



University
of Colorado
Boulder

$2 * n$ is better than n^2 : Decomposing Event Coreference Resolution into Two Tractable Problems

Shafiuddin Rehan Ahmed, Abhijnan Nath, James H. Martin, Nikhil Krishnaswamy

{shah7567, james.martin}@colorado.edu, {abhijnan.nath, nkrishna}@colostate.edu




Colorado
State
University

Abstract

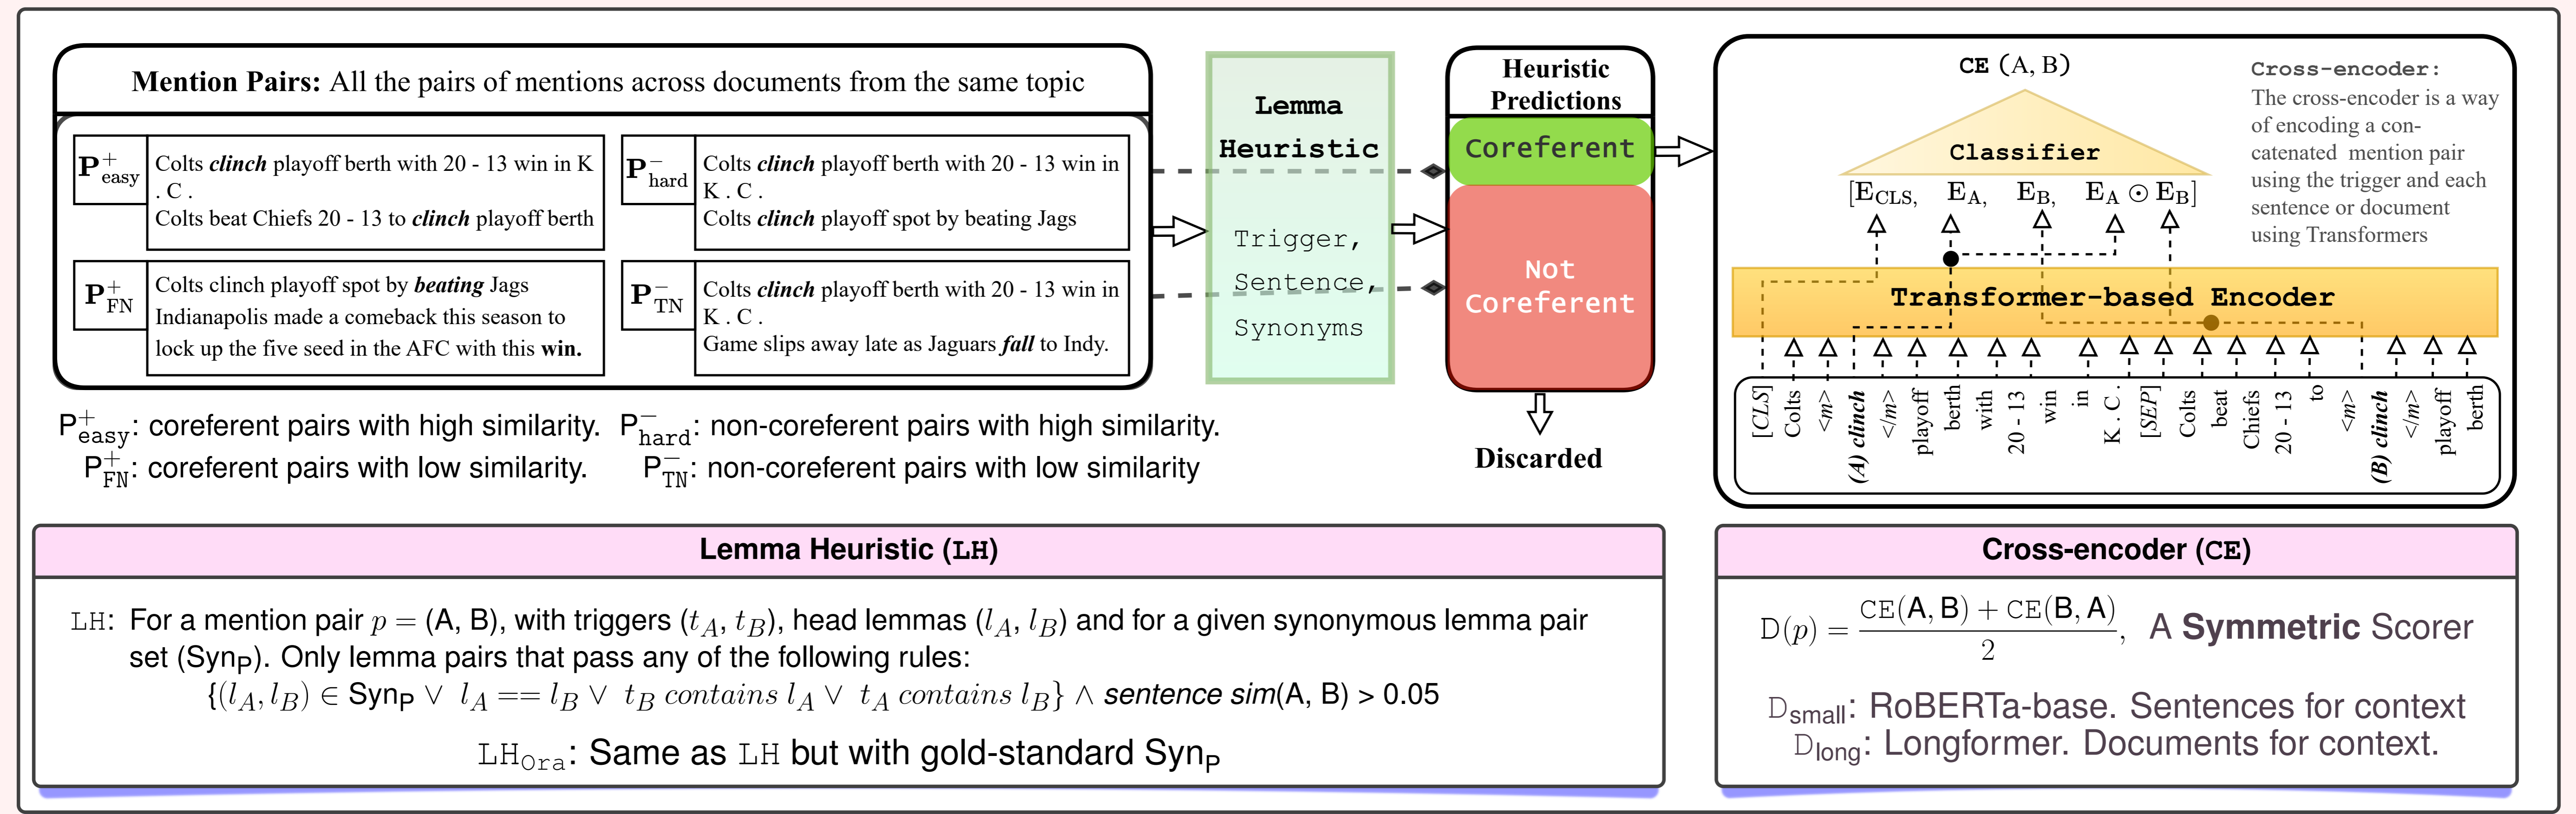


Existing methods for training Event Coreference Resolution (ECR) systems sample from a largely skewed distribution, making it difficult for the algorithm to learn coreference beyond surface matching. Additionally, these methods are intractable because of the quadratic operations needed. To address these challenges, we break the problem of ECR into two parts: a) a heuristic to efficiently filter out a large number of non-coreferent pairs, and b) a training approach on a balanced set of coreferent and non-coreferent mention pairs. By following this approach, we show that we get comparable results to the state of the art on two popular ECR datasets while significantly reducing compute requirements. We also analyze the mention pairs that are "hard" to accurately classify as coreferent or non-coreferent

