

Summary

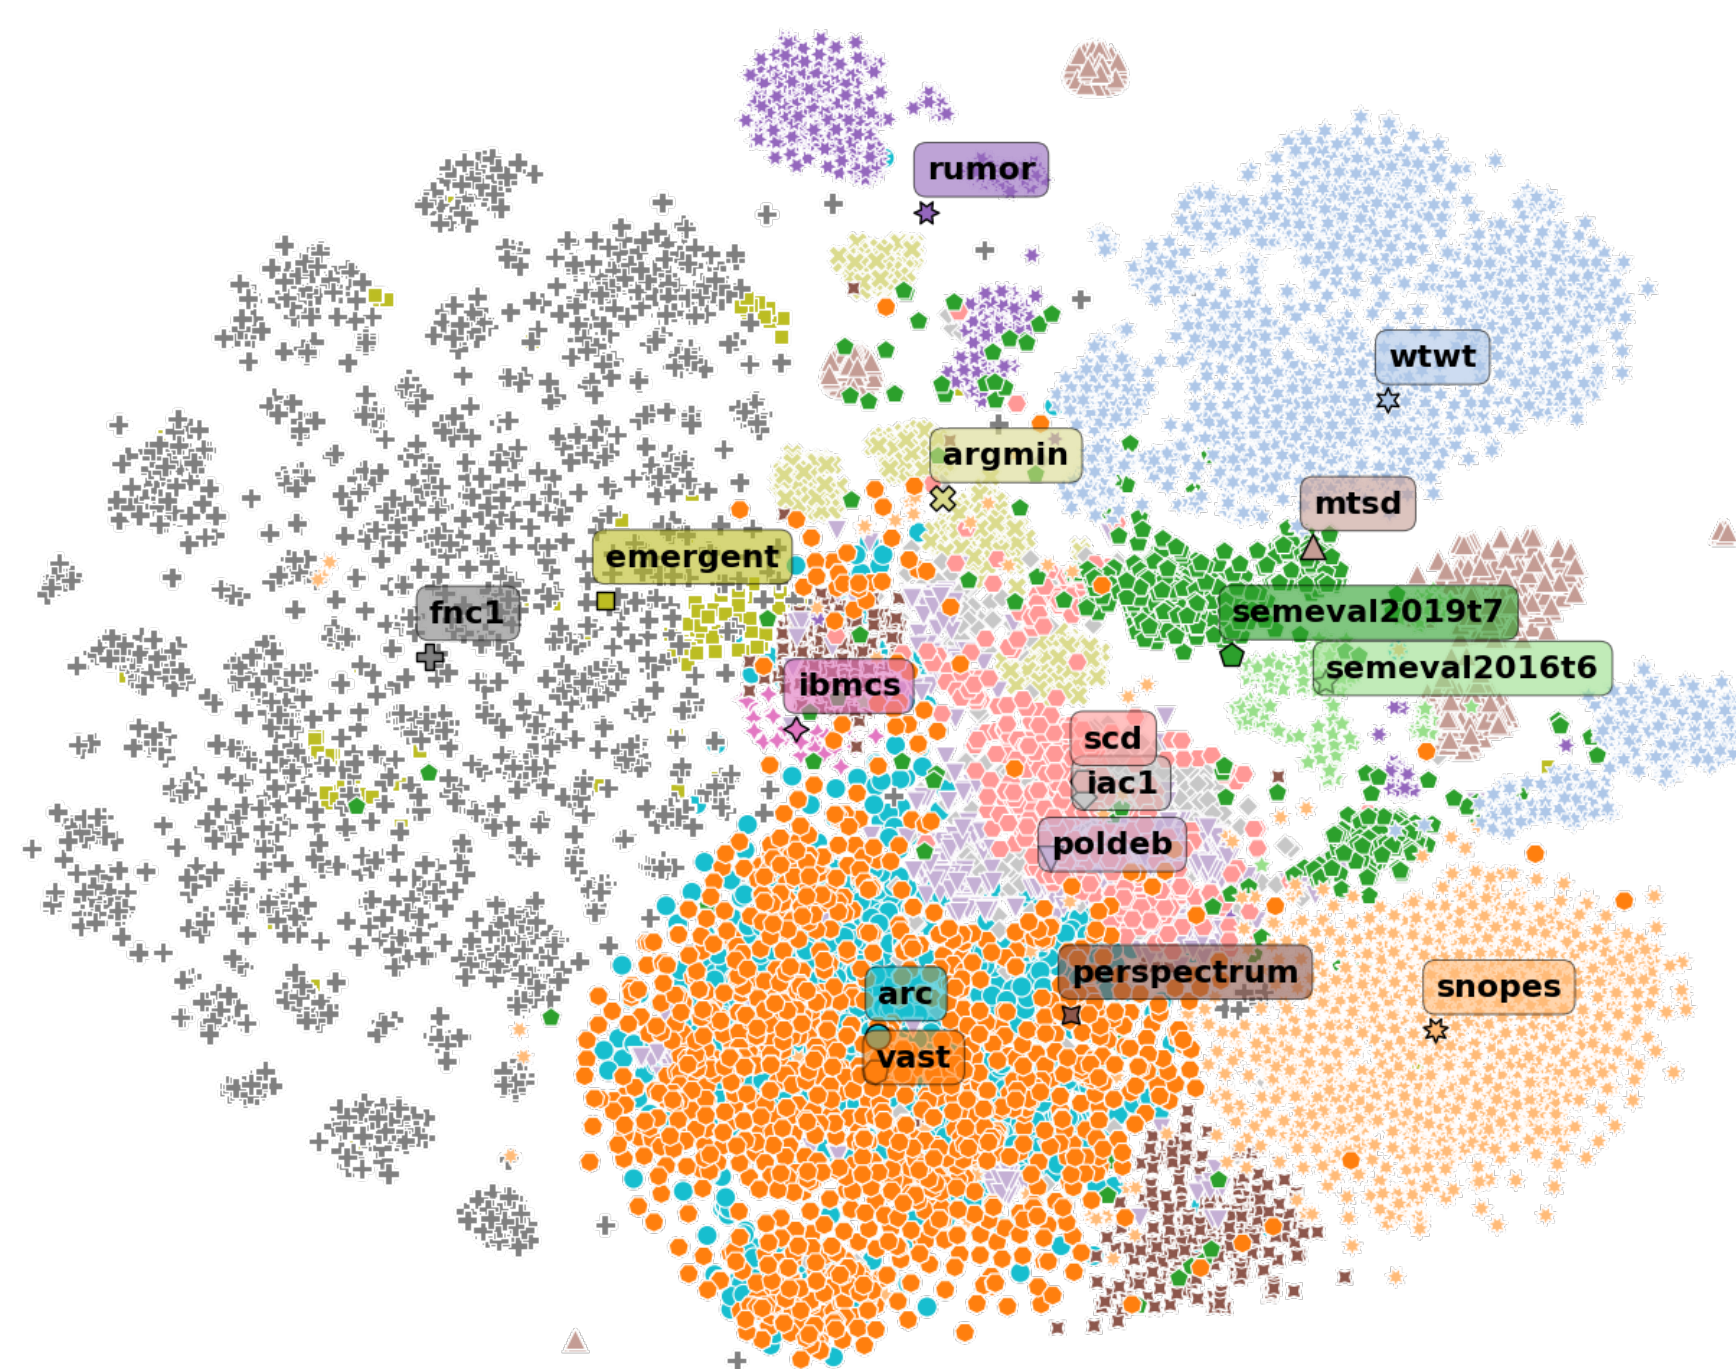
- Stance expressions can take **different forms**, e.g.
 - a tweet targeting a person
 - an article addressing a claim
 - a comment for a headline
- Datasets use **different labels** (*agree, support, argument for*, etc.)
- Data comes from **various domains** (rumours, news, politics, etc.)
- The above issues limit generalisation and cross-domain studies
- We introduce MoLE, which combines
 - Mixture of Experts
 - Domain-Adversarial Training
 - Label Embeddings
- MoLE can adapt to new domains and can work with unseen labels
- Strong performance in both in-domain and out-of-domain scenarios
- Exhaustive correlation analysis of performance vs. dataset characteristics

