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*, New Orleans, USA, June 2018.
- [44] Dongyeop Kang, Anusha Balakrishnan, Pararth Shah, Paul Crook, Y-Lan Boureau, and Jason Weston. Recommendation as a communication game: Self-supervised bot-play for goal-oriented dialogue. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, November 2019.
- [45] Dongyeop Kang, Varun Gangal, and Eduard Hovy. (male, bachelor) and (female, Ph.D) have different connotations: Parallely annotated stylistic language dataset with multiple personas. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, November 2019.
- [46] Dongyeop Kang, Varun Gangal, Ang Lu, Zheng Chen, and Eduard Hovy. Detecting and explaining causes from text for a time series event. In *Conference on Empirical Methods on Natural Language Processing (EMNLP)*, 2017.
- [47] Dongyeop Kang, Hiroaki Hayashi, Alan W Black, and Eduard Hovy. Linguistic versus latent relations for modeling coherent flow in paragraphs. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, November 2019.
- [48] Dongyeop Kang, Andrew Head, Risham Sidhu, Kyle Lo, Daniel Weld, and Marti A. Hearst. Document-level definition detection in scholarly documents: Existing models, error analyses, and future directions. In *Proceedings of the First Workshop on Scholarly Document Processing*, pages 196–206, Online, November 2020. Association for Computational Linguistics.
- [49] Dongyeop Kang and Eduard Hovy. Plan ahead: Self-supervised text planning for paragraph completion task. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, November 2020.
- [50] Dongyeop Kang and Eduard Hovy. Style is NOT a single variable: Case studies for cross-stylistic language understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2376–2387, Online, August 2021. Association for Computational Linguistics (ACL).
- [51] Dongyeop Kang*, Taehee Jung*, Lucas Mentch, and Eduard Hovy. Earlier isn’t always better: Sub-aspect analysis on corpus and system biases in summarization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, November 2019.
- [52] Dongyeop Kang, Tushar Khot, Ashish Sabharwal, and Peter Clark. Bridging knowledge gaps in neural entailment via symbolic models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, November 2018.
- [53] Dongyeop Kang, Tushar Khot, Ashish Sabharwal, and Eduard Hovy. AdvEntuRe: Adversarial training for textual entailment with knowledge-guided examples. In *The 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, Melbourne, Australia, July 2018.
- [54] Jihyung Kil, Farideh Tavazoei, Dongyeop Kang, and Joo-Kyung Kim. II-MMR: Identifying and improving multi-modal multi-hop reasoning in visual question answering. In *Annual Meeting of the Association for Computational Linguistics (ACL) Findings*, 2024.

- [55] Haechan Kim, Junho Myung, Seoyoung Kim, Sungpah Lee, Dongyeop Kang, and Juho Kim. LearnerVoice: A dataset of non-native English learners' spontaneous speech. In *Conference of the International Speech Communication Association (Interspeech)*, 2024.
- [56] Jaehyung Kim, Shin Jinwoo, and Dongyeop Kang. Prefer to classify: Improving text classifiers via auxiliary preference learning. In *International Conference on Machine Learning (ICML)*, 2023.
- [57] Jaehyung Kim, Dongyeop Kang, Sungsoo Ahn, and Jinwoo Shin. What makes better augmentation strategies? augment difficult but not too different. In *International Conference on Learning Representations (ICLR)*, 2022.
- [58] Jaehyung Kim, Yekyung Kim, Karin de Langis, Jinwoo Shin, and Dongyeop Kang. infoVerse: A universal framework for dataset characterization with multidimensional meta-information. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [59] Katie Kim, Bugil Chang, Debarati Das, Smitha Sudheendra, Jisu Huh, Dongyeop Kang, and Jaideep Srivastava. Rebuilding social connection and enhancing advertising effects through the nostalgic appeal during the pandemic. In *Association for Education in Journalism and Mass Communication (AEJMC)*, 2023.
