Let me Finish

Automatic Conflict Detection Using Speaker Overlap

Felix Grezes

Justin Richards

Prof. Andrew Rosenberg

CITY UNIVERSITY OF NEW YORK **THE GRADUATE CENTER**

INTERSPEECH LYON France 25-29 AUGUST 2013

Q QUEENS COLLEGE

Baseline Recap

- Kitchen-sink approach: ~3k low-level features

- Results fairly good — 79.1 % UAR*

*How can we improve performance while
clarifying the model?*

* on train + dev data

Our Approach

- Motivated by prior work, linguistic knowledge

- Hypothesis:

  The proportion of overlapping speech in recorded discussions should be a strong predictor of conflict level.

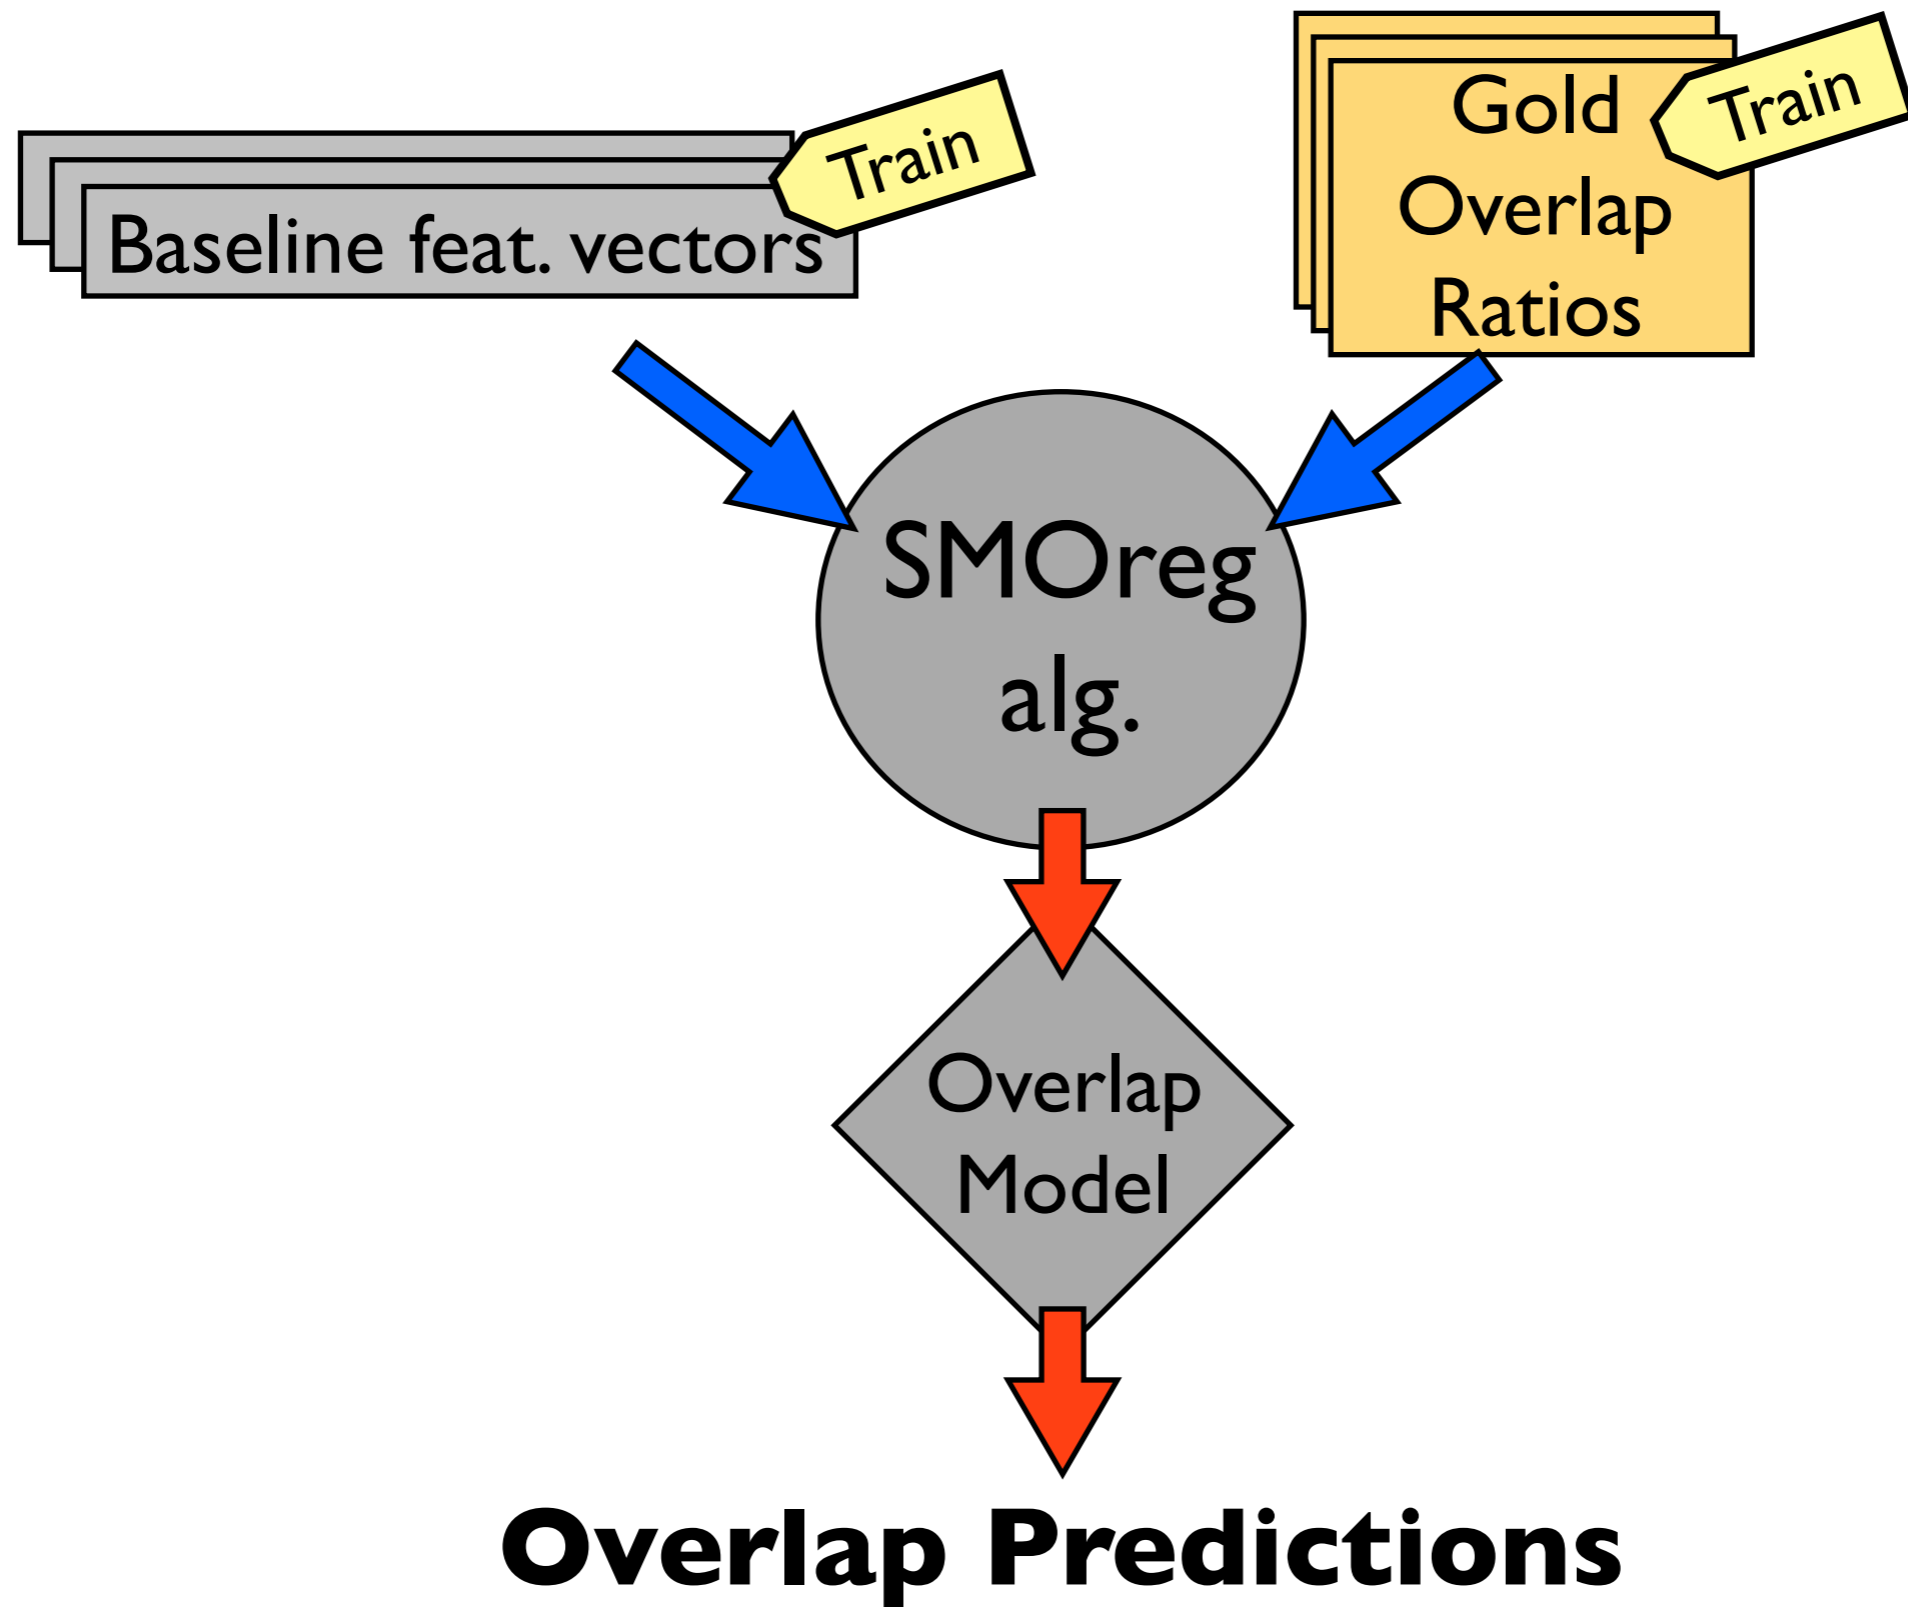- Testable and interpretable

Gold Overlap

Added feature:

- Corpus contains hand-labeled speaker turns

- Gold (actual) speaker overlap ratios were computed from this meta-data

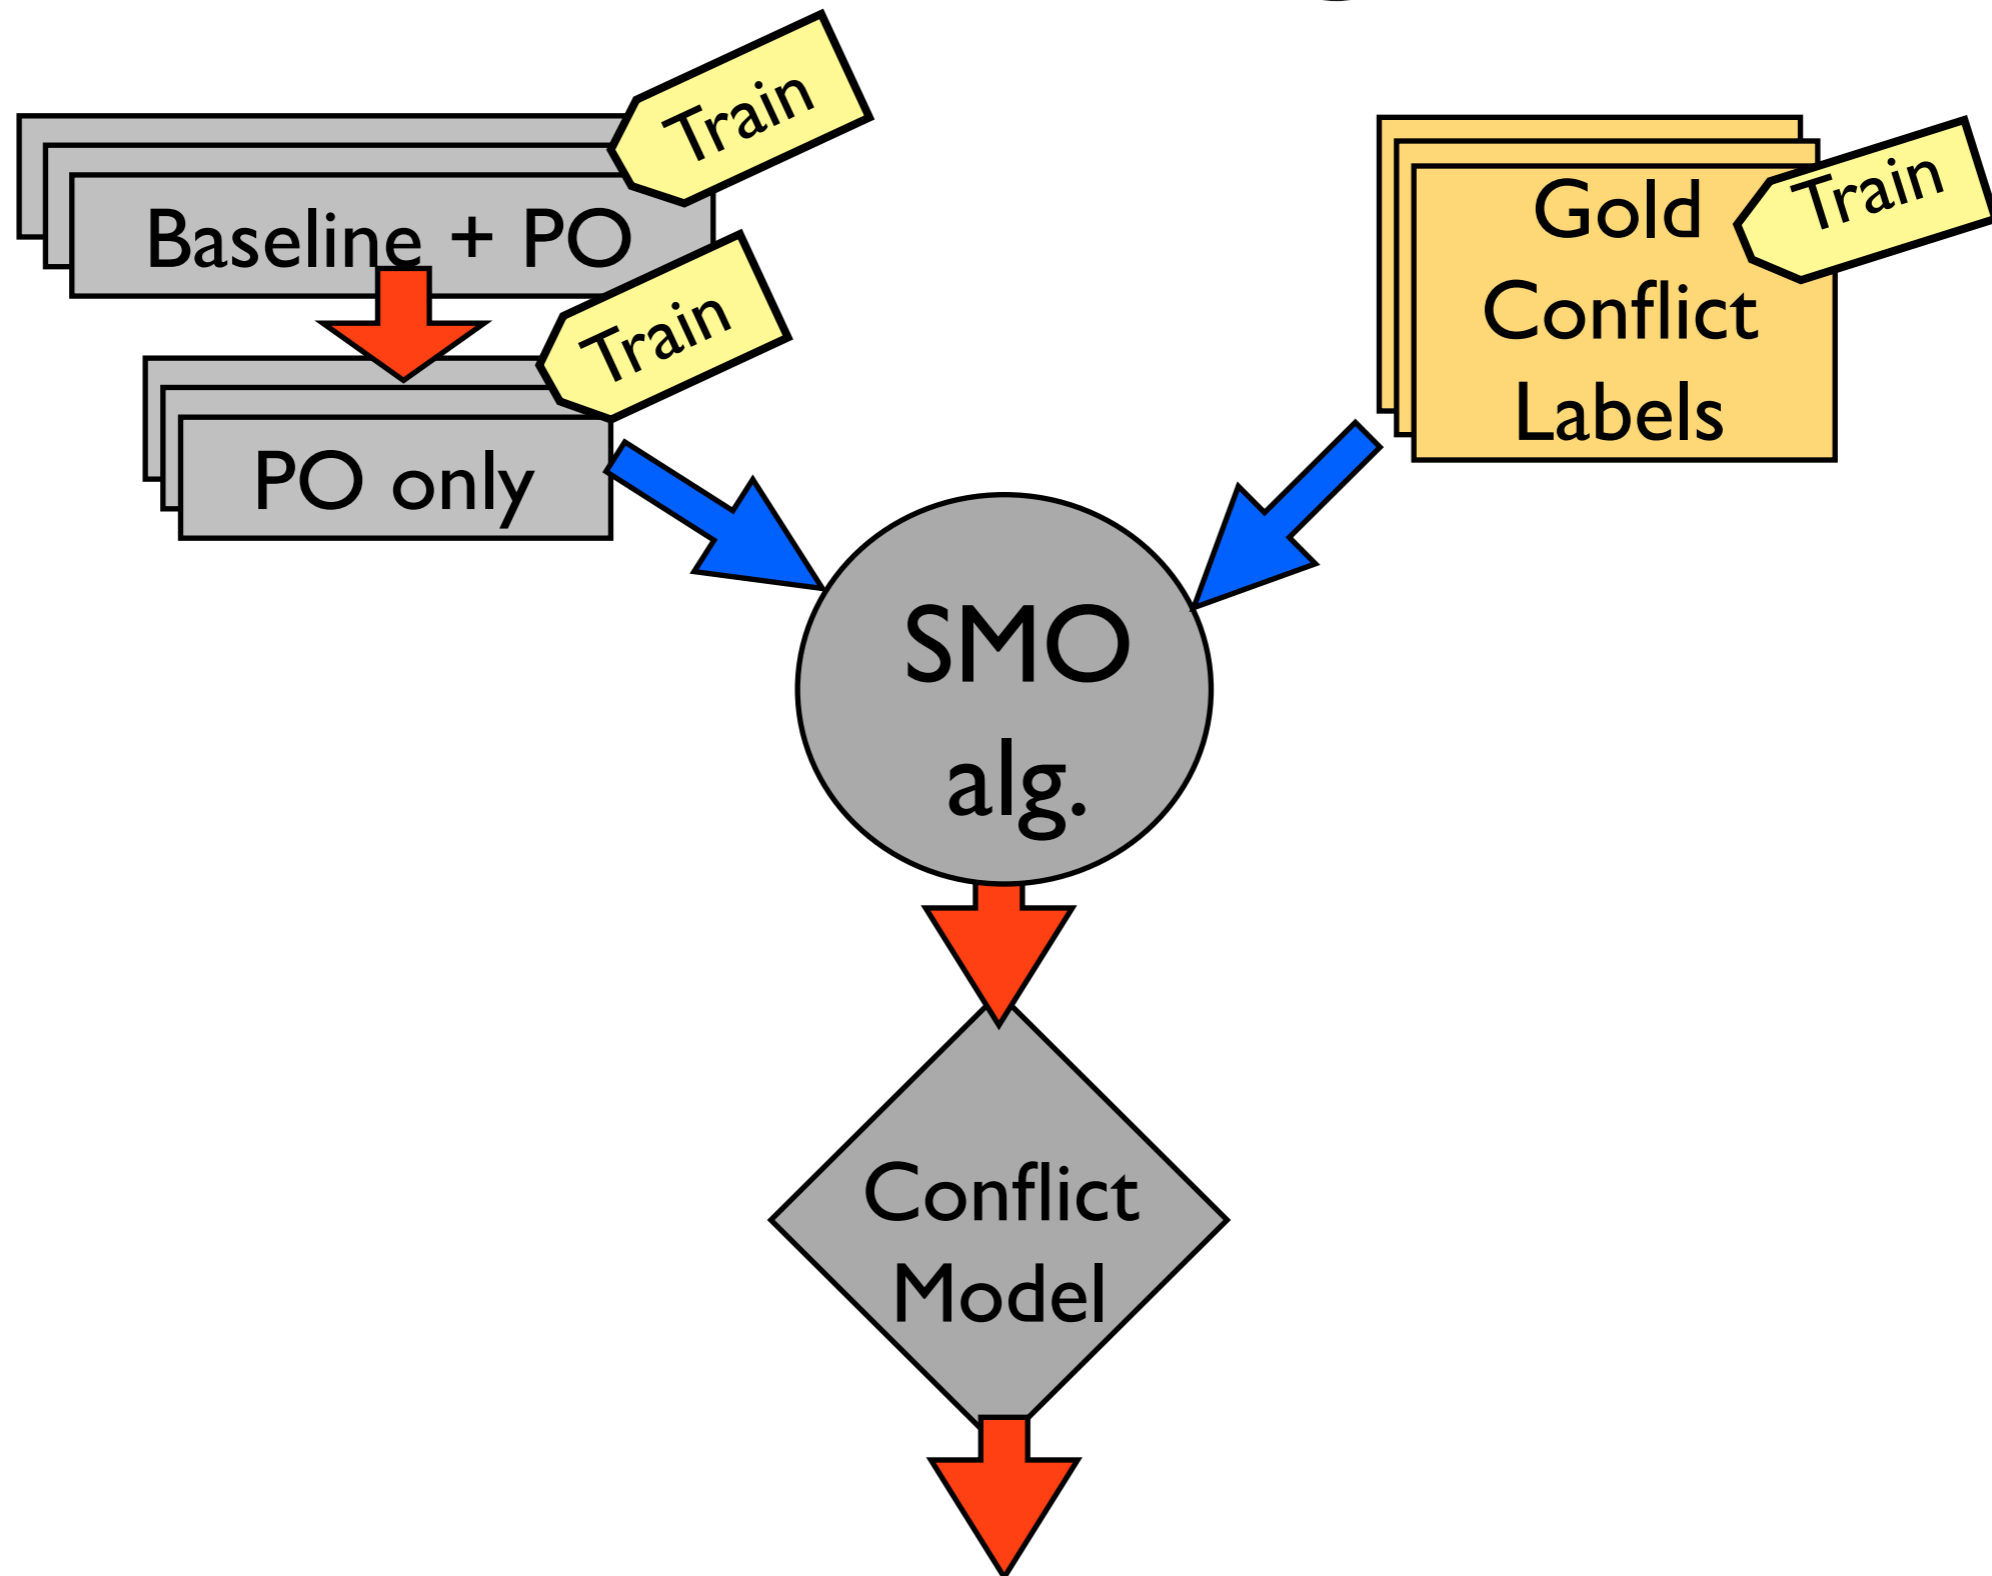- Conflict classifier built upon gold overlap ratio