Contributions

- The largest study of stance detection datasets to date
- Novel framework (**MoLE**) combining **Mixture of Experts and Label Embeddings**
- MoLE can adapt to **out-of-domain data** and can work with **unseen labels**
- MoLE outperforms both **in-domain** and **out-of-domain** baselines
- Correlation analysis** between performance and several dataset characteristics

Datasets

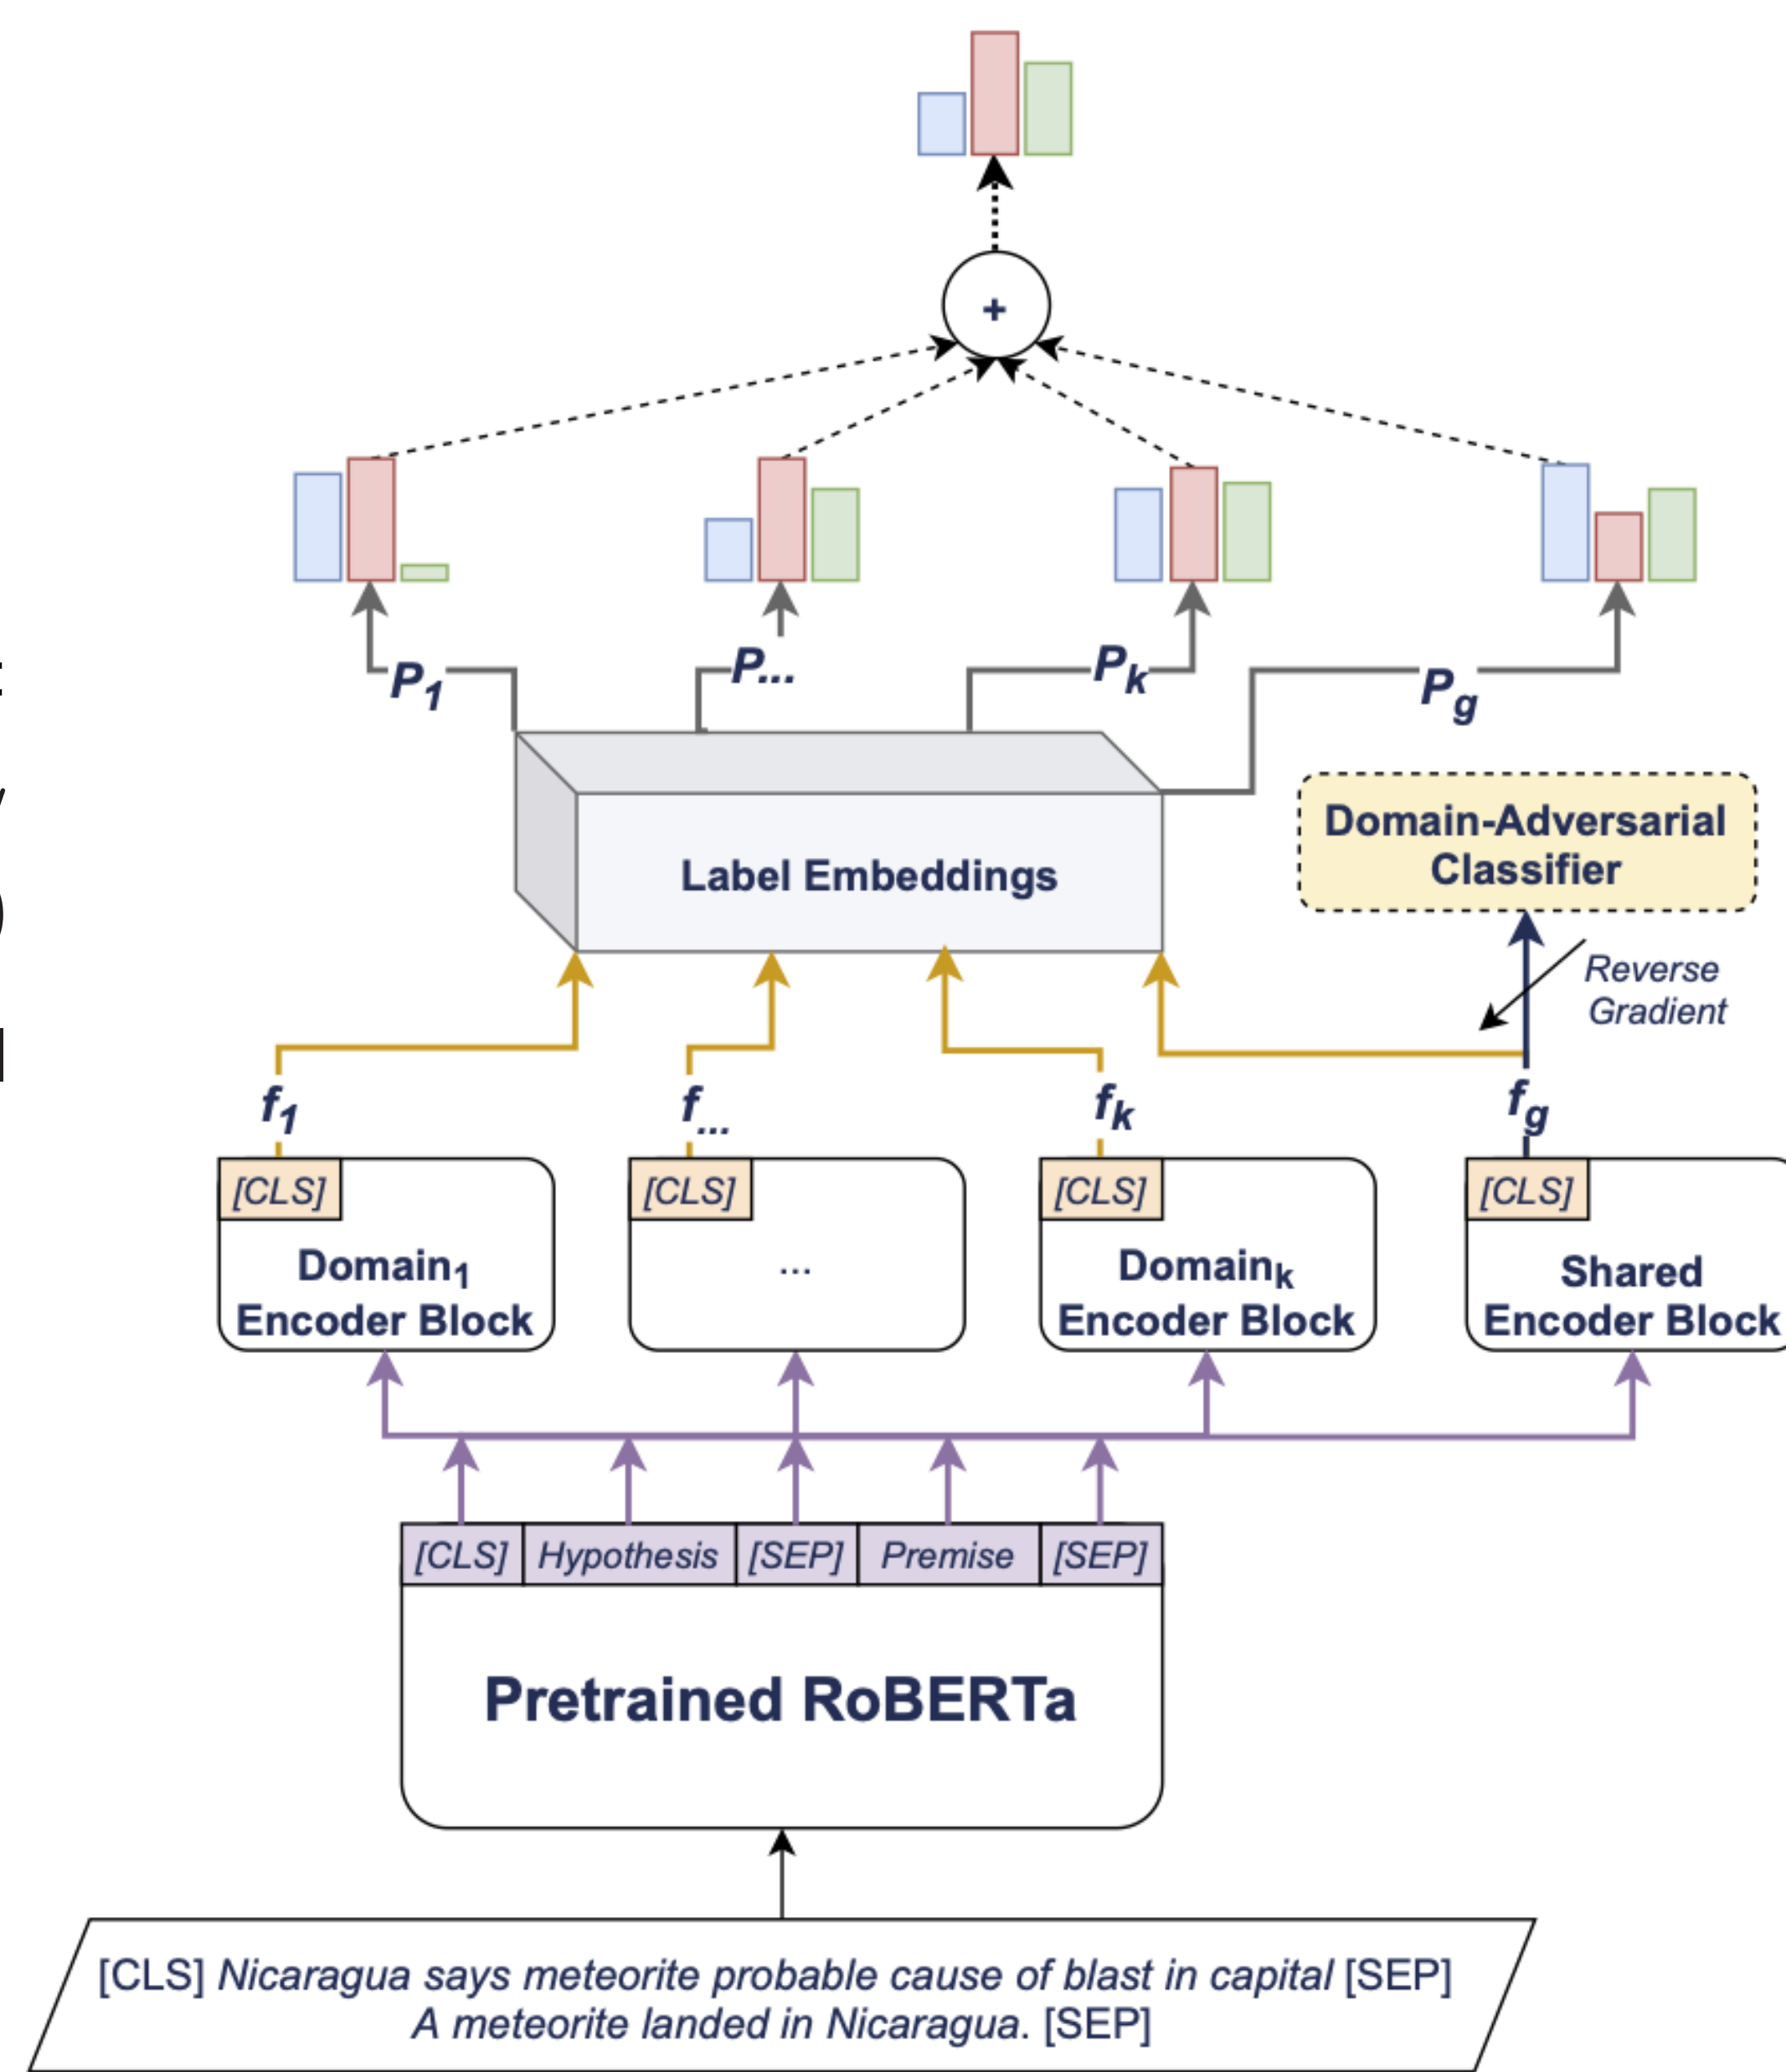
- Sixteen** diverse datasets
- Different **domains**
 - Fake News
 - Rumours
 - Politics
 - Company Acquisitions
- Over **fifty** labels
 - (well-known) Positive, Negative, etc.
 - (additional) Neutral, etc.
 - (unusual) Other, Question, etc.