- [60] Zae Myung Kim, Wanyu Du, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. Improving iterative text revision by learning where to edit from other revision tasks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9986–9999, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics.
- [61] Zae Myung Kim, Chanoo Kim, Vipul Raheja, and Dongyeop Kang. Toward evaluative thinking: Meta policy optimization with evolving reward models, 2025. under review.
- [62] Zae Myung Kim, Dongseok Lee, Jaehyung Kim, Vipul Raheja, and Dongyeop Kang. The amazing agent race: Strong tool users, weak navigators, 2026.
- [63] Zae Myung Kim, Kwang Hee Lee, Preston Zhu, Vipul Raheja, and Dongyeop Kang. Threads of subtlety: Detecting machine-generated texts through discourse motifs. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [64] Zae Myung Kim, Young-Jun Lee, Seungyeon Jwa, and Dongyeop Kang. Metaⁿ: Recursive self-improvement through emergent depth. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2026.
- [65] Zae Myung Kim, Vassilina Nikoulina, Dongyeop Kang, Didier Schwab, and Laurent Besacier. Visualizing Cross-Lingual discourse relations in multilingual TED corpora. In *Proceedings of the 2nd Workshop on Computational Approaches to Discourse*, pages 165–170, Punta Cana, Dominican Republic and Online, November 2021. Association for Computational Linguistics.
- [66] Zae Myung Kim, Anand Ramachandran, Farideh Tavazoei, Joo-Kyung Kim, Oleg Rokhlenko, and Dongyeop Kang. Align to structure: Aligning large language models with structural information. In *Fortieth AAAI Conference on Artificial Intelligence (AAAI)*, 2026.
- [67] Ryan Koo, Yekyung Kim, Dongyeop Kang, and Jaehyung Kim. Meta-crafting: Improved detection of out-of-distributed texts via crafting metadata space (student abstract). In *Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI) Student Abstract*, 2024.
- [68] Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. Benchmarking cognitive biases in large language models as evaluators. In *Annual Meeting of the Association for Computational Linguistics (ACL) Findings*, 2024.
- [69] Ryan Koo, Anna Martin, Linghe Wang, and Dongyeop Kang. Decoding the end-to-end writing trajectory in scholarly manuscripts. In *Proceedings of the Second Workshop on Intelligent and Interactive Writing Assistants (In2Writing) at CHI 2023*, 2023. https://minnesotanlp.github.io/REWARD_demo/.
- [70] Ryan Koo, Ian Yang, Vipul Raheja, Mingyi Hong, Kwang-Sung Jun, and Dongyeop Kang. Learning explainable dense reward shapes via Bayesian optimization, 2025.
- [71] Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task, 2025.
- [72] Khanh Chi Le, Linghe Wang, Minhwa Lee, Ross Volkov, Luan Tuyen Chau, and Dongyeop Kang. ScholaWrite: A dataset of end-to-end scholarly writing. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026.
- [73] Minhwa Lee, Zae Myung Kim, Vivek Khetan, Alex Kass, and Dongyeop Kang. CoTaxon: Enhancing domain expertise through human-AI collaborative taxonomy construction, 2025. under review.
- [74] Minhwa Lee, Zae Myung Kim, Vivek A. Khetan, and Dongyeop Kang. Human-AI collaborative taxonomy construction: A case study in profession-specific writing assistants. In *CHI 2024 In2Writing Workshop*, 2024.
- [75] Young-Jun Lee, Seungone Kim, Minki Kang, Alistair Cheong Liang Chuen, Zerui Chen, Seungho Han, Taehee Jung, and Dongyeop Kang. Evolution fine-tuning: Learning to discover across 371 optimization tasks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2026.
- [76] Chenliang Li, Siliang Zeng, Zeyi Liao, Jiaxiang Li, Dongyeop Kang, Alfredo Garcia, and Mingyi Hong. Joint reward and policy learning with demonstrations and human feedback improves alignment. In *International Conference on Learning Representations (ICLR)*, 2025. Spotlight.
