Thesis Defense: Hayley Ross

Date and time: Wednesday, April 16, 2:00 – 3:30 pm

Location: 103 Boylston Hall

Title: Artificial intelligence and fake reefs: what privative inferences and LLMs tell us about adjective-noun composition

Abstract: The fact that people understand completely novel phrases is often taken as an argument that linguistic meaning is composed from the meaning of its parts. Thus, a central concern for the study of meaning is how that meaning is composed, especially for open-class content words like adjectives and nouns. This dissertation studies meaning composition and its interaction with context through the lens of adjective-noun modification and the privative inferences that sometimes result (e.g., a fake gun is (usually) not a gun, and a stone lion is not a (living) lion). This dissertation shows that privativity is not limited to a particular class of adjectives, which leads to a new, non-intersective semantics for adjective-noun modification which handles potential contradictions as part of composition. Further, we find that humans and modern large language models (LLMs) can generalize to the inferences of adjective-noun combinations that they have not seen before. Working with LLMs foregrounds the possibility that these inferences could be drawn by other means than meaning composition, such as memorization or analogy. The resulting adaptation in experiment design as well as reflection on our standards of evidence feeds into the broader, currently emerging discussion about how to study compositionality in humans and language models alike.