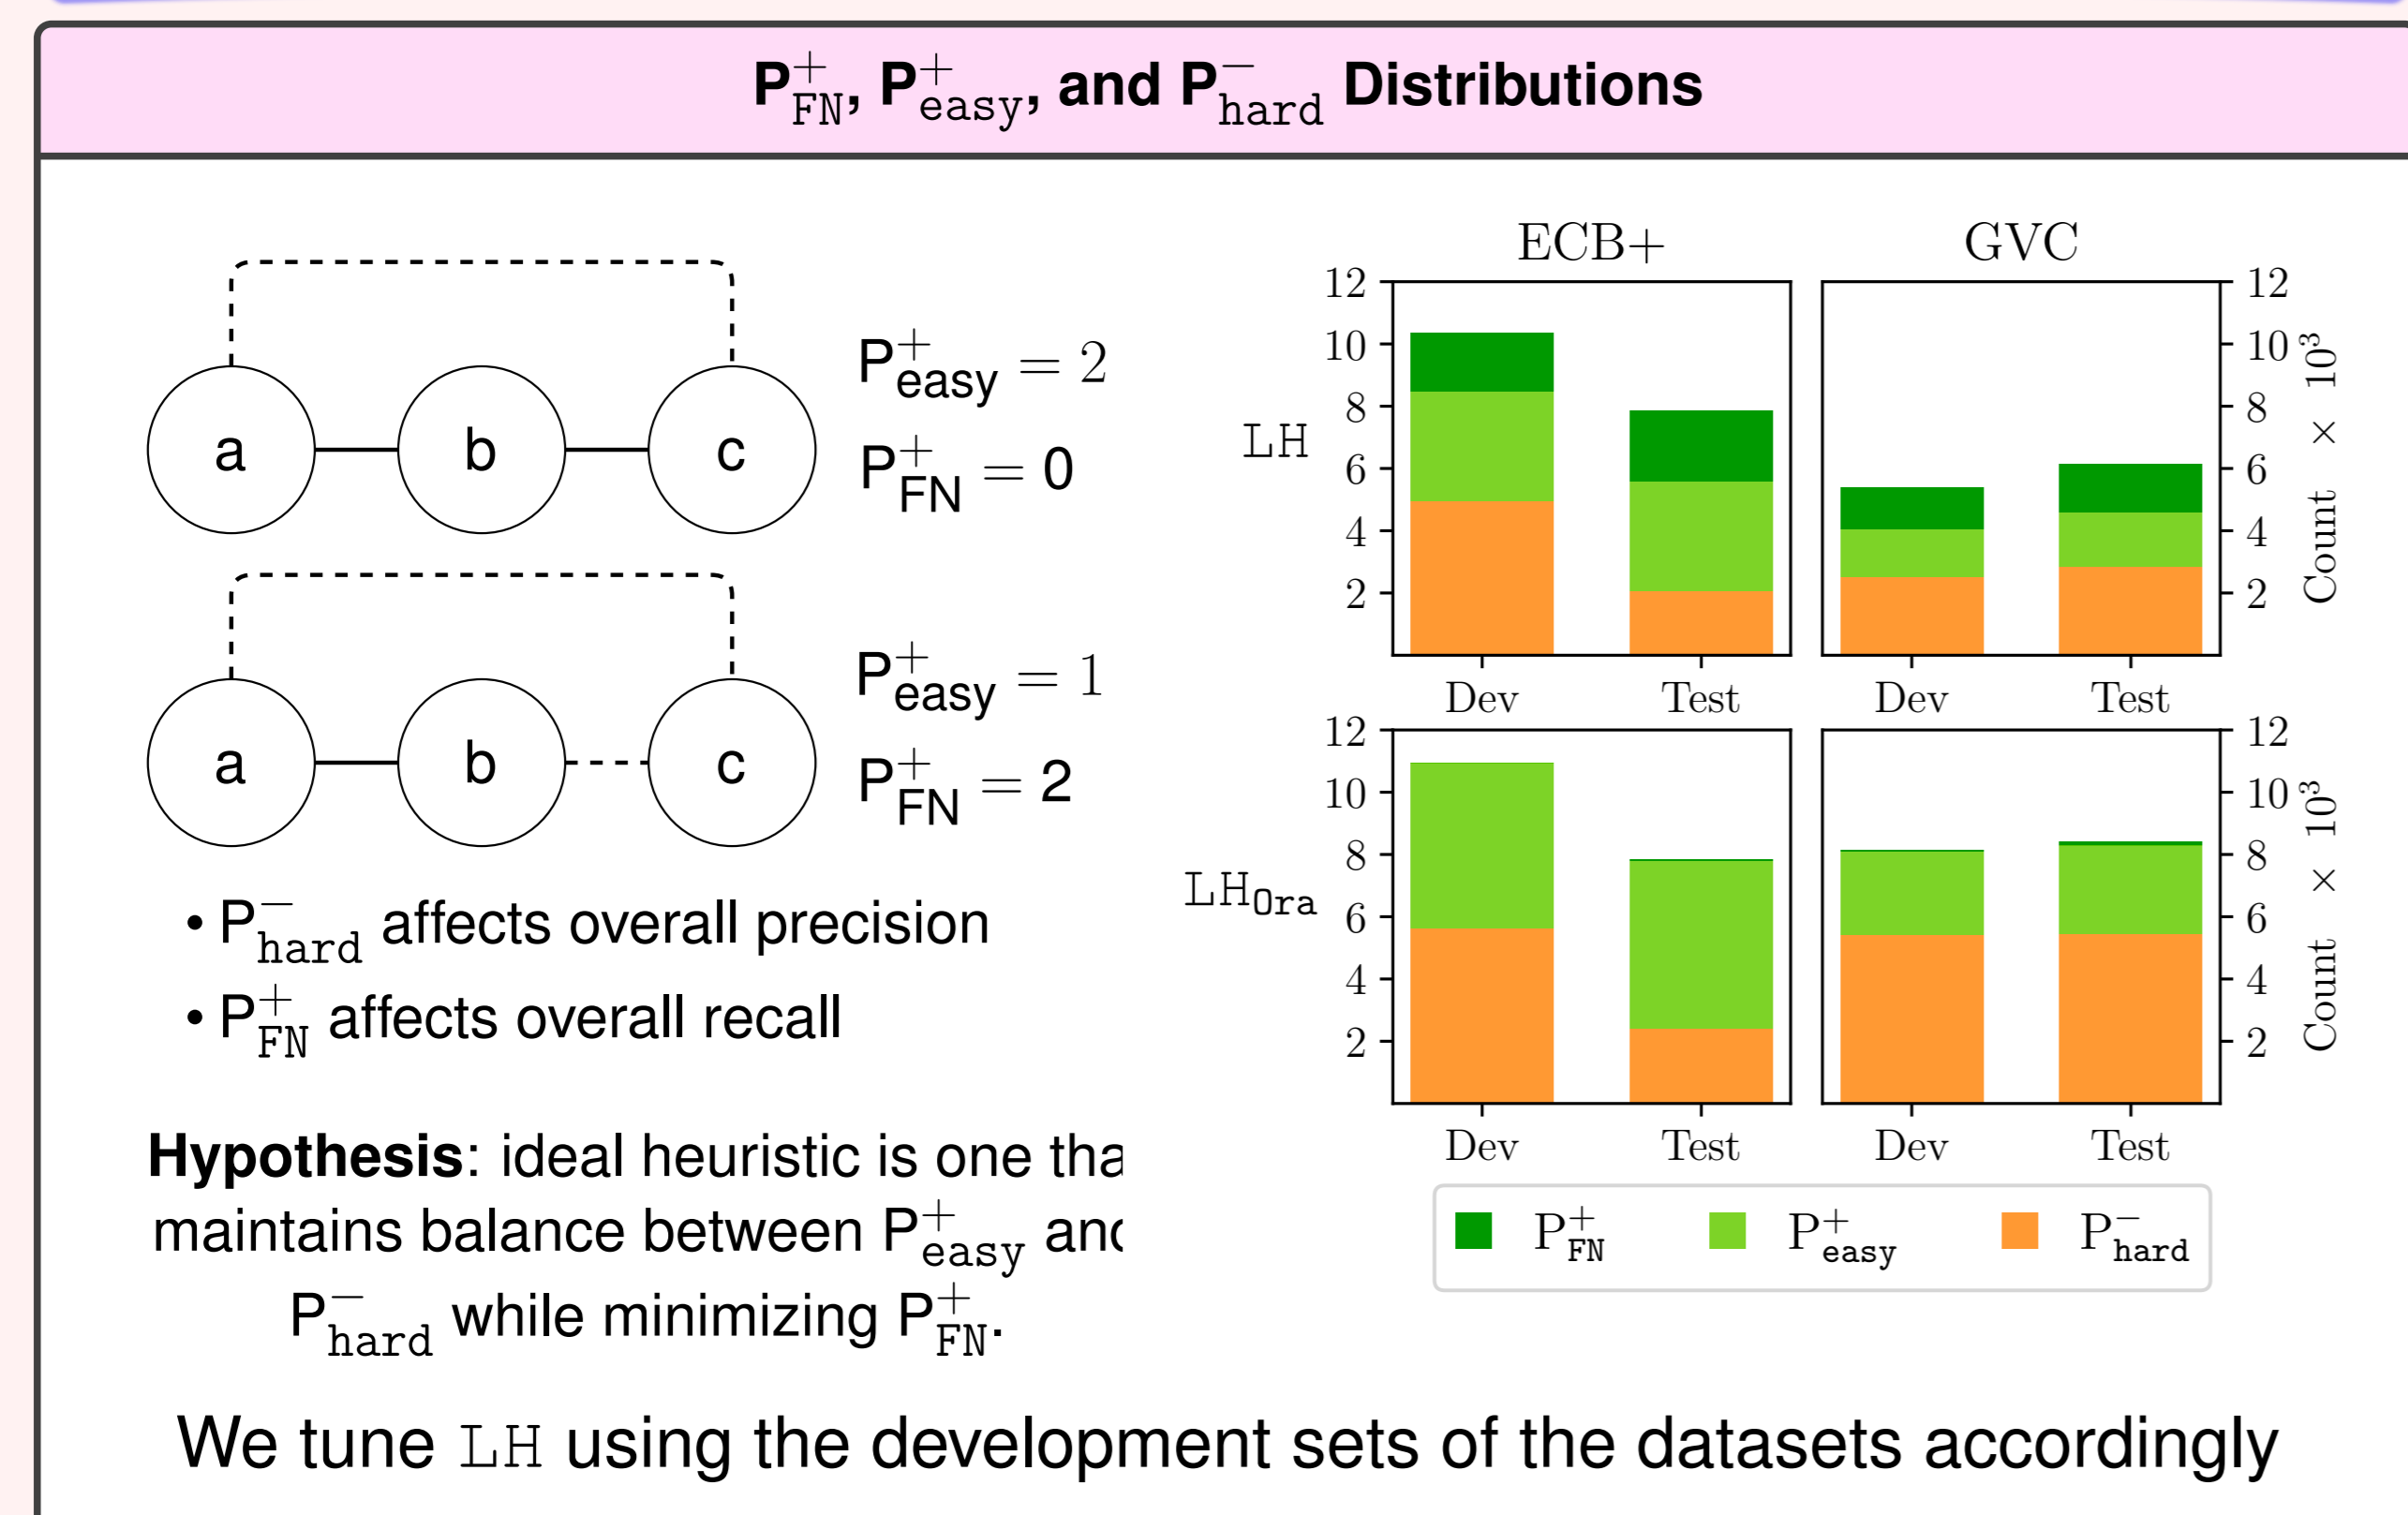