- [77] Chu-Cheng Lin, Dongyeop Kang, Michael Gamon, Madian Khabsa, Ahmed Hassan Awadallah, and Patrick Pantel. Actionable email intent modeling with reparametrized RNN. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2018.
- [78] Kyle Lo, Joseph Chee Chang, Andrew Head, Jonathan Bragg, Amy X. Zhang, Cassidy Trier, Chloe Anastasiades, Tal August, Russell Authur, Danielle Bragg, Erin Bransom, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Yen-Sung Chen, Evie Yu-Yen Cheng, Yvonne Chou, Doug Downey, Rob Evans, Raymond Fok, Fangzhou Hu, Regan Huff, Dongyeop Kang, Tae Soo Kim, Rodney Kinney, Aniket Kittur, Hyeonsu Kang, Egor Klevak, Bailey Kuehl, Michael Langan, Matt Latzke, Jaron Lochner, Kelsey MacMillan, Eric Marsh, Tyler Murray, Aakanksha Naik, Ngoc-Uyen Nguyen, Srishti Palani, Soya Park, Caroline Paulic, Napol Rachatasumrit, Smita Rao, Paul Sayre, Zejiang Shen, Pao Siangliulue, Luca Soldaini, Huy Tran, Madeleine van Zuylen, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamarron, Marti A. Hearst, and Daniel S. Weld. The semantic reader project: Augmenting scholarly documents through AI-powered interactive reading interfaces. In *Communications of ACM*, 2023.

- [79] London Lowmanstone, Ruyuan Wan, Risako Owan, Jaehyung Kim, and Dongyeop Kang. Annotation imputation to individualize predictions: Initial studies on distribution dynamics and model predictions. In *Workshop on Perspectivist Approaches to NLP (NLPerspectives) @ECAI 2023*, 2023.
- [80] Anna Martin-Boyle, Aahan Tyagi, Marti A. Hearst, and Dongyeop Kang. Shallow synthesis of knowledge in GPT-generated texts: A case study in automatic related work composition, 2024.
- [81] Minnesota NLP. SciFlow: A longitudinal data collection for human–AI scientific workflows. <https://minnesotanlp.github.io/sciflow/>, 2026.
- [82] James Mooney, Zae Myung Kim, Young-Jun Lee, and Dongyeop Kang. Structure liberates: How constrained sensemaking produces more novel research output, 2026.
- [83] James Mooney, Josef Woldense, Zheng Robert Jia, Shirley Anugrah Hayati, My Ha Nguyen, Vipul Raheja, and Dongyeop Kang. Are LLM agents behaviorally coherent? latent profiles for social simulation, 2025.
- [84] Smitha Muthya Sudheendra, Maral Abdollahi, Dongyeop Kang, Jisu Huh, and Jaideep Srivastava. SkOTaPA: A dataset for skepticism detection in online text after persuasion attempt. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue, editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14871–14876, Torino, Italia, May 2024. ELRA and ICCL.
- [85] Jinwoo Nam, Daechul Ahn, Dongyeop Kang, Seong Jong Ha, and Jonghyun Choi. Zero-shot natural language video localization. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2021.
- [86] Rose Neis, Karin de Langis, Zae Myung Kim, and Dongyeop Kang. An analysis of reader engagement in literary fiction through eye tracking and linguistic features. In *Workshop on Narrative Understanding (WNU) @ACL 2023*, 2023.
- [87] C. Thi Nguyen. The seductions of clarity. *Royal Institute of Philosophy Supplements*, 89:227–255, 2021.
- [88] Jong Inn Park, Maanas Taneja, Qianwen Wang, and Dongyeop Kang. Stealing creator’s workflow: A creator-inspired agentic framework with iterative feedback loop for improved scientific short-form generation, 2025.
- [89] Ritik Sachin Parkar, Jaehyung Kim, Jong Inn Park, and Dongyeop Kang. SelectLLM: Can LLMs select important instructions to annotate?, 2024.
