Generalizing Natural Language Analysis through Span-relation Representations

Zhengbao Jiang[1], Wei Xu[2], Jun Araki[3], Graham Neubig[1]

Carnegie Mellon University[1], Ohio State University[2], Bosch Research North America[3]

zhengbaj@cs.cmu.edu

Motivation

- NLP tasks are different
  - Capture different information

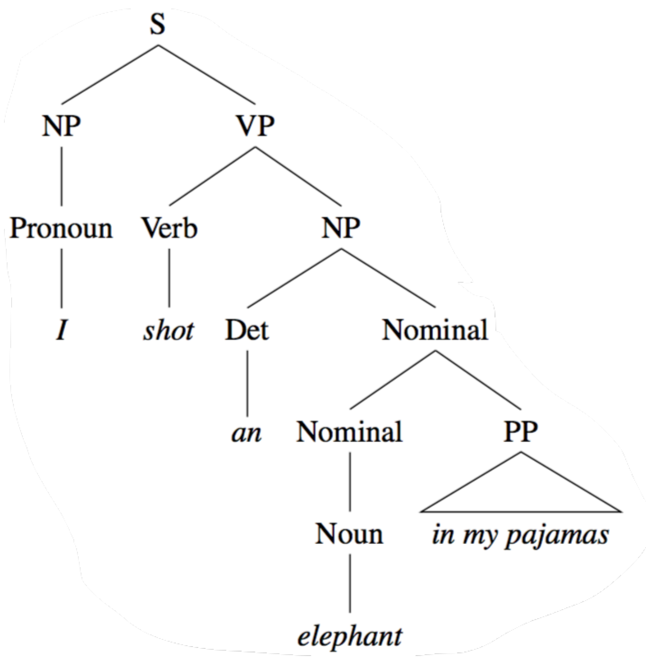

syntax

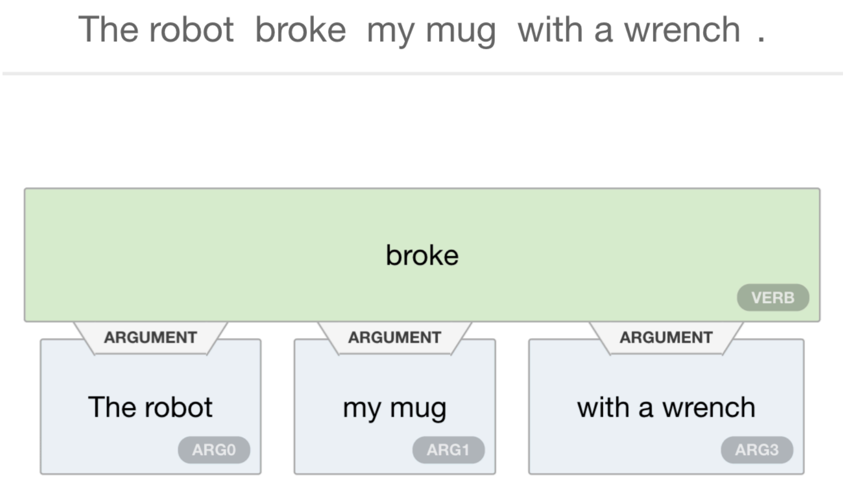

semantics

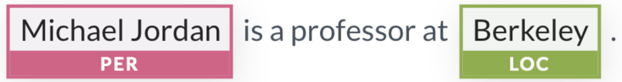

Information content

Motivation

- NLP tasks are different
  - Capture different information
  - Differ in output structure

