

# **Pardon the Interruption: Automatic Analysis of Gender and Competitive Turn-Taking in United States Supreme Court Hearings**

**Haley Lepp**

University of Washington

hlepp@uw.edu

## **Abstract**

The United States Supreme Court plays a key role in defining the legal basis for gender discrimination throughout the country, yet there are few checks on gender bias within the court itself. In conversational turn-taking, interruptions have been documented as a marker of bias between speakers of different genders. The goal of this study is to automatically differentiate between respectful and disrespectful turn-taking during official hearings, which could help in detecting bias and finding remediation techniques for discourse in the courtroom. In this paper, I present a corpus of turns annotated by legal professionals, and describe the design of a semi-supervised classifier that will use acoustic and lexical features to analyze turn-taking at scale. On completion of annotations, this classifier will be trained to extract the likelihood that turns are respectful or disrespectful for use in studies of speech trends.

## **1 Introduction**

The United States Supreme Court is a influential and controversial government institution that shapes American civil rights. As spoken discourse is critical to representation and influence in the courts, determining linguistic features of inequality or bias and automatically detecting them could be valuable to improving inclusion. Jacobi and Schweers (2017) find that interruptions correlating with gender within Supreme Court hearings have occurred consistently over time, and are not necessarily due to political polarization or personalities of certain justices. Tannen (1994) warns that speech overlap does not necessarily constitute impropriety, suggesting that speech overlaps may be both competitive and cooperative in different contexts. Determining what distinguishes a disrespectful turn from natural or encouraging turns is therefore necessary for the evaluation of turn-taking in hearings. Analysis specific to this domain is critical, because turn-taking that may be considered competitive or disrespectful in another domain may be ordinary in the context of a hearing. This paper describes preliminary results from small-scale studies of turn-taking in hearings, discusses the in-progress annotation of hearing recordings, and proposes the design of a domain-specific classifier that will use acoustic and lexical features to automatically predict the appropriateness of a turn.

## **2 Data Collection and Annotation of Competitive and Cooperative Overlaps**

All Supreme Court hearings since 2006 are publicly available. The hour-long hearings have all been professionally transcribed. Using a corpus of manually annotated turns from *Apple v. Pepper* (2018) and *Mont v. United States* (2018), the *inaSpeechSegmenter* (Doukhan et al., 2018) can automatically predict turns and speaker gender with an F-1 score of .744. The author, who has no legal training, manually labeled 99 extracted turns from the two different hearings as competitive, in which the second speaker taking the turn from the first speaker is interpreted as disrespectful; or cooperative, in which the speaker taking the turn is respectful. In this annotated set, over 90% of turns that included speech overlap fell

into a competitive category. The frequency of speech overlap appears to correlate with speaker role in the hearing (justice or attorney). The distribution of competitive and cooperative overlaps was entirely correlated with gender: 100% of cooperative overlaps occur between male speakers.

To expand and improve the quality of annotations, 10 American legal professionals will listen to the automatically-detected speech overlaps and non-overlapping turn changes, plus one second of speech on either side. In perception studies, Wichmann and Caspers (2001) similarly categorized speech overlaps, and asked listeners to predict based on whether the first speaker in an exchange had expected a turn change or planned to continue speaking. The annotators will rank the speech segments on a five point scale as to whether or not the turn change was expected and welcome. Each overlap will be annotated by four participants and checked for inter-annotator agreement.

### 3 Automatic Classification of Overlap Type

**Related Work** While automatically assessing turn-taking using text-based cues, such as lack of syntactic completion, has been highly studied, there are few studies that automate analysis of turn-taking using phonetic cues. Yang (2001) and Truong (2013) found a correlation between pitch contours and competitive interruptions. Amplitude and duration of speech overlaps have also been used to distinguish between when an interruption is competitive or cooperative (Kurtic, 2013). Due to the nature of the Supreme Court hearings recordings, amplitude may not be accurately reflected in the corpus and therefore will not be measured in this study.

