



DISSERTATION DEFENSE

Shuyang Cao



Towards Controllability, Efficiency, and Trustworthiness of Text Generation Systems

Monday, May 12, 2025

12:00pm – 2:00pm

3725 Beyster / Hybrid - [Zoom](#)

ABSTRACT: As automated text generation becomes more integrated into real-world applications, it is crucial to understand and address its limitations. This thesis explores three important aspects of text generation systems—controllability, trustworthiness, and efficiency.

The first part of this thesis focuses on *controllability*, where generated text must align with specific user requirements. We present a type-controlled question generation framework that uses exemplar templates and optionally generated fine-grained templates to guide generation more precisely. Second, we address *trustworthiness*. We propose a contrastive learning framework that trains models to distinguish correct outputs from automatically collected incorrect ones, leading to more factual model-generated summaries. Beyond generating trustworthy content, we explore verifiable generation with fine-grained citations, where model outputs are augmented with span-level references to source documents, enabling users to verify claims efficiently. Finally, we investigate *efficiency*, particularly in processing long documents. We propose a divide-and-conquer training framework that reduces memory requirements by summarizing document chunks individually, while maintaining coherence through memory mechanisms and shared salient content. To systematically assess long-context models on inputs of any length, we also develop an adaptive evaluation benchmark consisting of synthetic tasks targeting varying levels of model capabilities, all sharing the same context input to enable controlled analysis. These contributions represent key steps towards controllability, trustworthiness, and efficiency in the design of modern text generation systems.

CHAIR: Prof. Lu Wang