

NERO: A Neural Rule Grounding Framework for Label-Efficient Relation Extraction

Wenxuan Zhou, Hongtao Lin, Bill Yuchen Lin, Ziqi Wang, Junyi Du, Leonardo Neves, and Xiang Ren

University of Southern California, Tsinghua University, Snapchat Inc.



Relation Extraction

Microsoft was founded by **Bill Gates**.

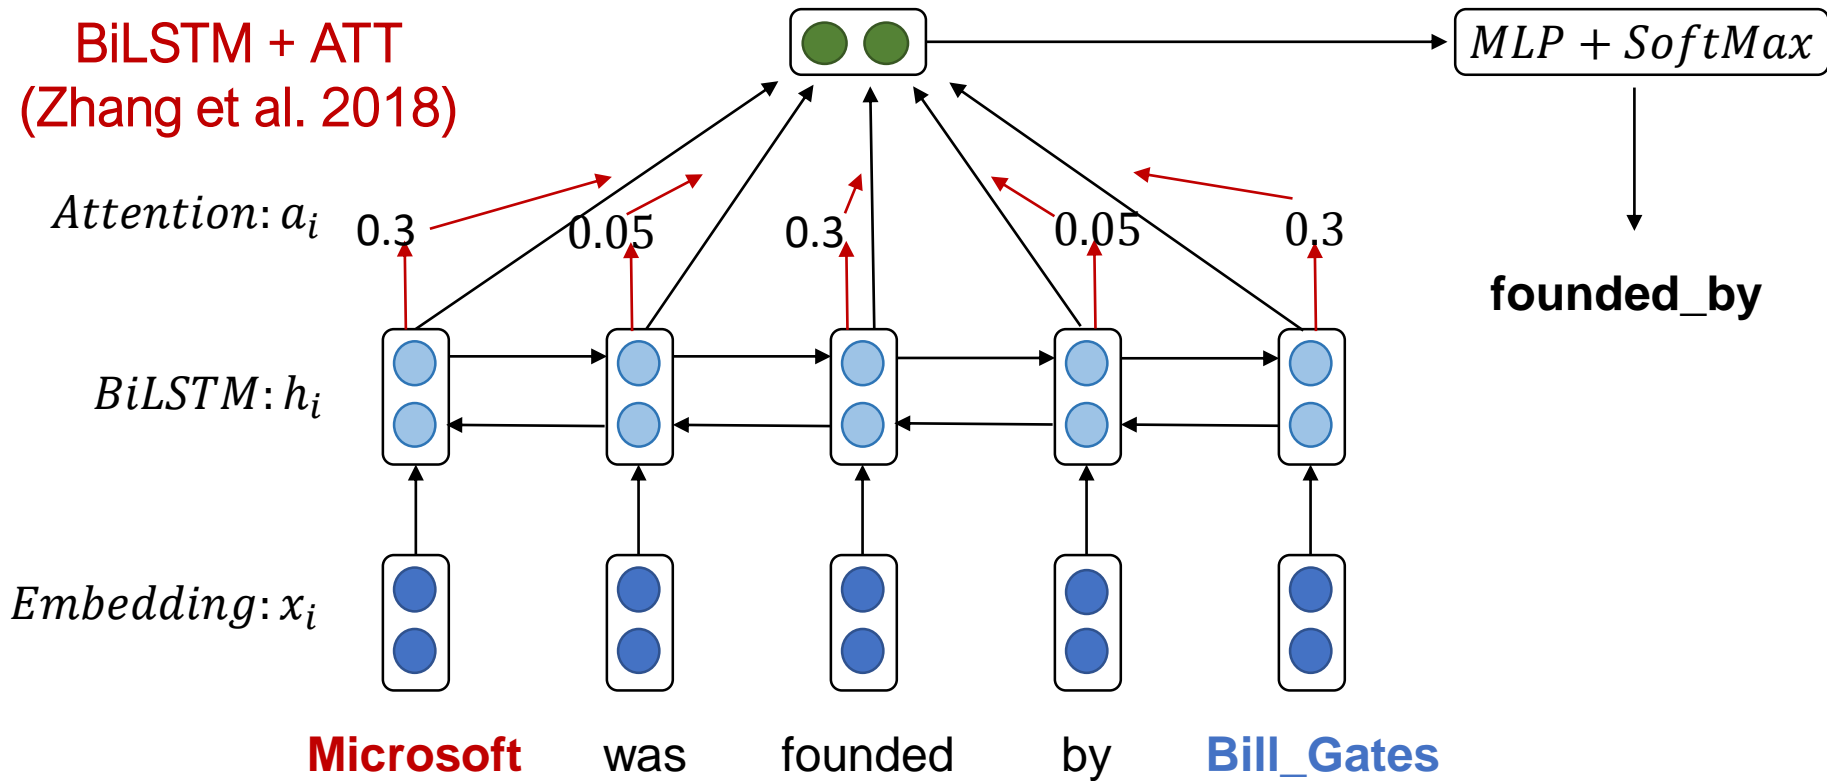
Relation: founded_by

Mike was born March 26 , 1965 , in **US**.

Relation: origin

What is the **semantic relationship** between the given entities?