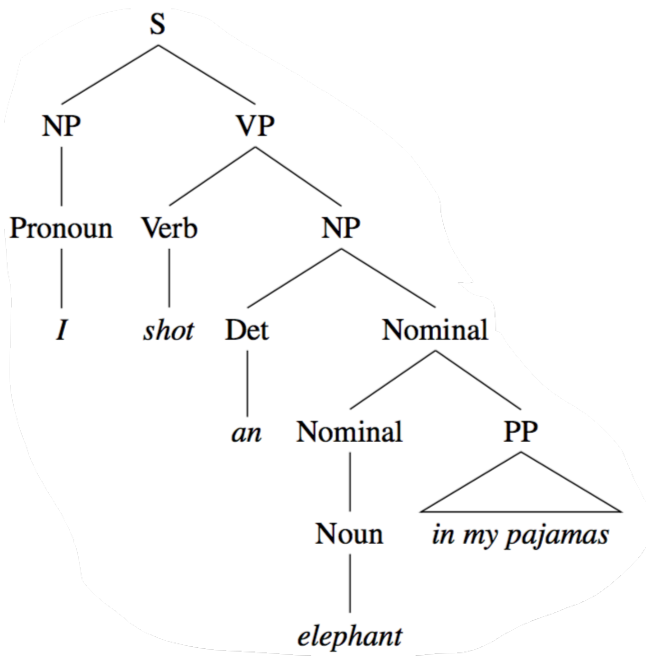

tree

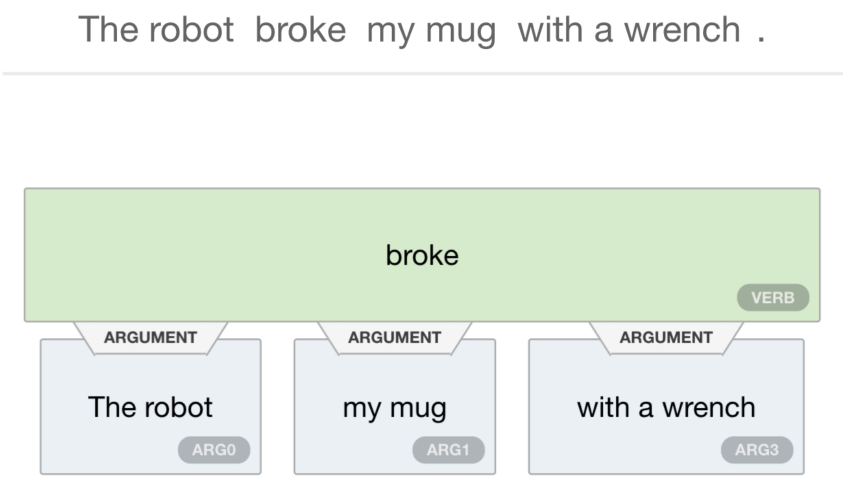

linked spans

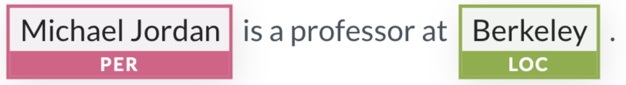

label sequence

Motivation

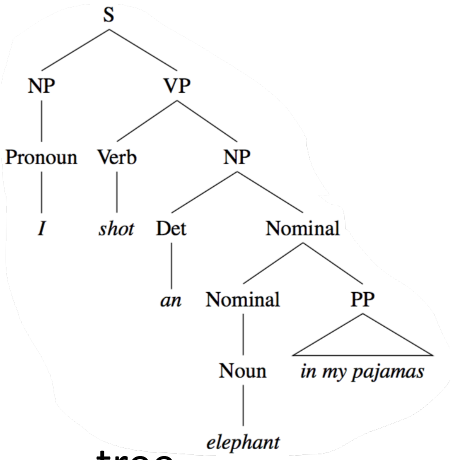

tree

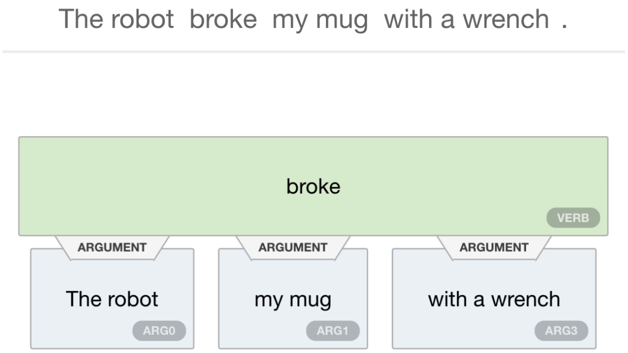

linked spans

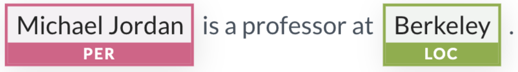

label sequence

Motivation

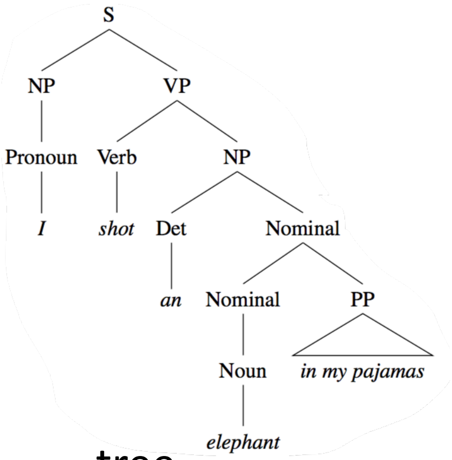

