



Dr. Subham Chatterjee

Assistant Professor

Department of Computer Science

Missouri University of Science and Technology

**Beyond Relevance: What Information Matters to AI Systems
That Search, Reason, and Remember?**

Abstract. Modern AI systems increasingly do more than generate text. They retrieve external information, reason over evidence, search iteratively, and maintain memory across interactions. These capabilities raise a fundamental question: what does it mean for a piece of information to be useful to an AI system? Traditional information retrieval has largely treated usefulness through the lens of relevance. But once retrieval becomes part of a larger AI system, relevance may no longer tell the whole story. In this talk, I will discuss our recent work examining information utility in retrieval-augmented generation and agentic search. In RAG systems, we find that retrieval signals designed to better capture answer-supporting evidence do not always translate into improved downstream behavior, highlighting a gap between seemingly better retrieval objectives and the outcomes we ultimately care about. In multi-step agentic search, this gap becomes even more pronounced: information that appears unhelpful to a static reader can be causally important because it changes what the agent searches for or does next. These findings suggest that information utility may need to be understood not as a fixed property of an individual document, but as something that depends on context, other available evidence, and the evolving trajectory of the system. I will use these results to motivate several directions we are currently exploring, including set-dependent notions of evidence utility, minimum sufficient evidence, and adaptive long-term memory for AI agents. More broadly, the talk will argue that as AI systems evolve from models that answer questions into systems that search, act, and remember, we need richer ways to measure which information actually matters.

Biographical Sketch. Shubham Chatterjee is an Assistant Professor of Computer Science at Missouri University of Science and Technology. His research lies at the intersection of information retrieval and artificial intelligence, with a particular focus on retrieval-augmented generation, neural information retrieval, agentic search, conversational systems, personalization, and memory for AI systems. His recent work examines how retrieved information affects downstream generation and agent behavior, including how to distinguish information that merely appears relevant from information that causally influences an AI system's decisions and search trajectory.