Neural Model for Relation Extraction

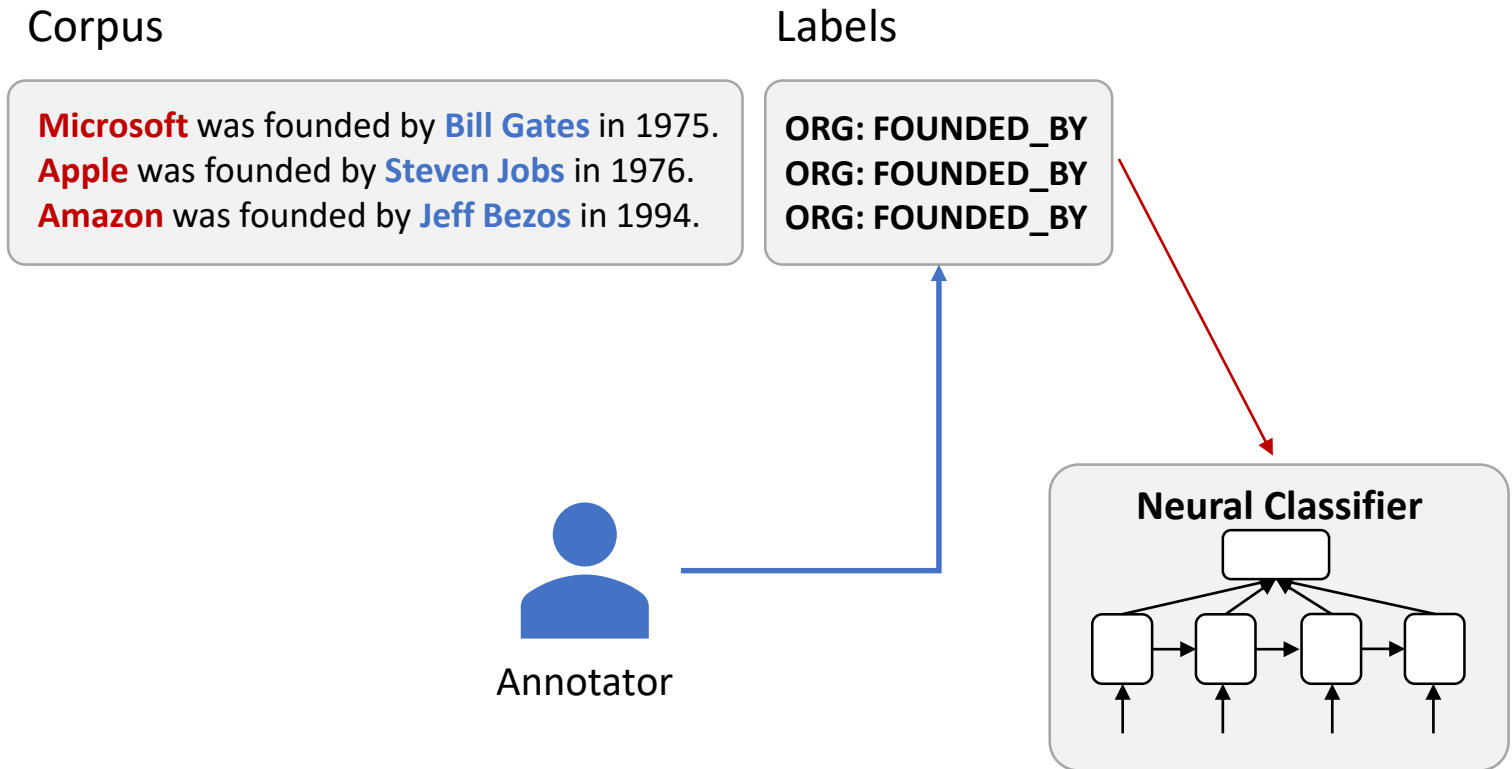


Need a lot of human-annotated labels!

Research question:

How to train a good neural model with less human labor?