Results:

- 74.2 % UAR* using gold overlap + baseline features

- 79.4 % UAR* using ONLY gold overlap

* on train + dev data

Model: Predicted Overlap

Model: Predicting Conflict

Conflict Predictions

Experiment Results

Train + Dev Data:

- 80.5 % UAR using **only** predicted overlap

  o Baseline: 79.1% UAR using ~3k features

  o Gold overlap: 79.4% UAR using one feature

- Adding even one extra feature, performance *decreases*.

Pred Overlap > Gold Overlap?

Theories:

- Finds instances of overlap that were missed by hand-annotators.

  o Interjections, failed interruptions

- Identifies a pressured or competitive quality in speech that is independent of overlap
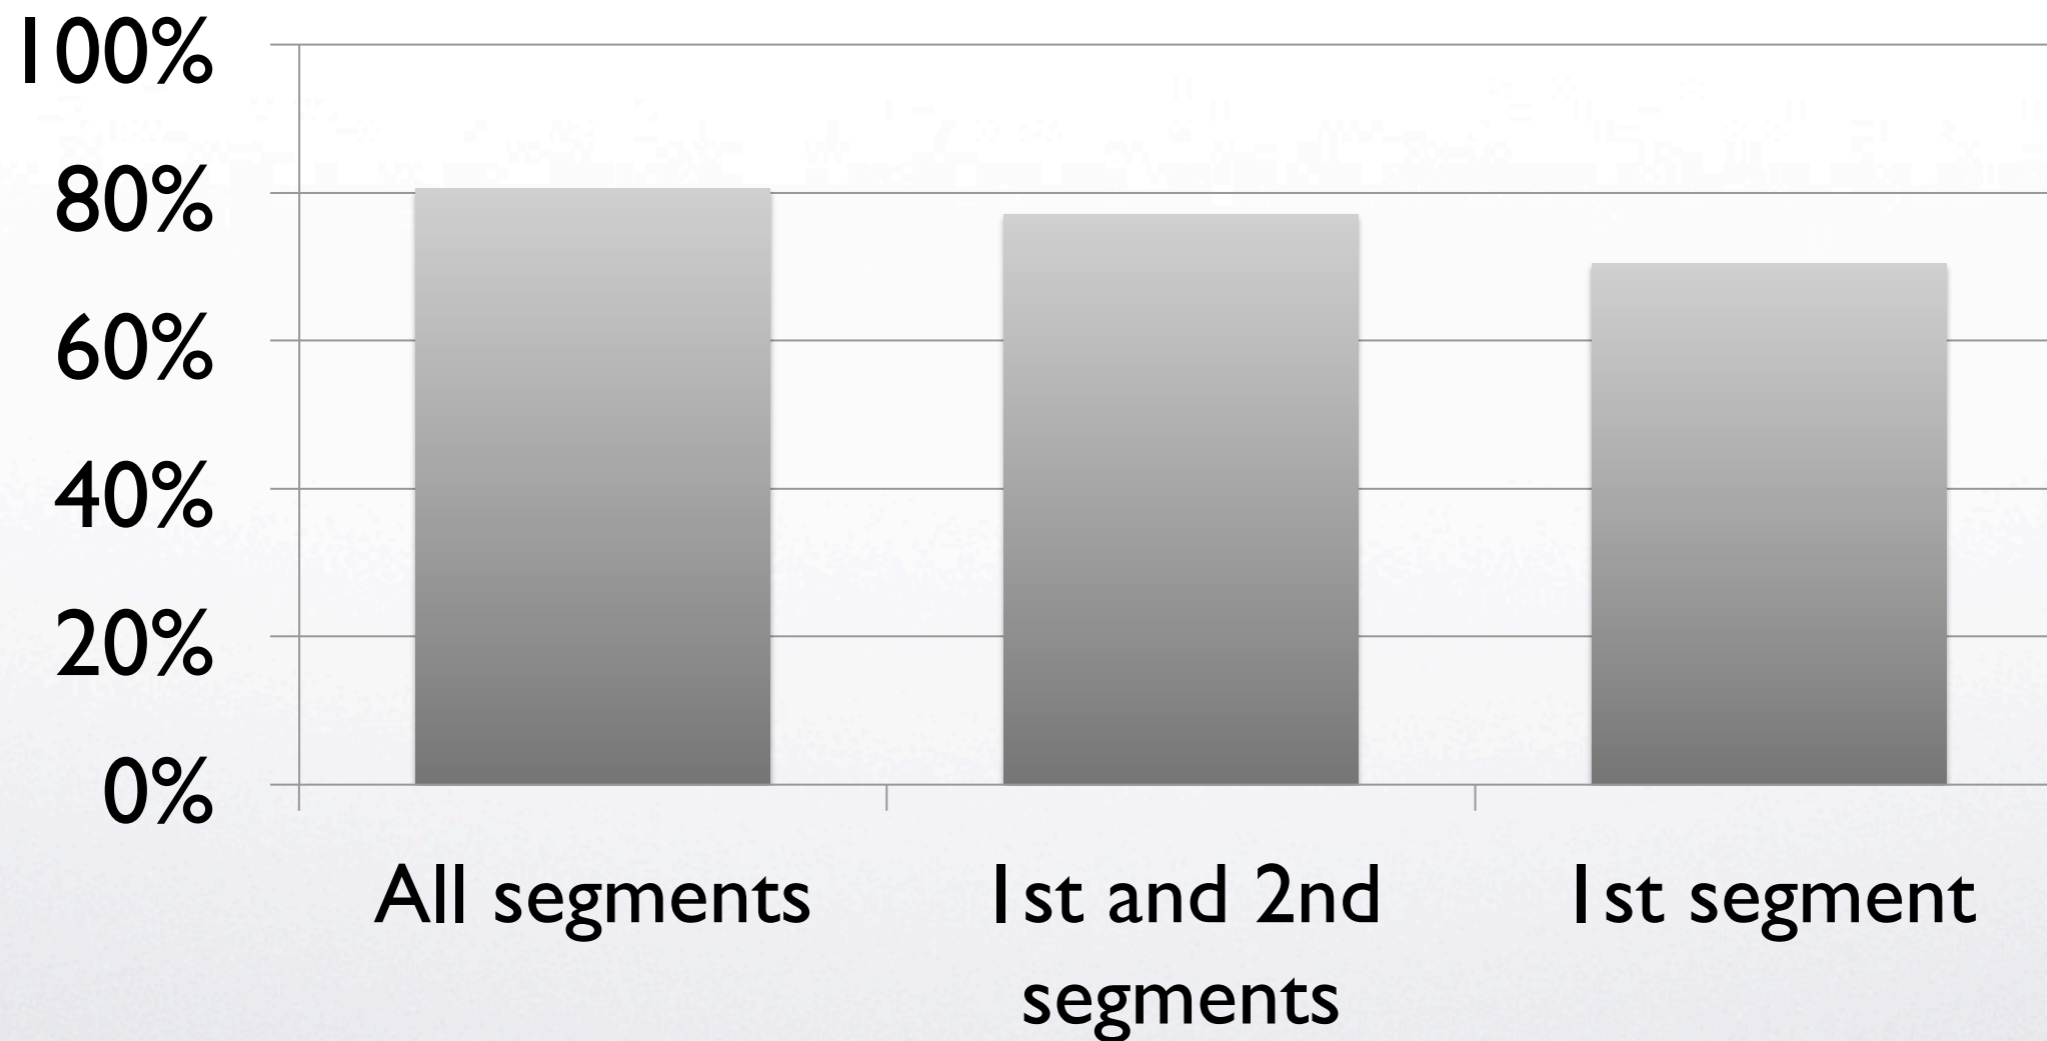
Performance

Test Data:

- 83.1 % UAR using **only** predicted overlap

- Baseline: 80.8 % UAR using ~3k features

We improved performance using a clearer, more intuitive model.

Segmentation & Degradation

Thank You!

Contact:

fgrezes@gc.cuny.edu
jrichards@gc.cuny.edu
andrew@cs.qc.cuny.edu

The Speech Lab at Queens College:

speech.cs.qc.cuny.edu