(MoLE) Mixture of Experts with Label Embeddings

Key Characteristics (Our Model):

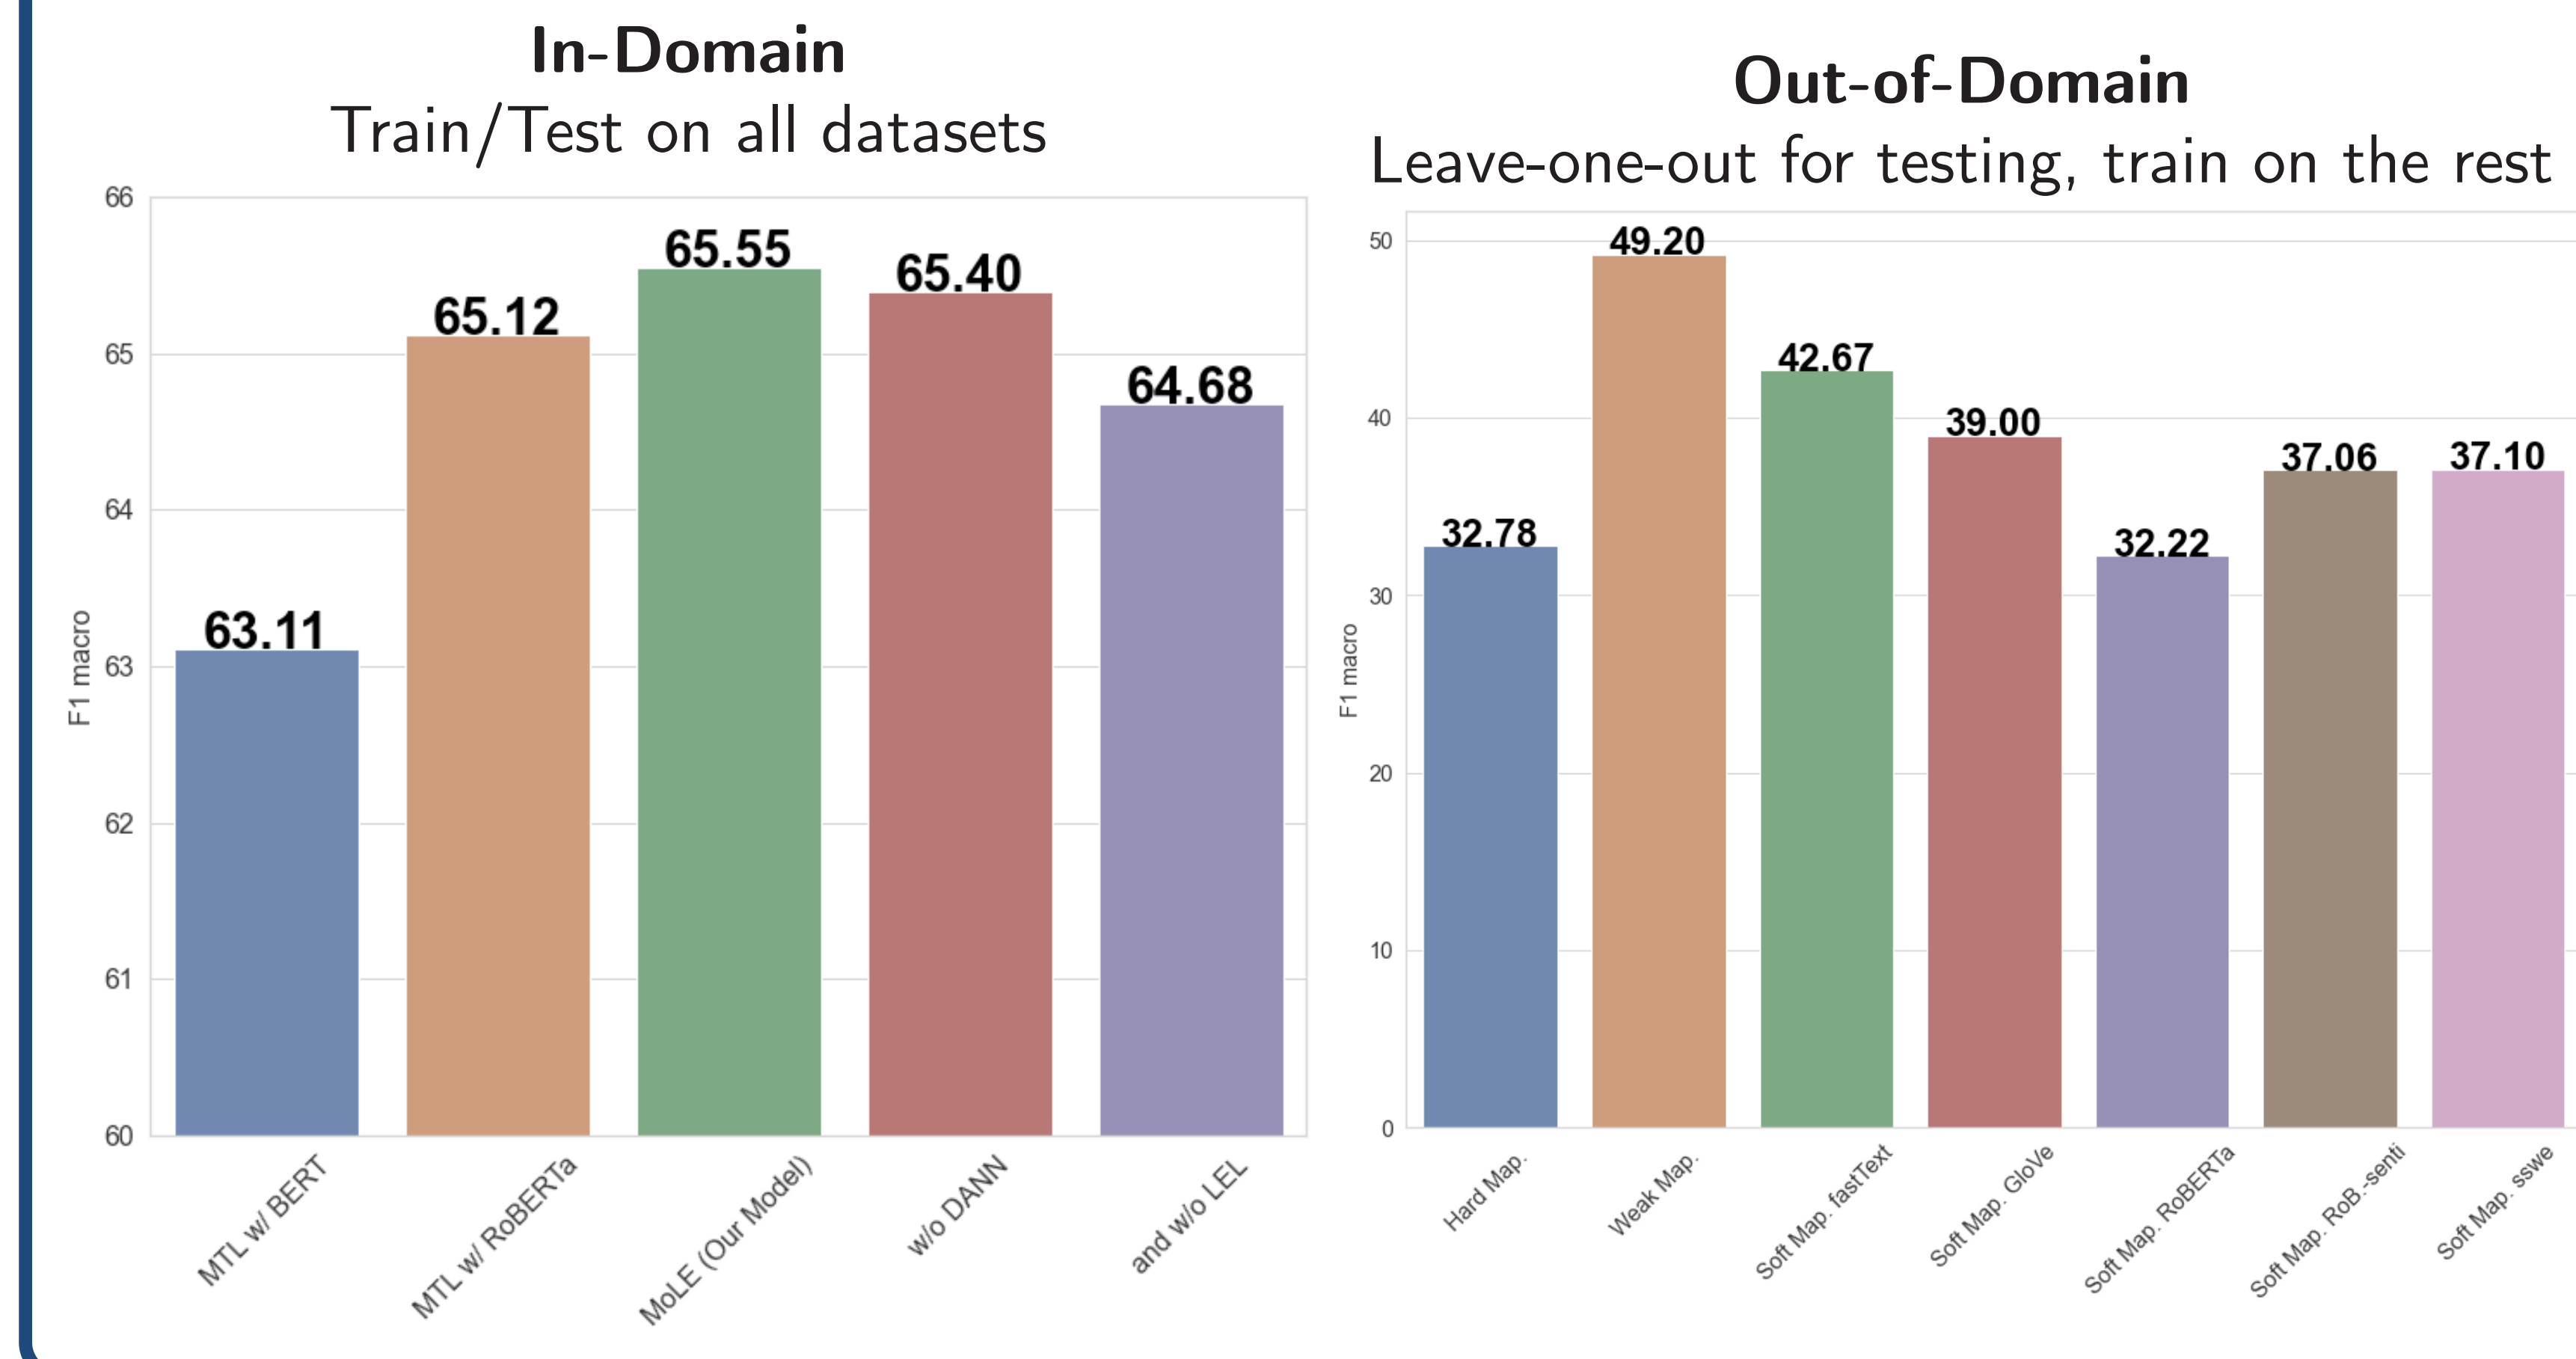
- Mixture-of-Experts (MoE) w/ Label Embeddings Layer (LEL)
- Domain-Adversarial Neural Network (DANN)
- Joint Label Space Voting
- Parameter-Efficient Model



Repository

Download the data, and try your own models!
<https://github.com/checkstep/mole-stance>

Results



Correlation Analysis

(Pearson) Correlation Analysis:

- Dataset size (more training data – **better**)
- Vocabulary size (richer vocab – **harder** task)
- Social Media, Tweets are the **hardest**
- News, Article, Claim as Context are the **easiest**