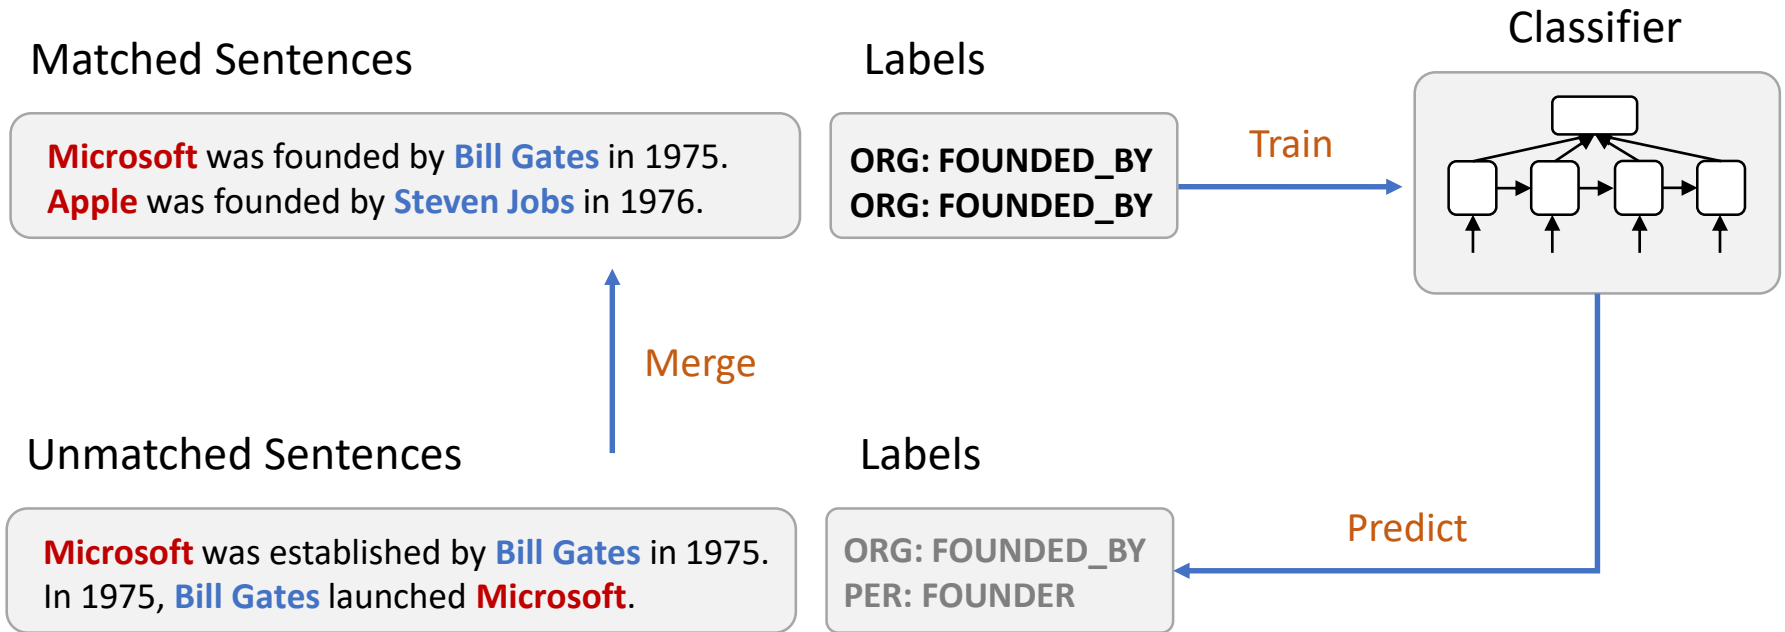
Standard Pipeline for Labeling Data



Standard Data Annotation

Slow, redundant annotation efforts
on similar instances!

Semi-supervised Learning: Self-Training



Can create pseudo labeled data, but will suffer from cascading error aggregation

Alternative Labeling Scheme: Surface Pattern Rules

Corpus

Microsoft was founded by **Bill Gates** in 1975.
Apple was founded by **Steven Jobs** in 1976.
Amazon was founded by **Jeff Bezos** in 1994.

Labels

ORG: FOUNDED_BY
ORG: FOUNDED_BY
ORG: FOUNDED_BY

SUBJ-ORG was founded by **OBJ-PER** → **ORG: FOUNDED_BY**



Annotator

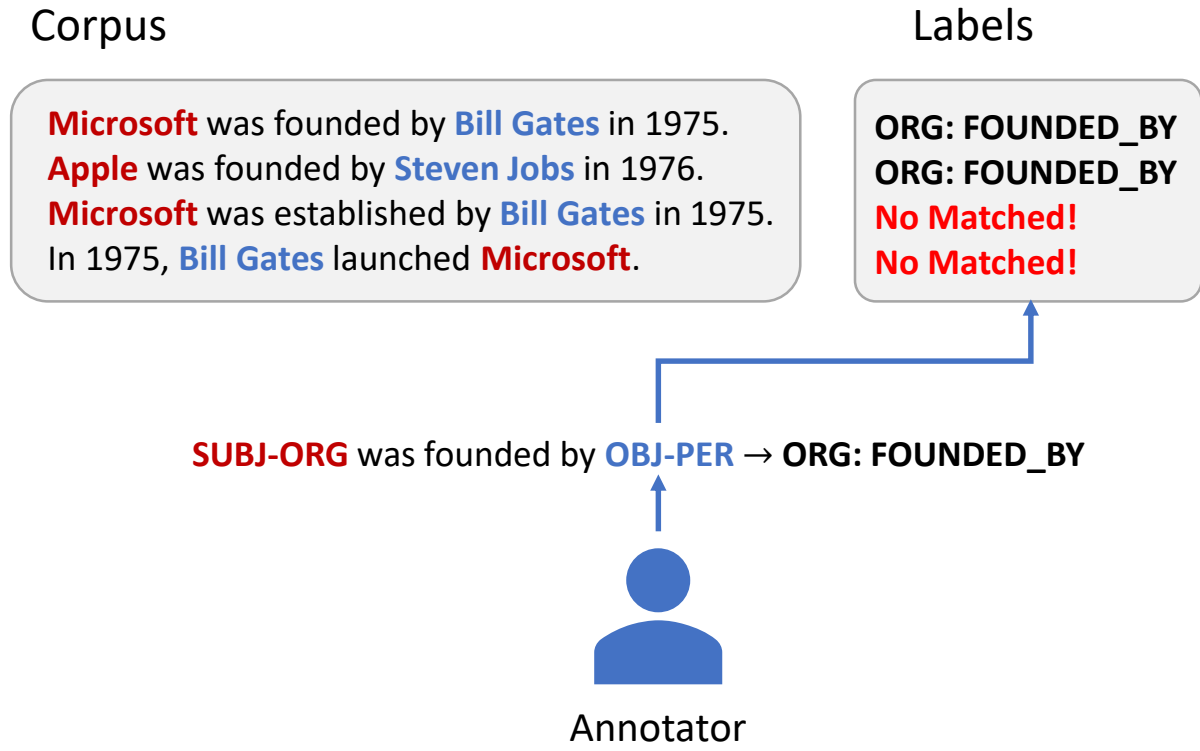
Labeling Rules

Annotate contextually similar instances
via much fewer rules



(Hearst, 1992)

Challenge: Language Variations



A lot of similar sentences cannot be matched ⇒ Not enough training data ⇒ Poor performance

Do we have to add more labeling rules?

Rethinking the Matching Process

Hard-matching

No Matched

Microsoft was established by Bill Gates

SUBJ-ORG was founded by OBJ-PER

Microsoft was established by Bill Gates in 1975.