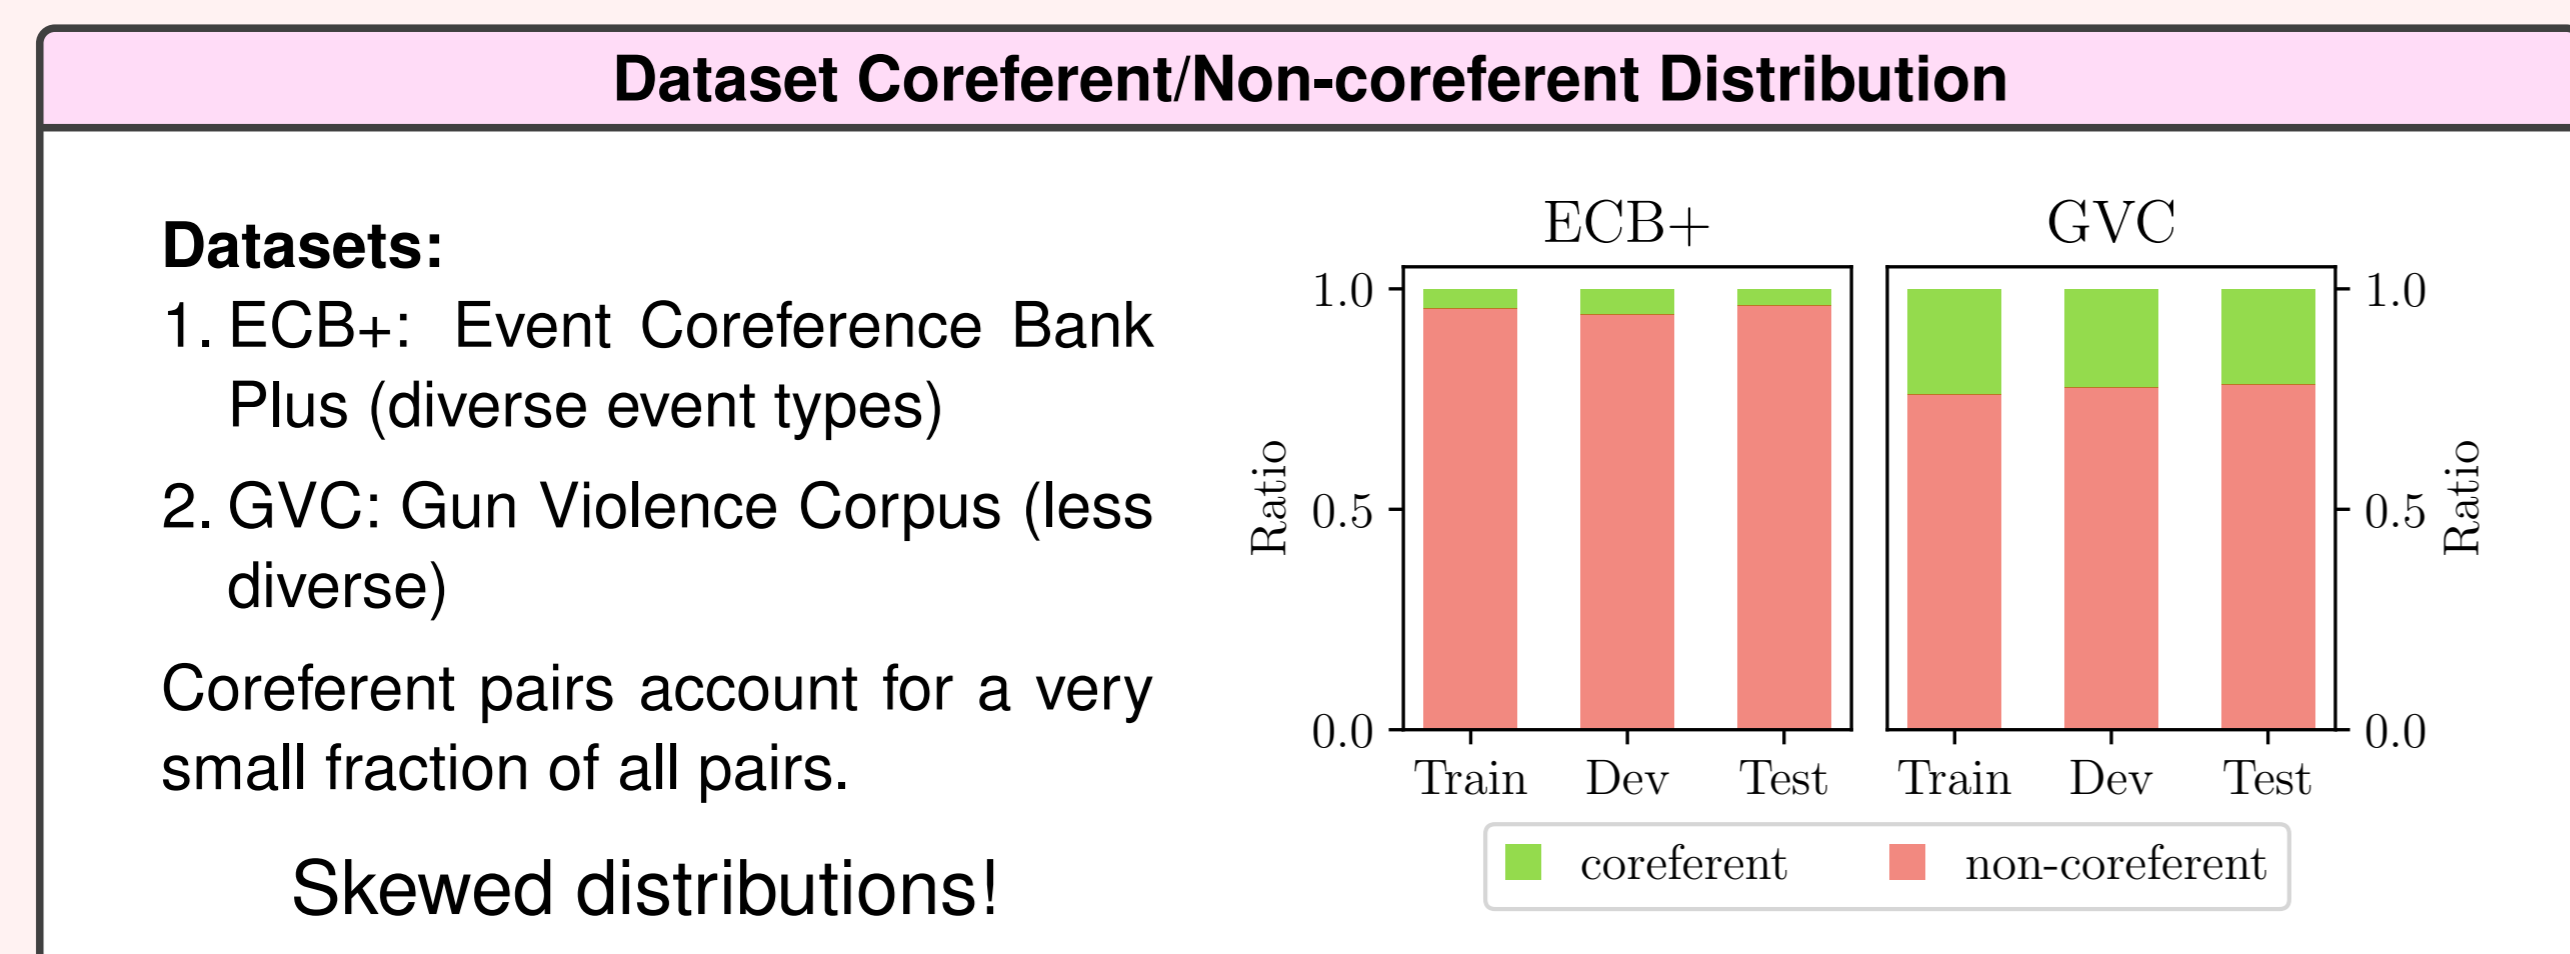
Author



