

Student Research Symposium 2004

Satanjeev (Bano) Banerjee

Automatically Detecting the Structure of Human Meetings

We are interested in automatically extracting the structure of meetings between humans. Such structure includes the state of a meeting (presentation, discussion, etc), the roles of each meeting participant (presenter, discussion participator, observer, etc), the onset/offset boundaries of agenda items, and the onset/offset boundaries of regions of decisions (such as action items). In this talk we will talk about our current research into detecting these various aspects of human meetings.

In particular, we will present a simple taxonomy of meeting states and participant roles. We trained a decision tree classifier that learns to detect these states and roles from simple speech-based features such as the number of speakers and the lengths of utterances and speech-overlaps. This classifier detects meeting states 18% absolute more accurately than a random classifier, and detects participant roles 10% absolute more accurately than a majority classifier. We will then report on the effect of adding more advanced features such as the words in the utterances as output by an automatic speech recognizer, as well as features drawn from other modalities such as the body positions and face directions of the various participants relative to each other as output by a camera-image processor.

Finally we will present initial research on agenda item and decision region boundary detection. Unlike meeting state and participant role detection, the problem of detecting agenda items and decision regions does not easily lend itself to a typical machine learning approach, since there are no clear pre-defined classes. However, preliminary observations of recorded meeting data suggest that different agenda items usually differ highly in both the pattern of words used in discussing them, as well as in the identities of the participants involved in the discussions thereof. We will report on our ongoing research where we draw upon ideas from the realm of topic tracking and leverage the above characteristics to perform agenda item/decision region detection.