SUBJ-ORG was founded by OBJ-PER →
ORG: FOUNDED_BY

Microsoft was established by Bill Gates

SUBJ-ORG was founded by OBJ-PER

Neural representation



Matching Score


Soft-matching

Rethinking the Matching Process

P : Set of labeling rules

S : Set of unlabeled sentences

label, score

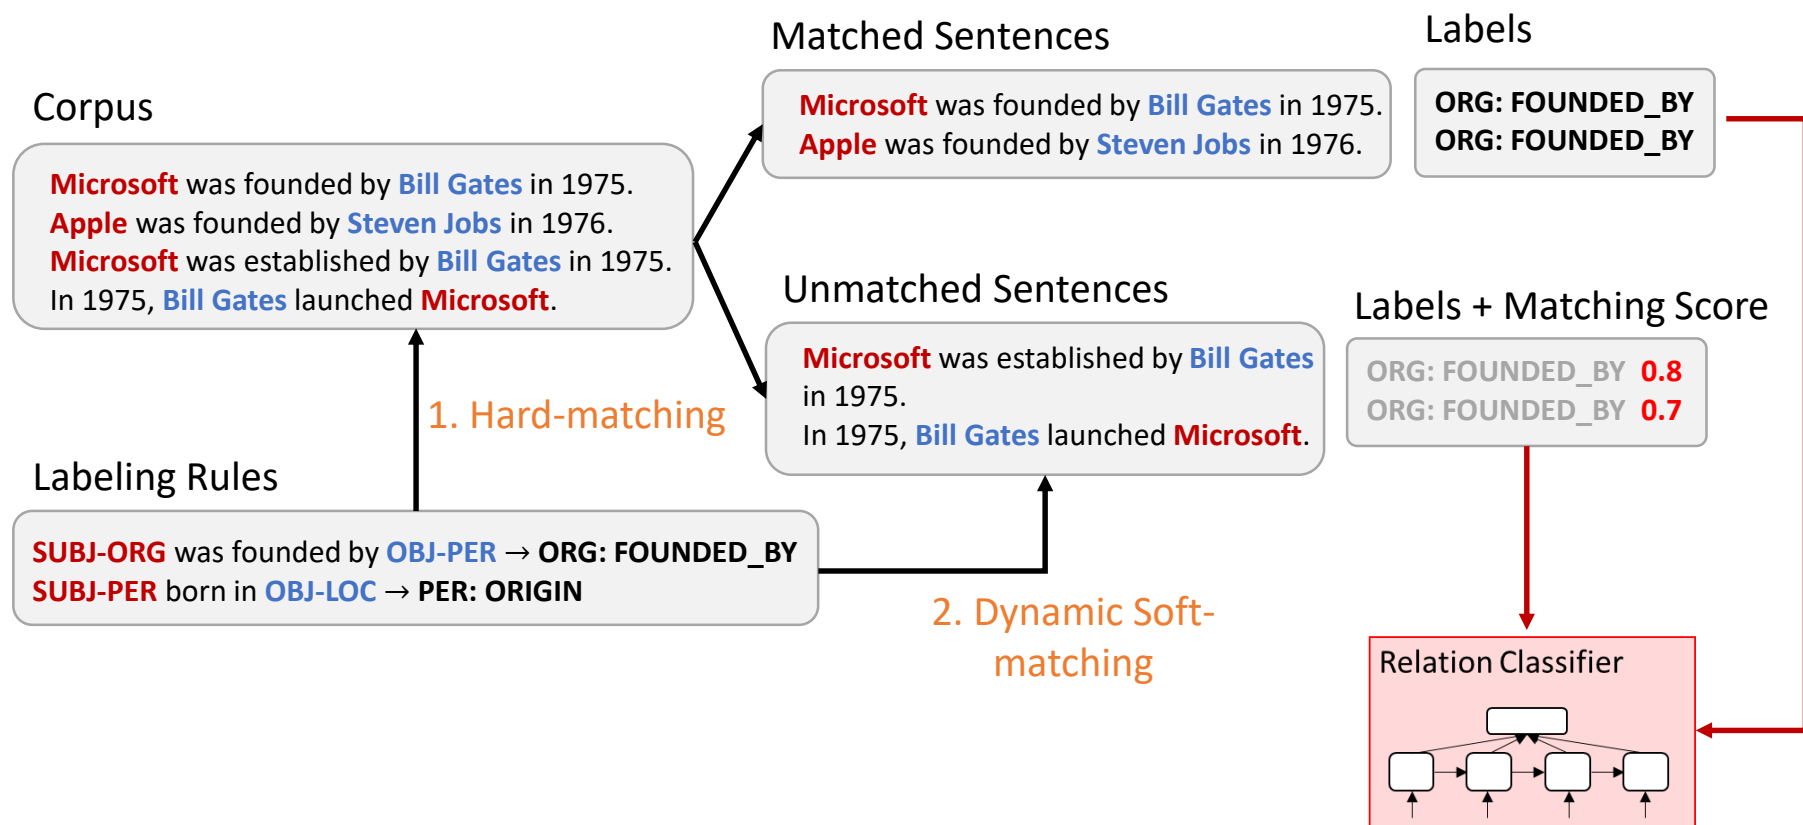