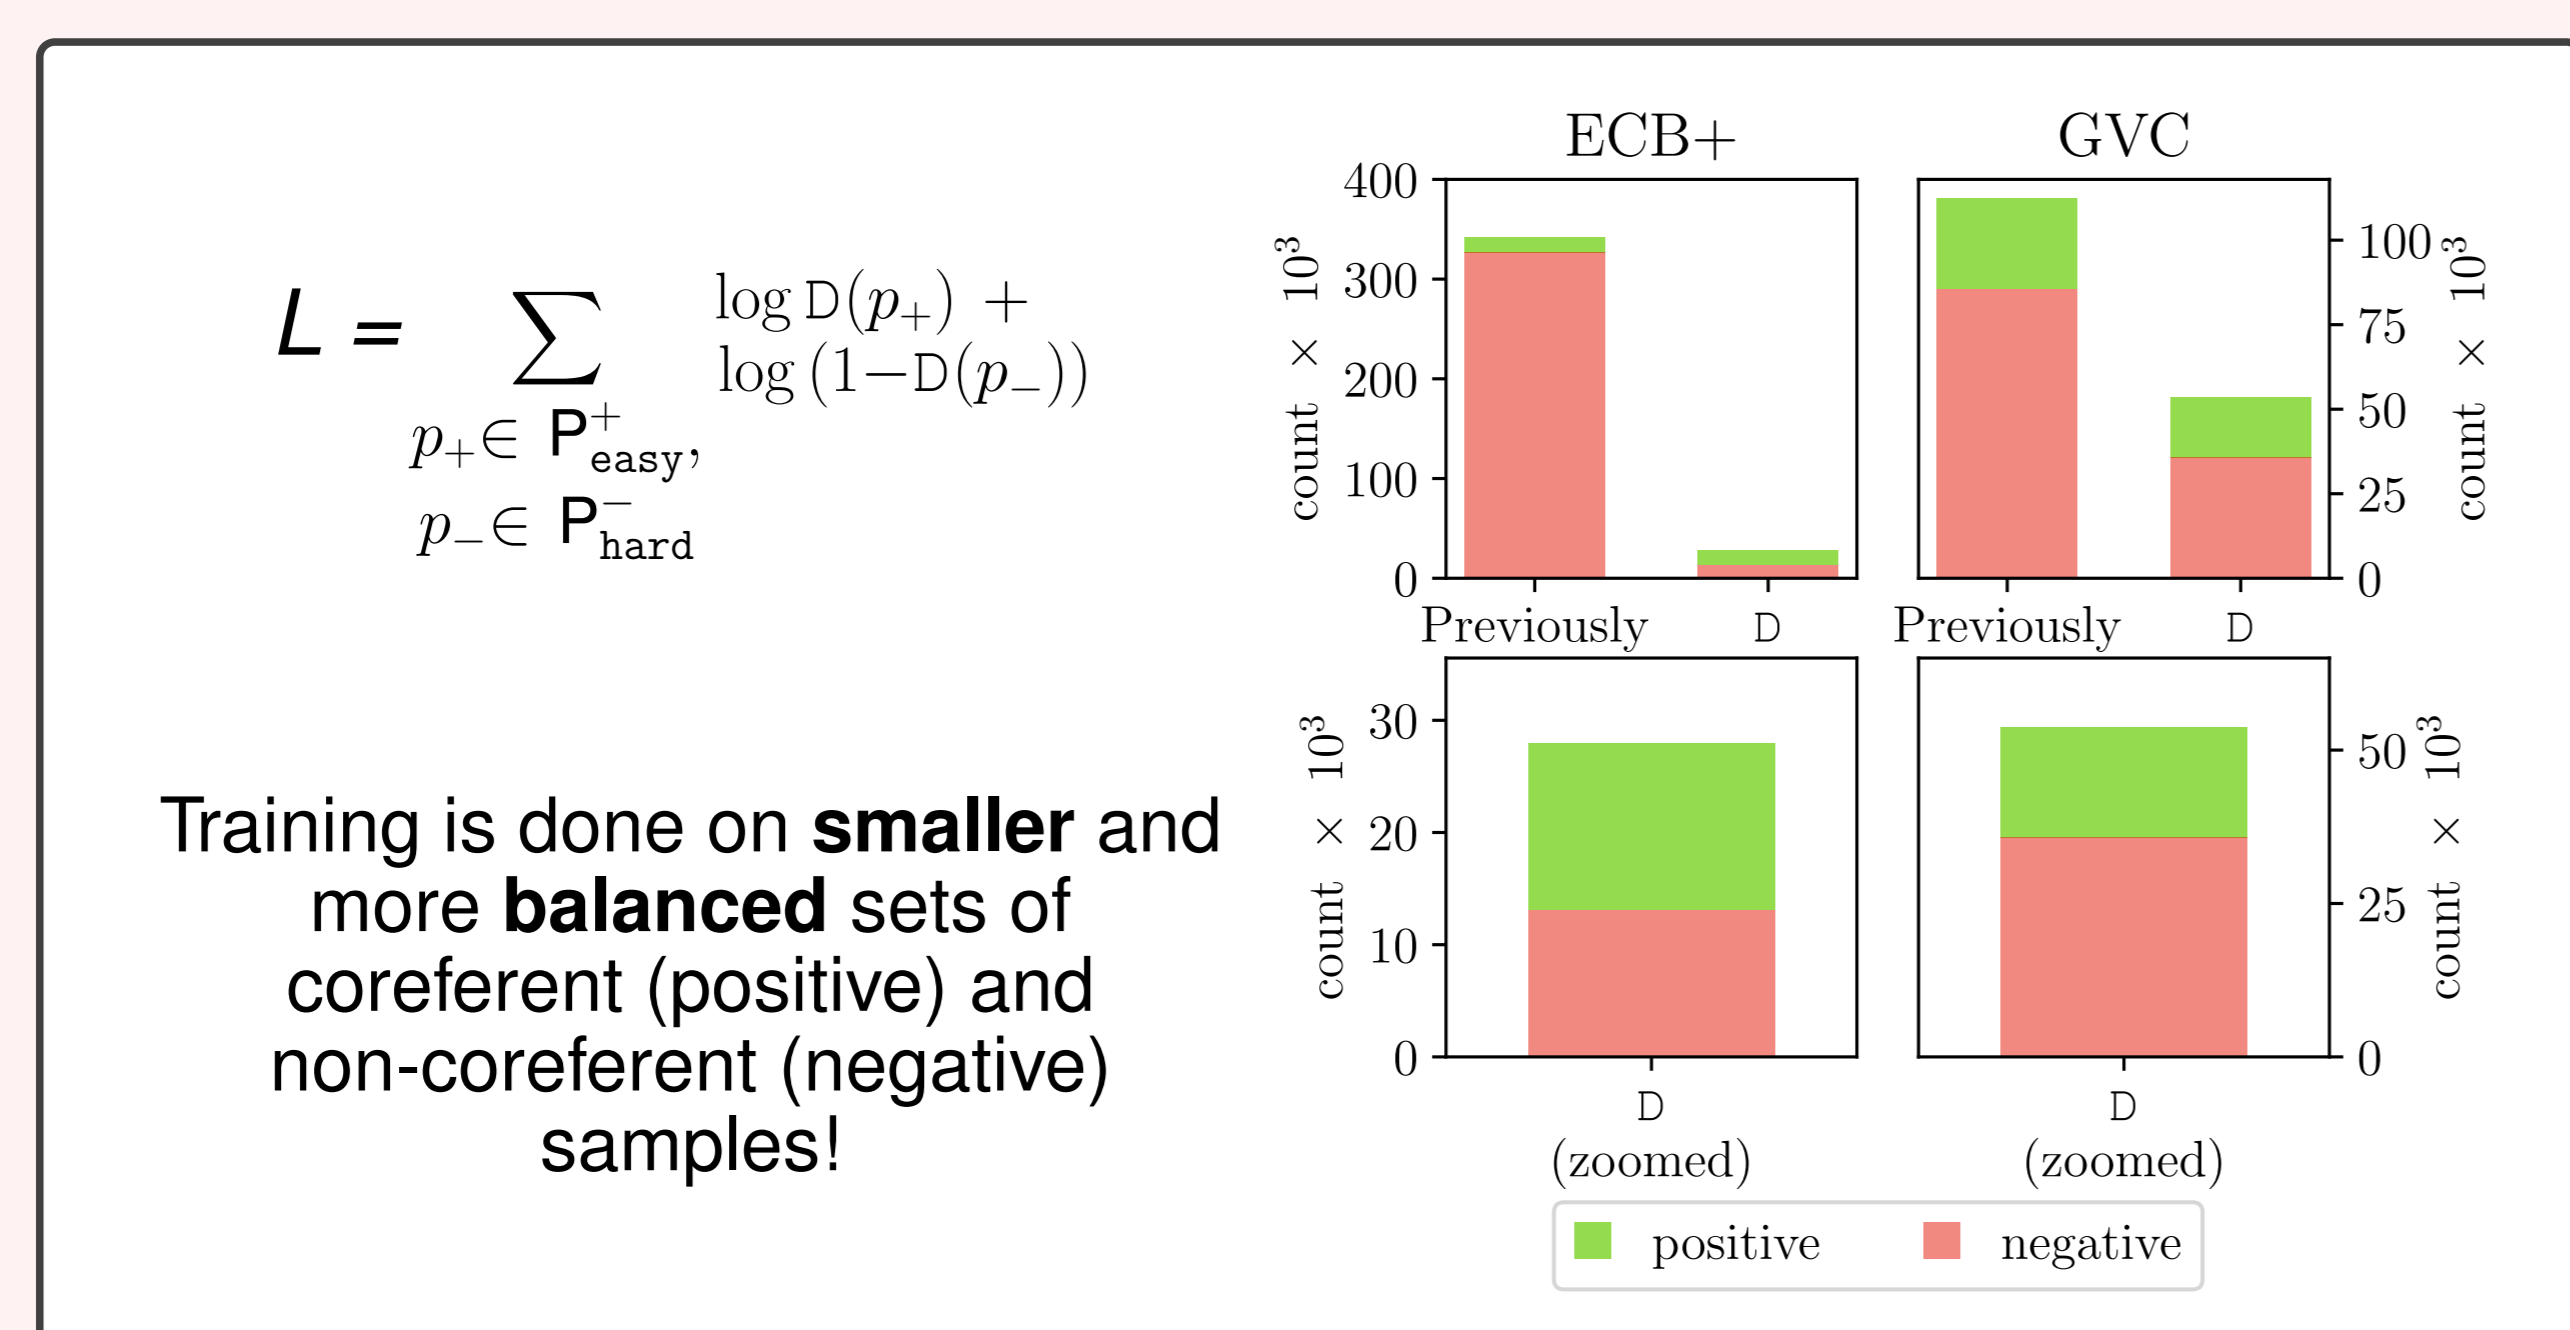
Lemma Heuristic (LH) & Cross-encoding-based Coreference Scorer (CE)



