

Carnegie Mellon University



MultiComp Lab



First Grand-Challenge and Workshop on Human Multimodal Language (Challenge-HML)

Organizers:



Amir Zadeh



Louis-Philippe Morency



Paul Pu Liang



Soujanya Poria



Erik Cambria



Stefan Scherer



Zhun Liu

Continuous Theories of (Multimodal) Language




Throughout evolution language and nonverbal behaviors developed together.



Cries and Imitations

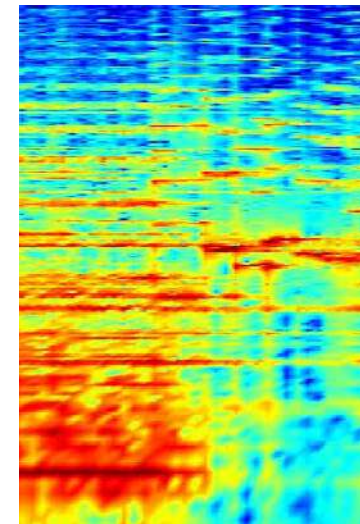


Modern Language



Multimodal Language Modalities

“I really like this vacuum cleaner, it was one of the few that has a large dust bag ...”



Language (words)

Vision (gestures)

Acoustic (voice)

Challenge-HML Structure

- Our goal in the first Challenge-HML is to build models of sentiment and emotions for human multimodal language through the proxy of in-the-wild speech videos.



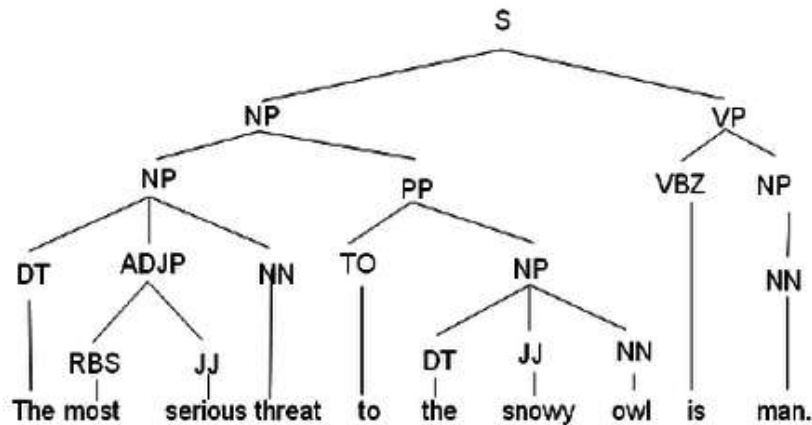
Challenge-HML Metrics

- Sentiment: Binary and Multiclass Sentiment Analysis
- Emotion Recognition: Binary and Multiclass Analysis of
 - Happiness
 - Sadness
 - Anger
 - Surprise
 - Disgust
 - Fear


Difficulties in Modeling Multimodal Language

- Modeling


- Intra-modal Dynamics: difficulties in modeling each modality.



Language (words)



Vision (gestures)



Acoustic (voice)

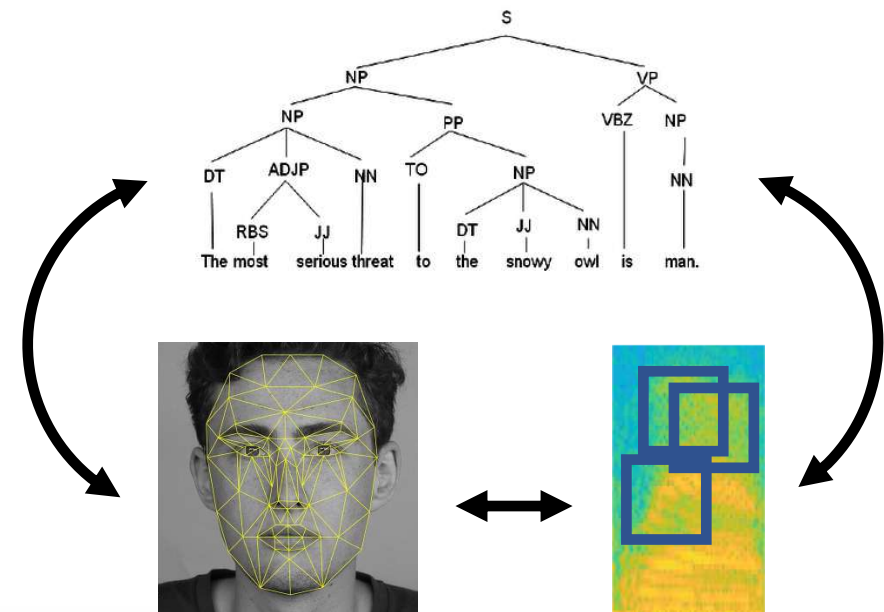
Difficulties in Modeling Multimodal Language

- Modeling

- Intra-modal Dynamics


- Difficulties in modeling each modality

- Inter-modal Dynamics (fusion): difficulties in modeling spatio-temporal relations between modalities



Difficulties in Modeling Multimodal Language

- Modeling
 - Intra-modal Dynamics
 - Difficulties in modeling each modality
 - Inter-modal Dynamics (fusion)
- Idiosyncratic signal
 - Speakers talk and behave differently



CMU-MOSEI Dataset

- CMU-Multimodal Sentiment and Emotion Intensity Analysis Dataset
 - Largest multimodal dataset of sentiment and emotion analysis.
 - More than 23,000 sentence utterances from YouTube.
 - More than 3,000 YouTube videos.
 - More than 1000 speakers.
 - More than 250 topics.
 - All three modalities.
 - Manual transcriptions.
 - Phoneme-level Alignment (semi-automatic)
 - Annotated for sentiment and emotions (Likert scales)
 - 7-scale sentiment (CMU-MOSI, SST)
 - 4-scale emotions (happiness, sadness, anger, disgust, surprise, fear)

Keynotes



Dr. Bing Liu
University of Illinois at Chicago (USA)



Dr. Sharon Oviatt
Monash University (Australia)



Dr. Roland Göcke
University of Canberra (Australia)

Thank you for joining us today