1. soft-matching of rule-sentence pairs

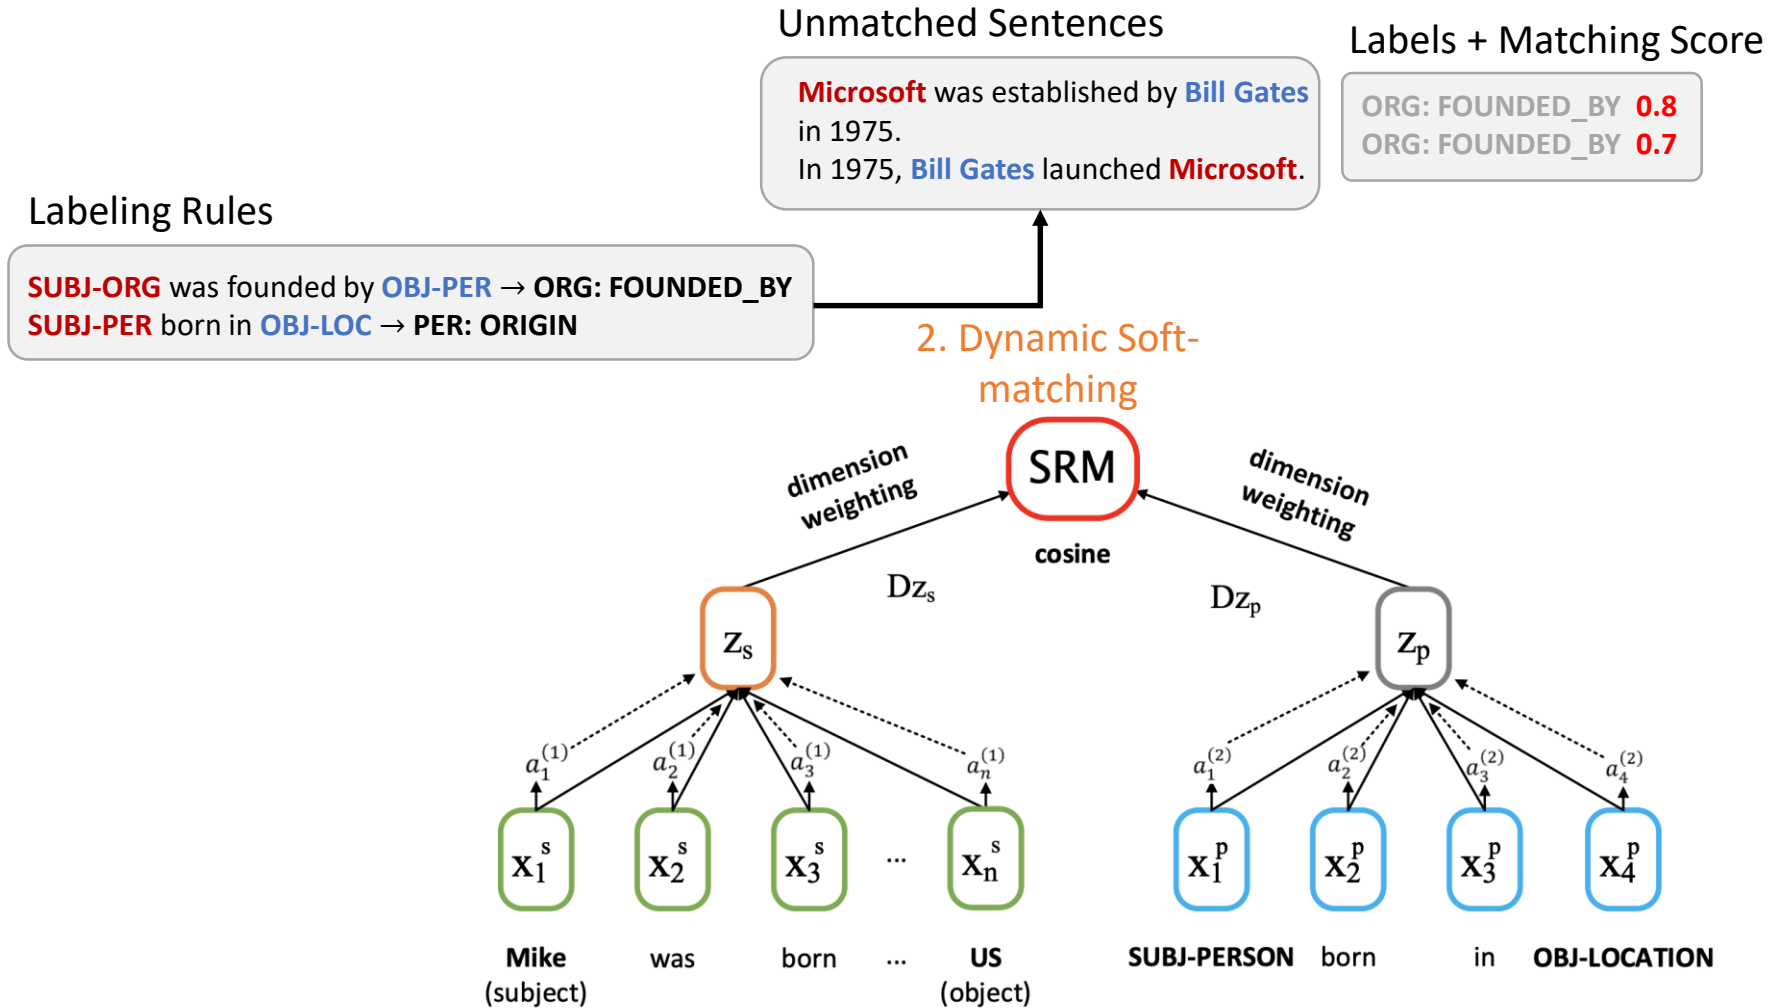
2. weighting labels by matching score

Our Idea: Neural Rule Grounding for Data Augmentation

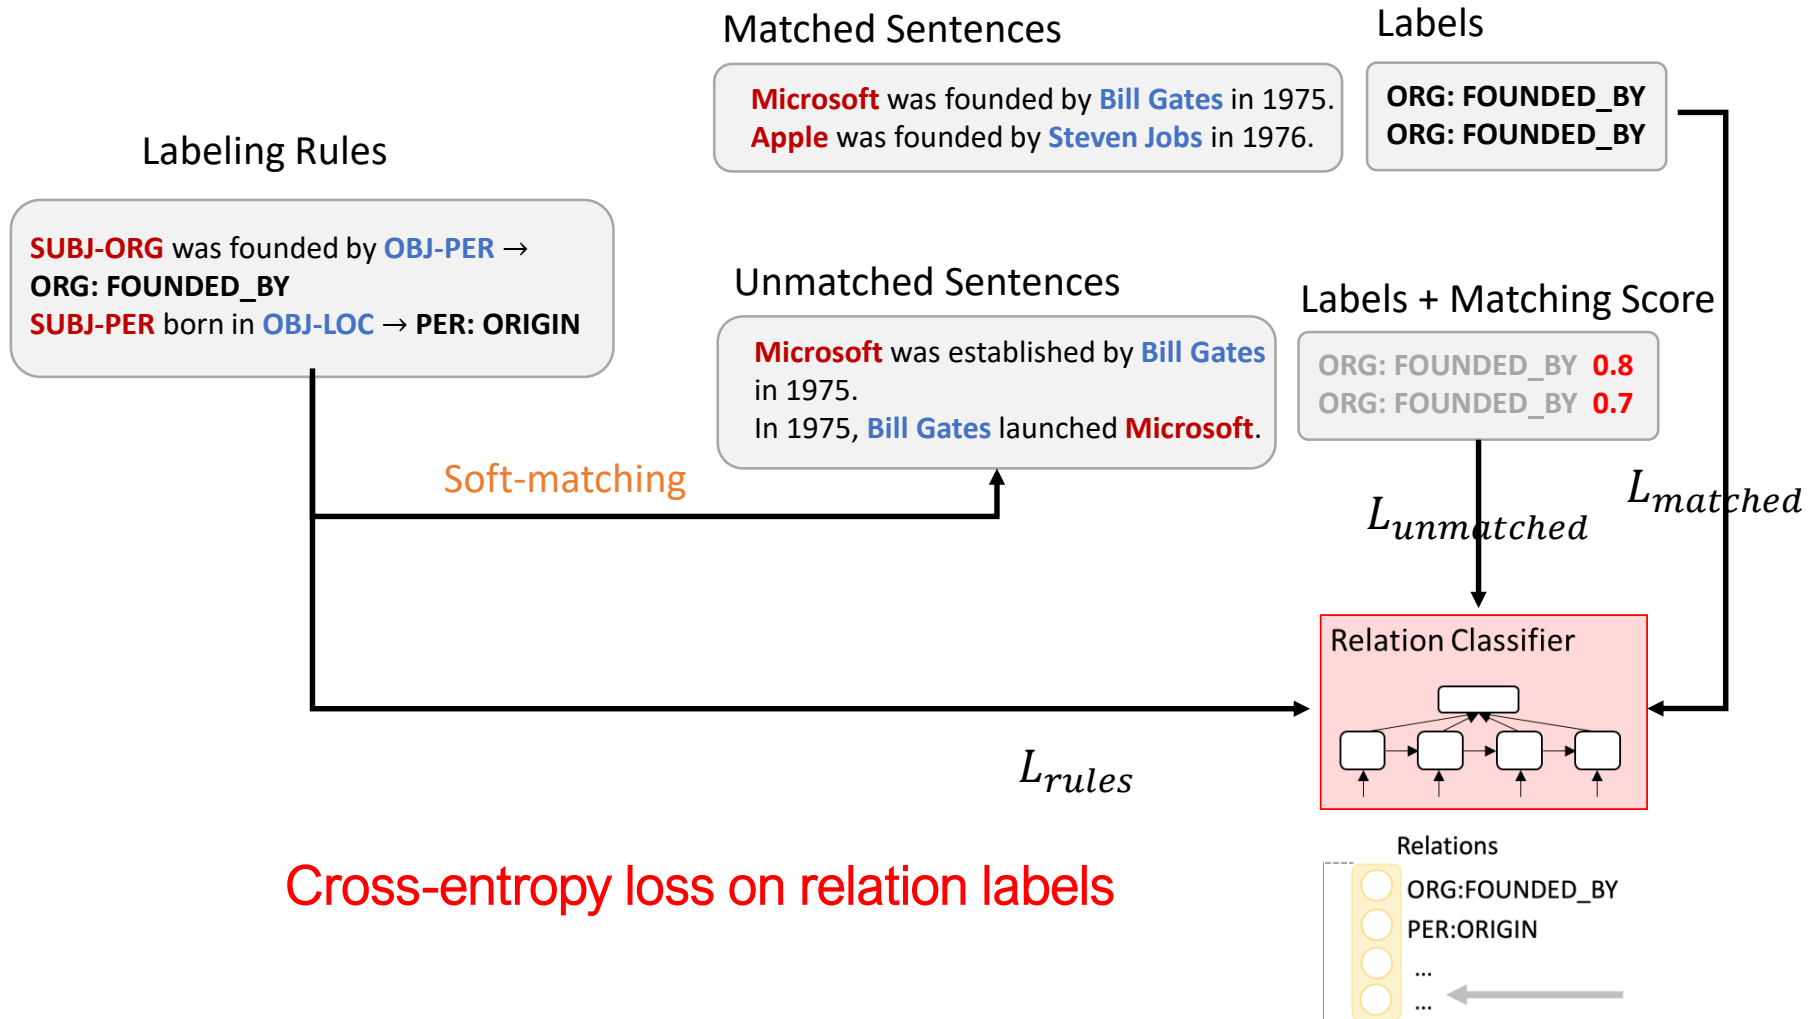
Generalizing rules' coverage via text similarity



Soft Rule Matcher: Architecture