**Preliminary Findings** Initial measurements of the small, layperson-annotated corpus of turns indicate that duration of overlap correlates highly with turn type. Duration was consistently longer for competitive than cooperative overlaps. Competitive overlaps were on average 0.480 seconds long. Cooperative overlaps were on average 0.070 seconds long. The use of a larger and professionally annotated corpus will help to verify this trend. Measurements of pitch also confirm the results of previous studies: F0 measurements of the second speaker in the turn follow a clear high-to-low contour in competitive overlaps, while pitch slope in cooperative overlaps was largely static. Determining when to start and end measurements of pitch automatically will require tuning, which will take place on a training set once the professional annotations of the corpus are complete.

### 4 Automatic Feature Detection and Classification

With limited annotated turns but a relatively large corpus of hearings, this classification problem can be solved using a multi-view semi-supervised learning approach. Acoustic and lexical feature sets will be integrated by using a co-training model, with each set acting as a distinct view of the data (Blum and Mitchell, 1998). To extract features at scale, Praat and Python scripts will automatically measure duration and pitch contour of overlaps from a subset of the corpus. Dependency parsers will extract lexical and syntactic text-based cues from court transcripts. Using a fraction of the annotated corpus, the model will assign probabilities to features associated with annotated overlaps. Un-annotated predictions of positive (cooperative) and negative (competitive) overlaps with high confidence will be added to the training set. The system will continue to iterate until the probability of predictions on labeled data stops increasing. I will analyze final accuracy levels using a held-out portion of the annotated corpus.

### 5 Conclusion

This is an in-progress study of the automatic classification of turns taken in U.S. Supreme Court hearings. Preliminary results indicate that turn type correlates highly with pitch contour and overlap duration. The annotation of a corpus by legal professionals over the coming months will allow for the training of a semi-supervised learning system to classify turn types. This annotated corpus and the predictions of the model can be used to demonstrate the distribution of inappropriate turns across variables such as gender, role in court, and even individual personality. In an era when bias within the U.S. courts is often discussed but rarely measured, it is critical that future studies continue to develop methods for detecting and mitigating bias.

## 6 Acknowledgments

I wish to thank the faculty at the University of Washington Department of Linguistics, especially Gina-Anne Levow and Richard Wright, for their guidance and feedback.

## References

- Apple v. Pepper. 2018. *Oyez*. [www.oyez.org/cases/2018/17-204](http://www.oyez.org/cases/2018/17-204). Accessed 15 Jun. 2019.
- Avrim Blum and Tom Mitchell. 1998. *Combining Labeled and Unlabeled Data with Co-Training*. Proceedings of the 11th Annual Conference on Computational Learning Theory, 92-100.
- David Doukhan, Jean Carrière, Felicien Vallet, Anthony Larcher, and Sylvain Meignier. 2018. *An Open-Source Speaker Gender Detection Framework for Monitoring Gender Equality*. Acoustics Speech and Signal Processing (ICASSP).
- Tonja Jacobi and Dylan Schweers. 2017. *Justice, Interrupted: The Effect of Gender, Ideology and Seniority at Supreme Court Oral Arguments*. Northwestern Law & Econ Research Paper No. 17-03.
- Emina Kurtic, Guy J. Brown, and Bill Wells. 2013. *Resources for turn competition in overlapping talk*. Speech Communication, Volume 55 Issue 5, 721-743.
- Mont v. United States. 2018. *Oyez*. [www.oyez.org/cases/2018/17-8995](http://www.oyez.org/cases/2018/17-8995). Accessed 15 Jun. 2019.
- Deborah Tannen. 1994. *Gender and discourse*. New York: Oxford University Press.
- Khiet Phuong Truong. 2013. *Classification of cooperative and competitive overlaps in speech using cues from the context, overlapper, and overlappee*. In INTERSPEECH-2013, 1404-1408.
- Anne Wichmann and Johanneke Caspers. 2001. *Melodic cues to turn-taking in English: evidence from perception*. Proceedings of the Second SIGdial Workshop on Discourse and Dialogue.
- Li-Chiung Yang. 2001. *Visualizing Spoken Discourse: Prosodic Form and Discourse Functions of Interruptions*. Proceedings of the Second SIGdial Workshop on Discourse and Dialogue.