

Learn more about Adobe Research's work on Agentic AI by viewing some of our research publications in this area.

Updated on August 7, 2025

Agent optimization & frameworks

[PromptRefine: Enhancing Few-Shot Performance on Low-Resource Indic Languages with Example Selection from Related Example Banks \(NAACL'25\)](#)

[LoRA-Contextualizing Adaptation of Large Multimodal Models for Multi-page Document Understanding \(ICLR'25\)](#)

[MoDS: Moderating a Mixture of Document Speakers to Summarize Debatable Queries in Document Collections \(NAACL'25\)](#)

[VipAct: Visual-Perception Enhancement via Specialized VLM Agent Collaboration and Tool-use](#)

[RCStat: A Statistical Framework for using Relative Contextualization in Transformers](#)

Planning and reasoning

[ToolChain*: Efficient Action Space Navigation in Large Language Models with A* Search \(ICLR'24\)](#)

[RaDA: Retrieval-augmented Web Agent Planning with LLMs \(ACL'24\)](#)

[Calibrating Reasoning in Language Models with Internal Consistency \(NeurIPS'24\)](#)

[DeCoT: Debiasing Chain-of-Thought for Knowledge-Intensive Tasks in Large Language Models via Causal Intervention \(ACL'24\)](#)

Attribution, hallucination prevention, and evaluation

[Hallucination Diversity-Aware Active Learning for Text Summarization \(NAACL'24\)](#)

[Token-level adversarial prompt detection based on perplexity measures and contextual information \(ICLR'25 Workshop\)](#)

[A Flash in the Pan: Better Prompting Strategies to Deploy Out-of-the-Box LLMs as Conversational Recommendation Systems \(Coling'25\)](#)

[Self-Debiasing Large Language Models: Zero-Shot Recognition and Reduction of Stereotypes \(NAACL'25\)](#)

[Personalizing Data Delivery: Investigating User Characteristics and Enhancing LLM Predictions \(The Web Conf'25\)](#)

[OCEAN: Offline Chain-of-thought Evaluation and Alignment in Large Language Models \(ICLR'25\)](#)

[Annotation Efficiency: Identifying Hard Samples via Blocked Sparse Linear Bandits](#)

[A Multi-LLM Debiasing Framework](#)

Agentic applications in marketing, creativity, and documents

[Continuous Treatment Effects with Surrogate Outcomes \(ICML'24\)](#)

[Doc-React: Multi-page Heterogeneous Document Question-answering \(ACL'25\)](#)

[Diversify-verify-adapt: Efficient and Robust Retrieval-Augmented Ambiguous Question Answering \(NAACL'25\)](#)

[Exploring Rewriting Approaches for Different Conversational Tasks](#)

[VisDoM: Multi-Document QA with Visually Rich Elements Using Multimodal Retrieval-Augmented Generation \(NAACL'25\)](#)

[Moderating a Mixture of Document Speakers to Summarize Debatable Queries in Document Collections \(NAACL'25\)](#)

[Multi-LLM Text Summarization](#)

[PosterMate: Audience-driven Collaborative Persona Agents for Poster Design \(UIST'25\)](#)

Code generation – NL2SQL, NL2Script, and NL2Action (creative)

[FINDR: A Fast Influential Data Selector for NL2Code Pretraining](#)

[CodeLutra: Boosting LLM Code Generation via Preference-Guided Refinement \(TMLR'25\)](#)

[Augment before You Try: Knowledge-Enhanced Table Question Answering via Table Expansion](#)

[Mixture of Structural-and-Textual Retrieval over Text-rich Graph Knowledge Bases \(ACL'25\)](#)

Agentic experiences

[Few-Shot Dialogue Summarization via Skeleton-Assisted Prompt Transfer in Prompt Tuning](#) (EACL'24)

[Knowledge-Aware Query Expansion with Large Language Models for Textual and Relational Retrieval](#) (NAACL'25)

[A Case Study of Human-Authored versus Automatic Dashboard Summaries](#) (CHI'25)

[Cache-Craft: Managing Chunk-Caches for Efficient Retrieval-Augmented Generation](#) (SIGMOD'25)

[Since U Been Gone: Augmenting Context-Aware Transcriptions for Re-engaging in Immersive VR Meetings](#) (CHI'25)

[Toward Living Narrative Reviews: An Empirical Study of the Processes and Challenges in Updating Survey Articles in Computing Research](#) (CHI'25)