Joint Parameter Learning: Relation Extractor + Soft Rule Matcher




Joint Parameter Learning: Relation Extractor + Soft Rule Matcher

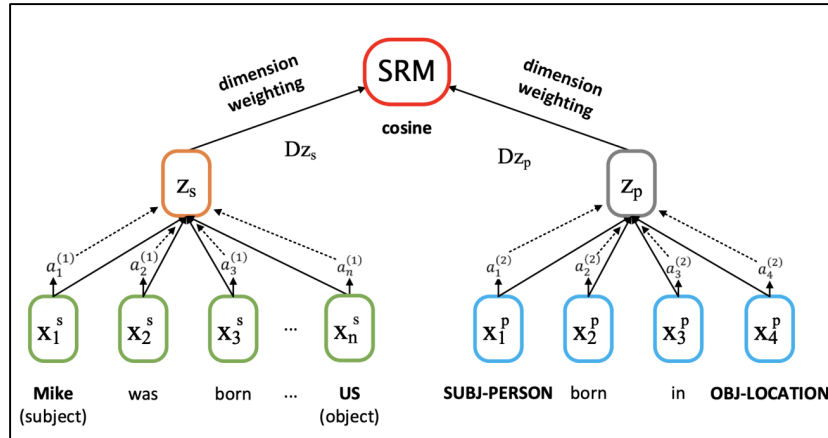
Labeling Rules

SUBJ-ORG was founded by **OBJ-PER** →
ORG: FOUNDED_BY
SUBJ-PER born in **OBJ-LOC** → **PER: ORIGIN**

