

GSAS Universals Workshops in Linguistics

Alexander Sasha Rush (SEAS, Harvard University)

February 15 at 3:30 pm

Sever Hall Room 306

Title: Controllable Text Generation

Abstract: Progress in deep learning has led to optimism for automatic text generation. Yet state-of-the-art systems still predict wildly inaccurate output on a non-trivial percentage of examples. Lack of user control to correct these issues makes it difficult to deploy these models in real applications. In this talk, I will argue that discrete latent-variable models provide a natural declarative framework for more controllable text models. I will present two recent works exploring this theme: (1) a method for learning neural template models that can be adjusted directly by users (Wiseman et al, 2018); (2) an approach to learning soft alignments as a latent variable (Deng et al, 2018). I will end by talking about the significant remaining challenges with this approach.