Data Analysis with LH and LH_{Ora}



Training using P_{easy}^+ & P_{hard}^- only



Results

Methods	CoNLL F_1	
	ECB+	GVC
Bugert et al. (2021)	-	59.4
Cattan et al. (2021)	81.0	-
Caciularu et al. (2021)	85.6	-
Held et al. (2021)	85.7	83.7
LH	76.4	51.8
LH + D_{small}	80.3	73.7
LH + D_{long}	81.7	75.0
LH_{Ora}	81.9	53.4
LH_{Ora} + D_{small}	85.9	75.4
LH_{Ora} + D_{long}	87.4	76.1

Key Takeaways

Comparable performance with a significant **reduction in complexity!**
 LH_{Ora} results for ECB+ a **jumping-off point** to **improve ECR** with heuristics
Synonymous Lemma Pair identification for **GVC** is **not useful**

Cross-encoder Error Analysis using LH_{Ora} (Dev sets)

Precision-based Errors	Recall-based Errors
Cluster impurity: # of non-coreferent mention pairs in a cluster. We trace the D predictions and analyze them. Key findings: Hard bridging pairs - major cause of errors Set-member relations cause a lot of errors Singletons appear in least impure clusters	Errors categorized as: <ul style="list-style-type: none"> • same-sentence pro-nouns • weak temporal reasoning • ambiguity due to corefering entities • lexically-different but semantically similar event mention lemmas

Conclusions

- We showed a simple LH + CE does comparable ECR while being computationally efficient.
 - Presented shortcomings of CE through extensive error analysis and suggested techniques to improve performance upon
- Future Work**
- Annotations/automated approaches for detecting synonymous lemma pairs.
 - LH + CE + visual transformers for multimodal ECR

References

Bugert et al. (2021): Generalizing Cross-Document Event Coreference Resolution Across Multiple Corpora

Cattan et al. (2021): Cross-document Coreference Resolution over Predicted Mentions

Caciularu et al. (2021): CDLM: Cross-Document Language Modeling

Held et al. (2021): Focus on what matters: Applying Discourse Coherence Theory to Cross Document Coreference