L_{clus}

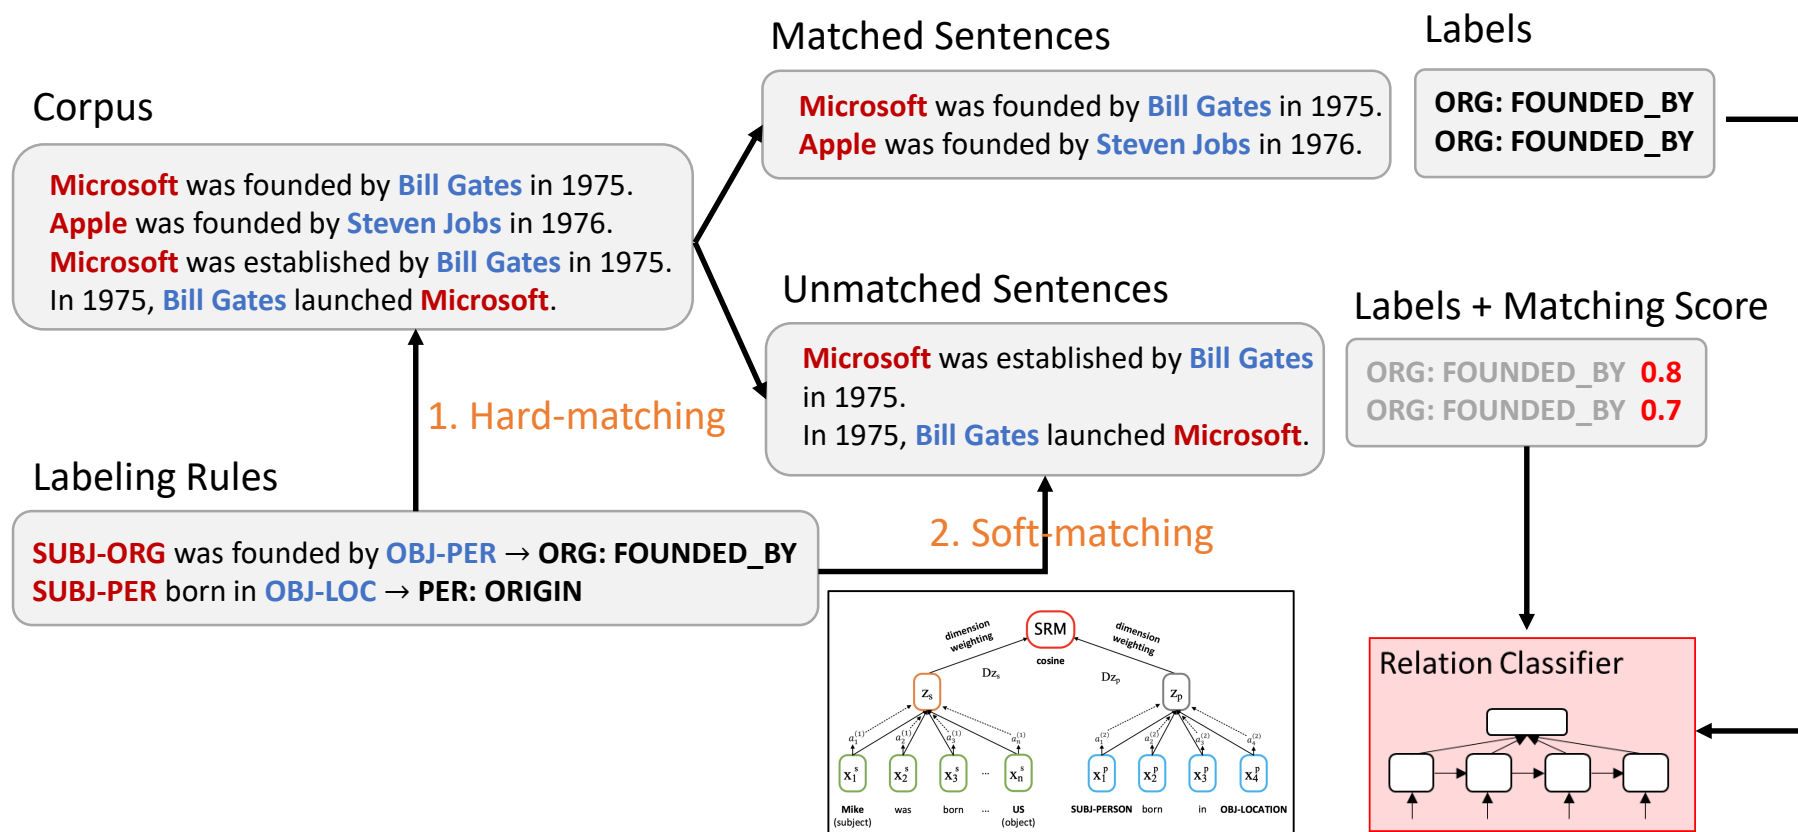


Contrastive loss for discriminating rule bodies



Joint Parameter Learning: Relation Extractor + Soft Rule Matcher

$$L = L_{matched} + \alpha \cdot L_{unmatched} + \beta \cdot L_{rules} + \gamma \cdot L_{clus}$$



Generating Labeling Rules

Corpus

Microsoft was founded by **Bill Gates** in 1975.
Apple was founded by **Steven Jobs** in 1976.
Microsoft was established by **Bill Gates** in 1975.
In 1975, **Bill Gates** launched **Microsoft**.

No need to write labeling rules
by yourself!