tree

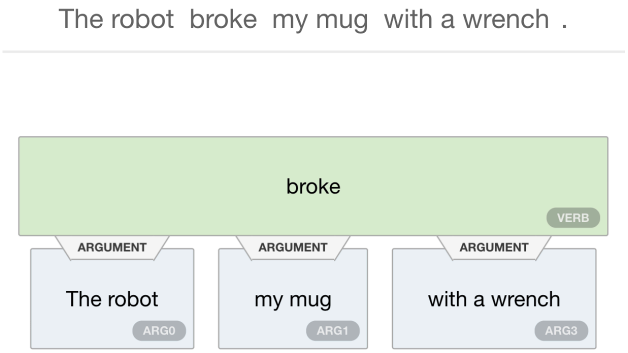

linked spans

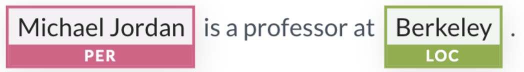

label sequence

unified representational formalism

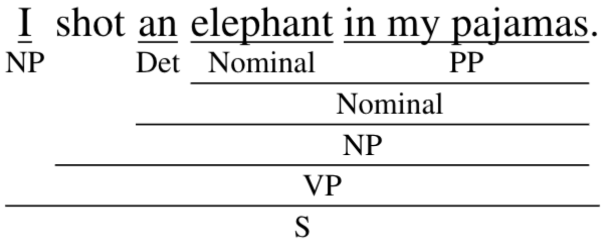

span-relation

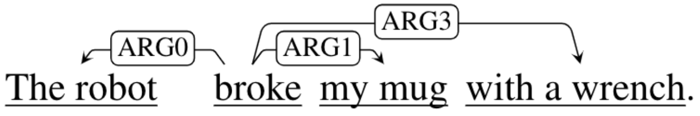

span-relation

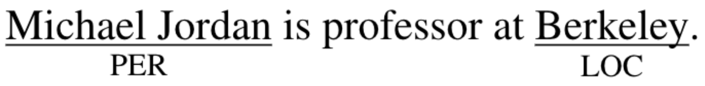

span-relation

Motivation

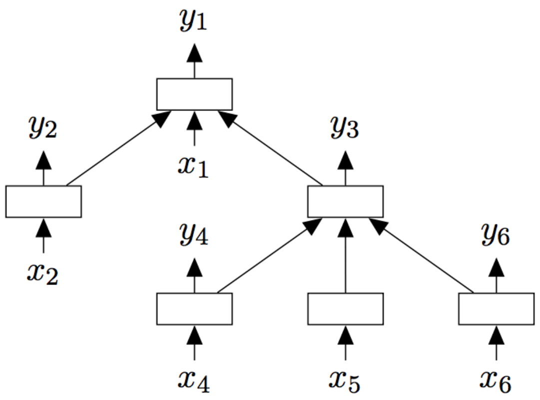

tree-based

