

Predicting Legal Citations

Leader: Jack Cushman

For this project we're asking students to predict legal citations: given a legal case truncated immediately before a citation to another case, predict the case that was cited. A baseline would be to simply predict a random case that was decided in the same jurisdiction within the previous 10 years, but we want to see how much better students can do through analysis of the text. Students will construct their own training and test sets based on collections of case law from the Caselaw Access Project. This problem statement has real-world applications such as natural language search, recommender systems, and drafting tools.