1. Automatic
Pattern Mining

Frequent Patterns

SUB-ORG was founded by **OBJ-PER**.

Labeling Rules

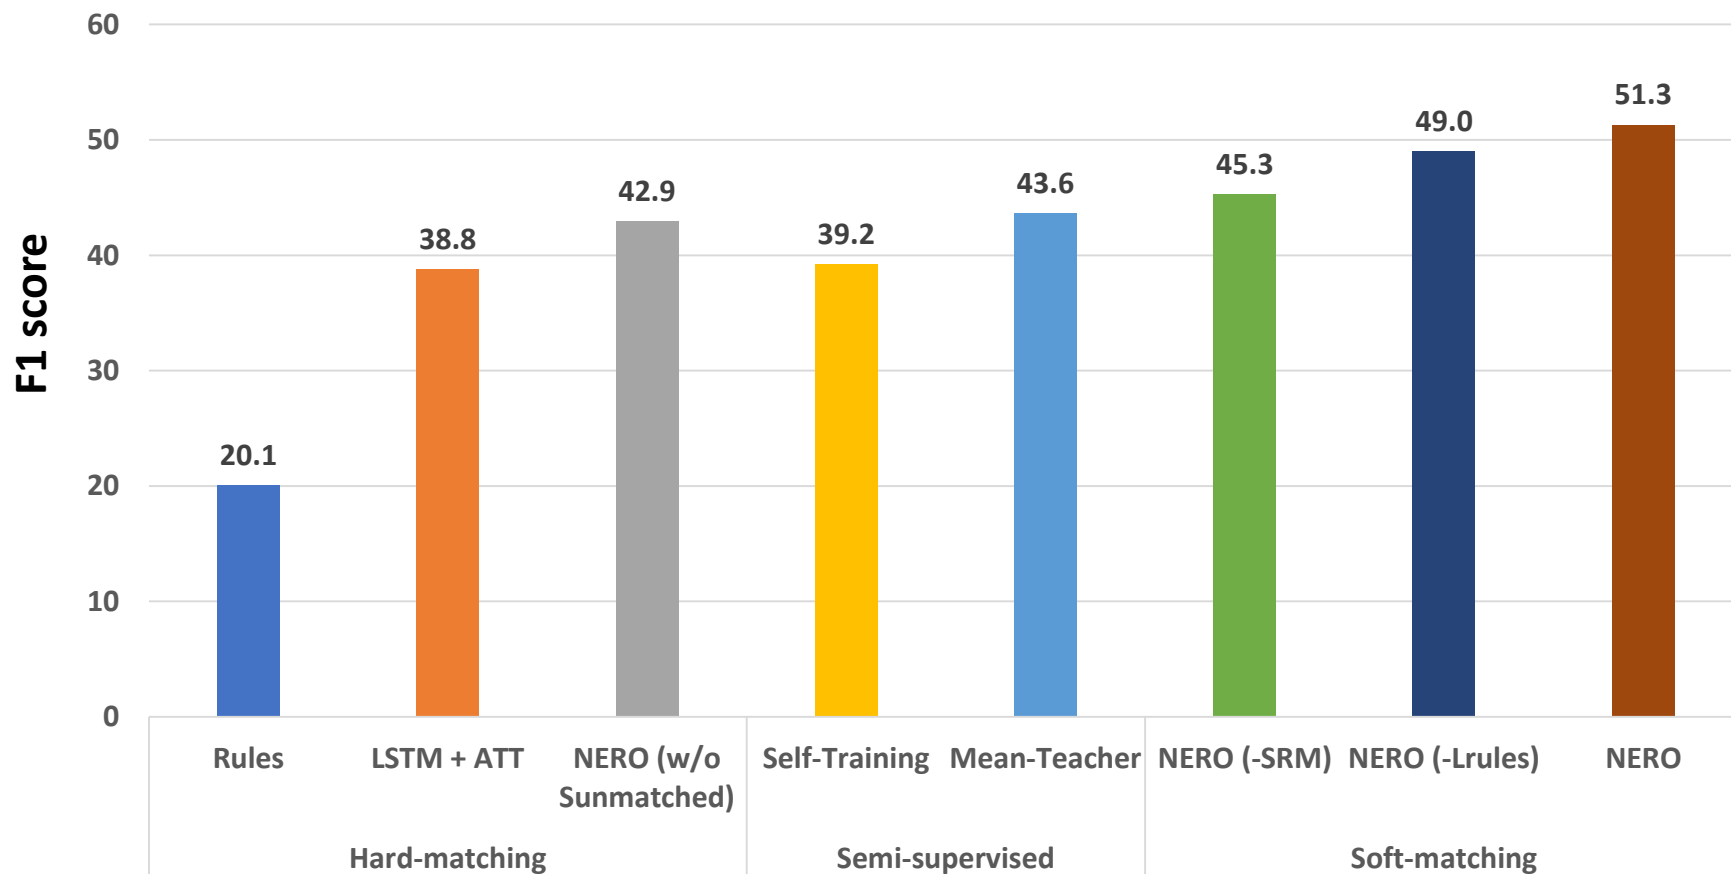
SUBJ-ORG was founded by **OBJ-PER** → **ORG:**
FOUNDED_BY

2. Annotate
Patterns



Annotator


Results on Relation Extraction



Relation Extraction Performance (in F1 score) on TACRED

Study on Label Efficiency

Spent 40min
on labeling
instances from
TACRED



Dashed: Avg # of **rules** / **sentences** labeled by annotators.

Solid: Avg **model F1** trained with corresponding annotations.

{Rules + Neural Rule Grounding} produces much more effective model with limited time!

Conclusion

- Soft-matching mechanism for increased coverage.
- A novel framework for label-efficient relation extraction.
 - Hard-matching + Dynamic Soft-matching
 - Joint Parameter Learning
- Code available at <https://github.com/INK-USC/NERO>