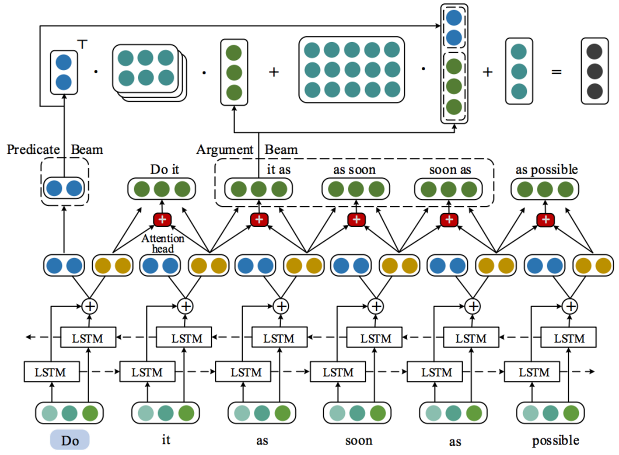

span-based

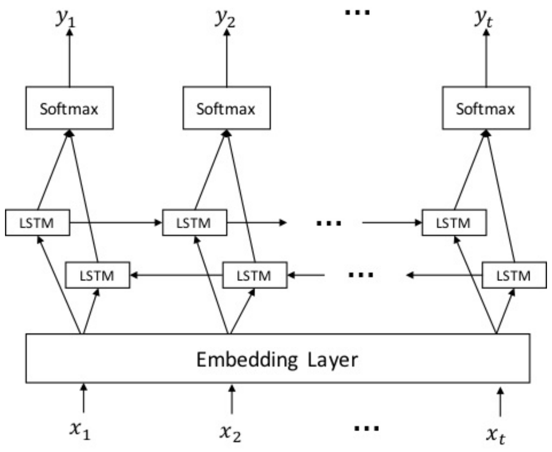

sequence labeling

Motivation

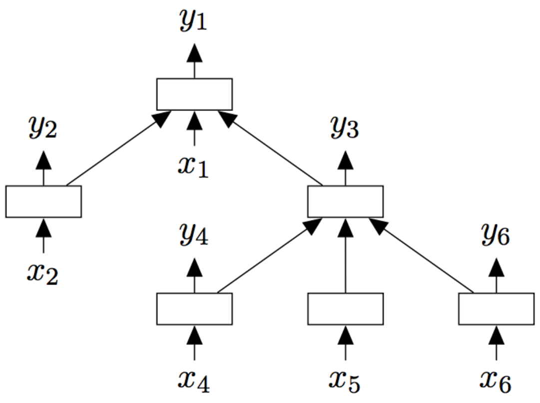

tree-based

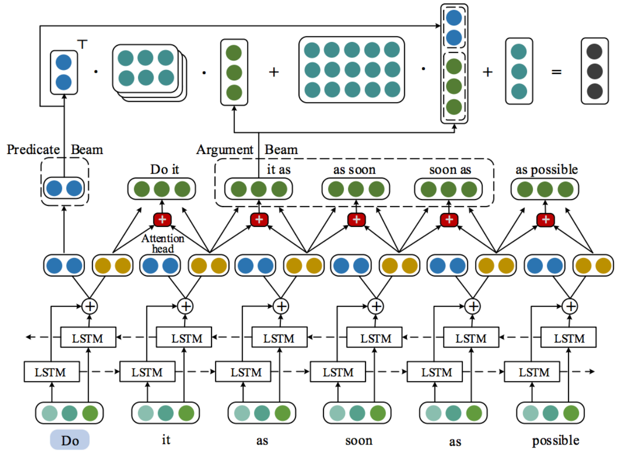

span-based

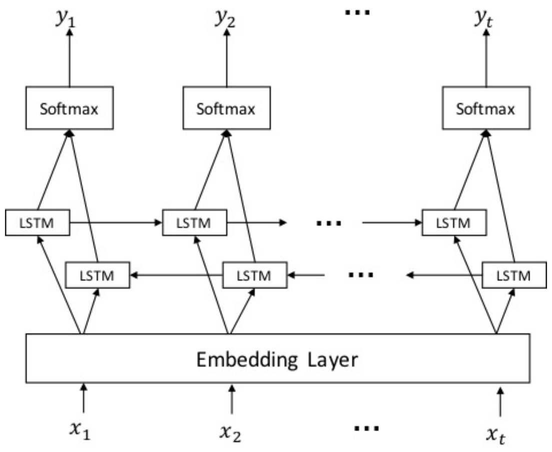

sequence labeling

unified model

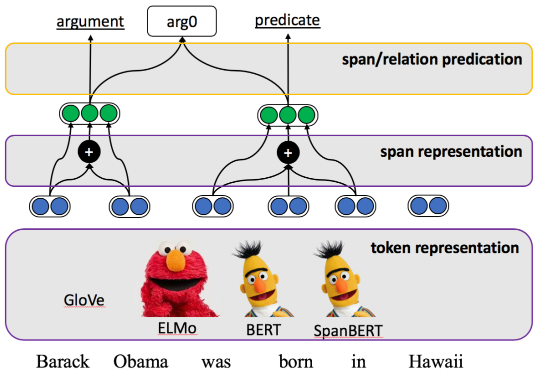

SpanRel

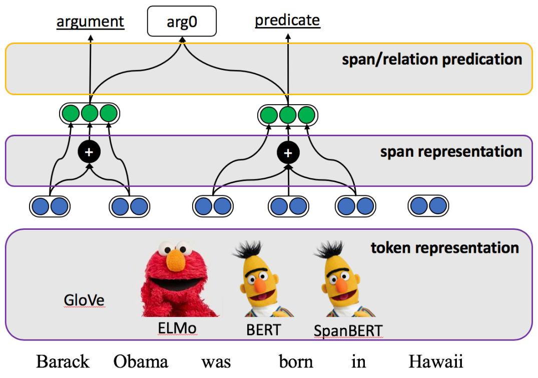