- [90] Andy J. Phu, Karin de Langis, James Mooney, Khanh Chi Le, and Dongyeop Kang. SERUM: State extraction and refinement for user modeling. In *Conference on Language Modeling (COLM)*, 2026.
- [91] Vipul Raheja, Dhruv Kumar, Ryan Koo, and Dongyeop Kang. CoEdit: Text editing by task-specific instruction tuning. In *The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP Findings)*, 2023.
- [92] Andreas Schramm, Karin de Langis, Jong Inn Park, Anh Thu Tong, Michael Mensink, Bin Hu, Khanh Chi Le, and Dongyeop Kang. How human is AI? a comparison of temporal meanings in 8 LLMs and 3 populations. In *American Association for Applied Linguistics (AAAL)*, 2025.
- [93] Daniel Schwarcz, Debarati Das, Dongyeop Kang, and Brett McDonnell. Thinking like a lawyer in the age of generative AI: Cognitive limits on AI adoption among lawyers. *Journal of Institutional and Theoretical Economics*, pages 6–26, 2026.
- [94] Zhecheng Sheng, Tianhao Zhang, Chen Jiang, and Dongyeop Kang. BBScore: A Brownian bridge based metric for assessing text coherence. In *Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI)*, 2024.
- [95] Ioannis Tsaknakis, Bingqing Song, Shuyu Gan, Jiangweizhi Peng, Dongyeop Kang, Alfredo Garcia, Gaowen Liu, Charles Fleming, and Mingyi Hong. Do LLMs recognize your latent preferences? a benchmark for latent information discovery in personalized interaction. In *Conference on Language Modeling (COLM)*, 2026.
- [96] Ruyuan Wan, Jaehyung Kim, and Dongyeop Kang. Everyone’s voice matters: Quantifying annotation disagreement using demographic information. In *Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
- [97] Quan Wei, Chung-Yiu Yau, Hoi To Wai, Yang Zhao, Dongyeop Kang, Youngsuk Park, and Mingyi Hong. RoSTE: An efficient quantization-aware supervised fine-tuning approach for large language models. In *International Conference on Machine Learning (ICML)*, 2025.
- [98] Carl Winge, Adam Imdieke, Bahaa Aldeeb, Dongyeop Kang, and Karthik Desingh. Talk through it: End user directed manipulation learning. In *IEEE Robotics and Automation Letters (RA-L)*, 2024.
- [99] Ruixin Yang, Dheeraj Rajagopal, Shirley Anugrah Hayati, Bin Hu, and Dongyeop Kang. Confidence calibration and rationalization for LLMs via multi-agent deliberation. In *ICLR 2024 Workshop on Reliable and Responsible Foundation Models*, 2024.
- [100] Hongseon Yeom, Jaeyoul Shin, Soojin Min, Jeongmin Yoon, Seunghak Yu, and Dongyeop Kang. Breaking determinism: Stochastic modeling for reliable off-policy evaluation in ad auctions. In *ACM International Conference on Web Search and Data Mining (WSDM)*, 2026. Oral.
- [101] Haneul Yoo, Rifki Afina Putri, Changyoon Lee, Youngin Lee, So-Yeon Ahn, Dongyeop Kang, and Alice Oh. Rethinking annotation: Can language learners contribute? In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [102] Skylar Zhai, Jingcheng Liang, and Dongyeop Kang. Abstain-R1: Calibrated abstention and post-refusal clarification via verifiable RL. In *Findings of the Association for Computational Linguistics (ACL Findings)*, 2026.
- [103] Tianhao Zhang, Zhecheng Sheng, Zhexiao Lin, Chen Jiang, and Dongyeop Kang. BBScoreV2: Learning time-evolution and latent alignment from stochastic representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
- [104] Hao Zou, Zae Myung Kim, and Dongyeop Kang. A survey of diffusion models in natural language processing, 2023.