SpanRel

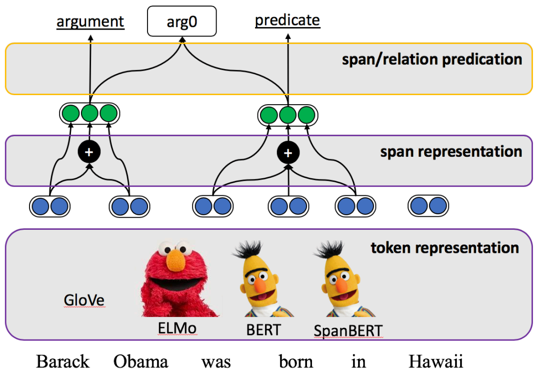

SpanRel

A Unified View: Span-relation Representation

- BRAT annotation interface

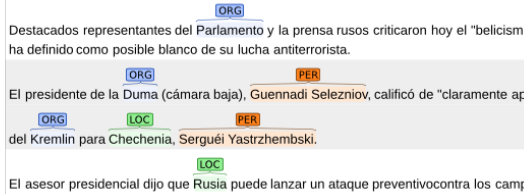

entity

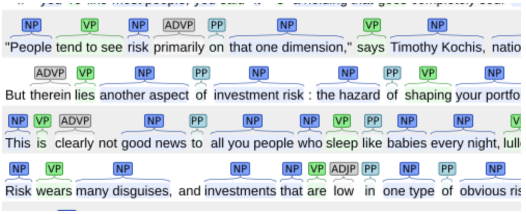

chunking

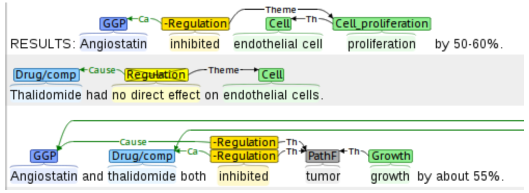

event

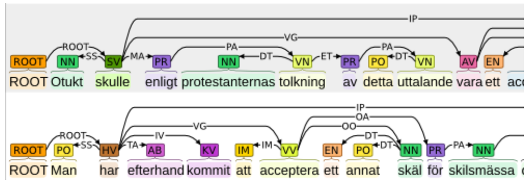

dependency

- Labeled relations between span pairs.

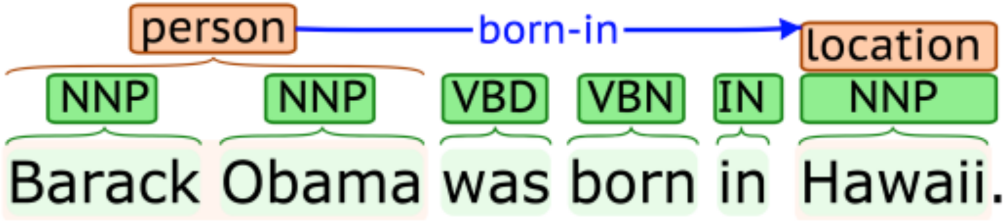

- Spans of one/multiple words with their labels

Span-oriented Tasks

Task	**Spans annotated with labels**
NER	$\underbrace{\text{Barack Obama}}_{\text{person}}$ was born in $\underbrace{\text{Hawaii}}_{\text{location}}$.
Consti.	$\underbrace{\text{And } \underbrace{\underbrace{\text{their suspicions}}_{\text{NP}}\ \underbrace{\text{of } \underbrace{\text{each other}}_{\text{NP}}}_{\text{PP}}}_{\text{NP}}\ \underbrace{\text{run } \underbrace{\text{deep}}_{\text{ADVP}}}_{\text{VP}}\text{.}}_{\text{S}}$
POS	$\underbrace{\text{What}}_{\text{WP}}\ \underbrace{\text{kind}}_{\text{NN}}\ \underbrace{\text{of}}_{\text{IN}}\ \underbrace{\text{memory}}_{\text{NN}}$?
ABSA	Great laptop that offers many great $\underbrace{\text{features}}_{\text{positive}}$!

Span-oriented Tasks

Task	Spans annotated with labels
NER	Barack Obama was born in Hawaii. person location
Consti.	And their suspicions of each other run deep. NP NP ADVP PP VP NP S
POS	What kind of memory? WP NN IN NN
ABSA	Great laptop that offers many great features! positive

1. Named entity recognition
Spans are named entities

Span-oriented Tasks

Task	Spans annotated with labels
NER	Barack Obama was born in Hawaii. person location
Consti.	And their suspicions of each other run deep. NP NP ADVP PP VP NP S
POS	What kind of memory? WP NN IN NN
ABSA	Great laptop that offers many great features! positive

1. Named entity recognition
Spans are named entities

2. Constituency parsing
Spans are (nested) constituents

Span-oriented Tasks

Task	Spans annotated with labels
NER	Barack Obama was born in Hawaii. Barack Obama: person; Hawaii: location
Consti.	And their suspicions of each other run deep. their suspicions: NP; each other: NP; deep: ADVP; of each other: PP; run deep: VP; their suspicions of each other: NP; And their suspicions of each other run deep.: S
POS	What kind of memory? What: WP; kind: NN; of: IN; memory: NN
ABSA	Great laptop that offers many great features! features: positive

1. Named entity recognition
Spans are named entities

2. Constituency parsing
Spans are (nested) constituents

3. Part-of-speech tagging
Spans are single-token words

Span-oriented Tasks

Task	Spans annotated with labels
NER	Barack Obama was born in Hawaii. Barack Obama: person; Hawaii: location
Consti.	And their suspicions of each other run deep. their suspicions: NP; each other: NP; deep: ADVP; of each other: PP; run deep: VP; their suspicions of each other: NP; And their suspicions of each other run deep.: S
POS	What kind of memory? What: WP; kind: NN; of: IN; memory: NN
ABSA	Great laptop that offers many great features! features: positive

1. Named entity recognition
Spans are named entities

2. Constituency parsing
Spans are (nested) constituents

3. Part-of-speech tagging
Spans are single-token words

4. Aspect-based sentiment analysis
Spans are aspects

Relation-oriented Tasks

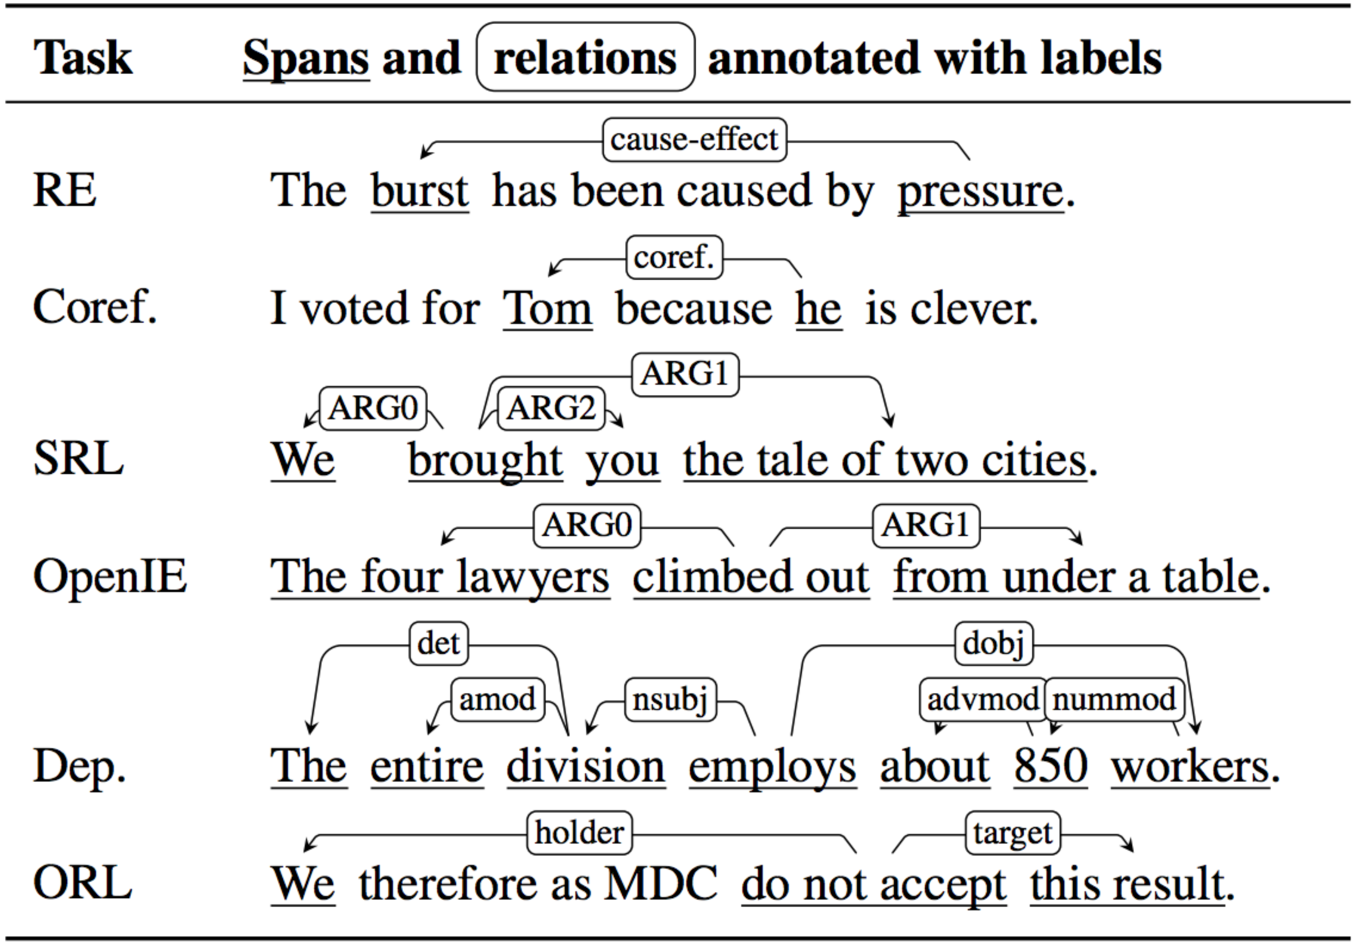

Relation-oriented Tasks

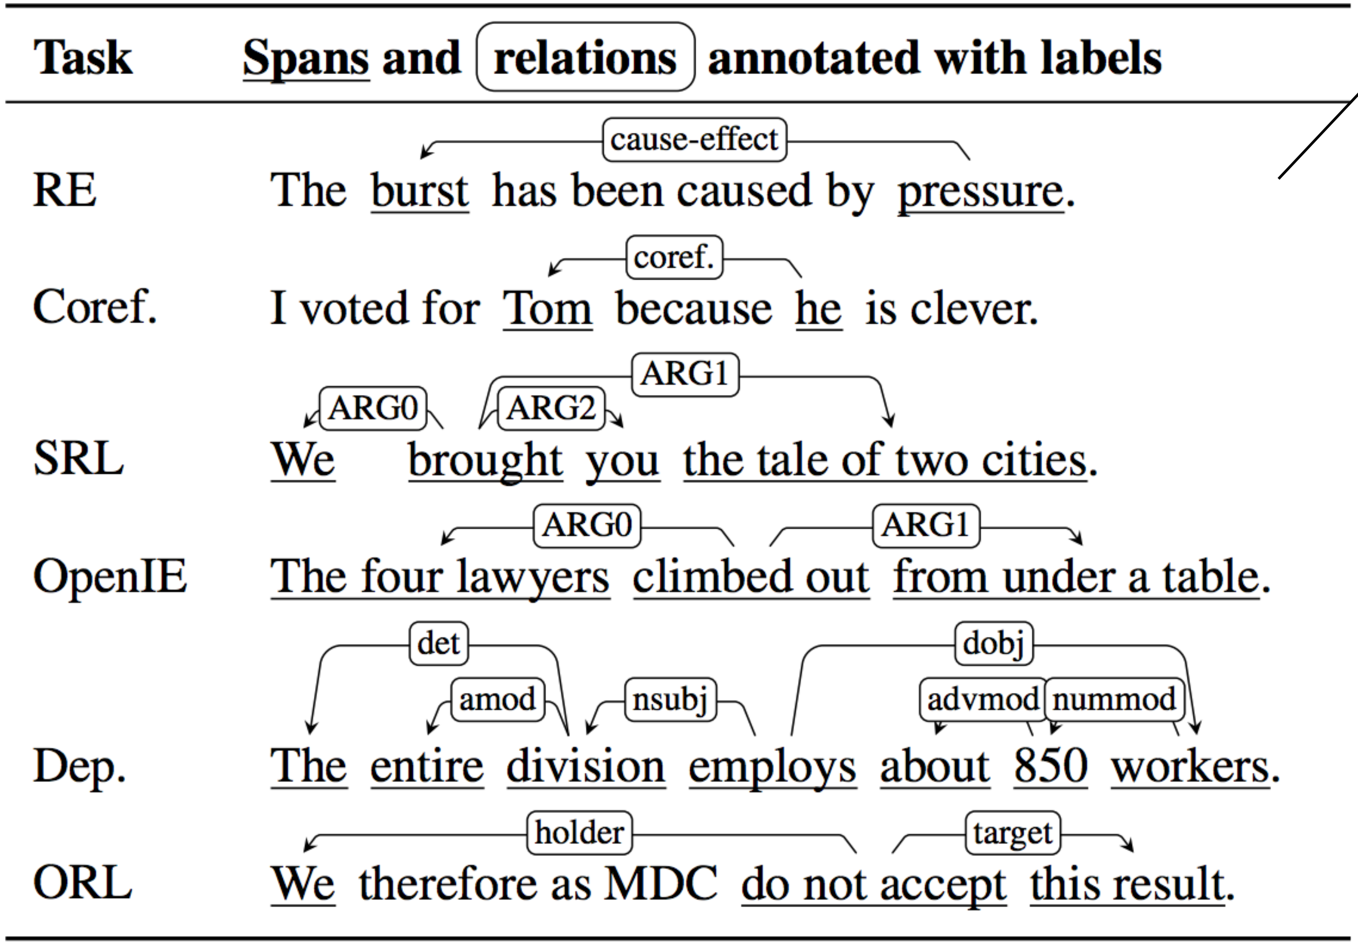

5. Relation extraction

Spans are entities.

Relations are their relationships.

Relation-oriented Tasks

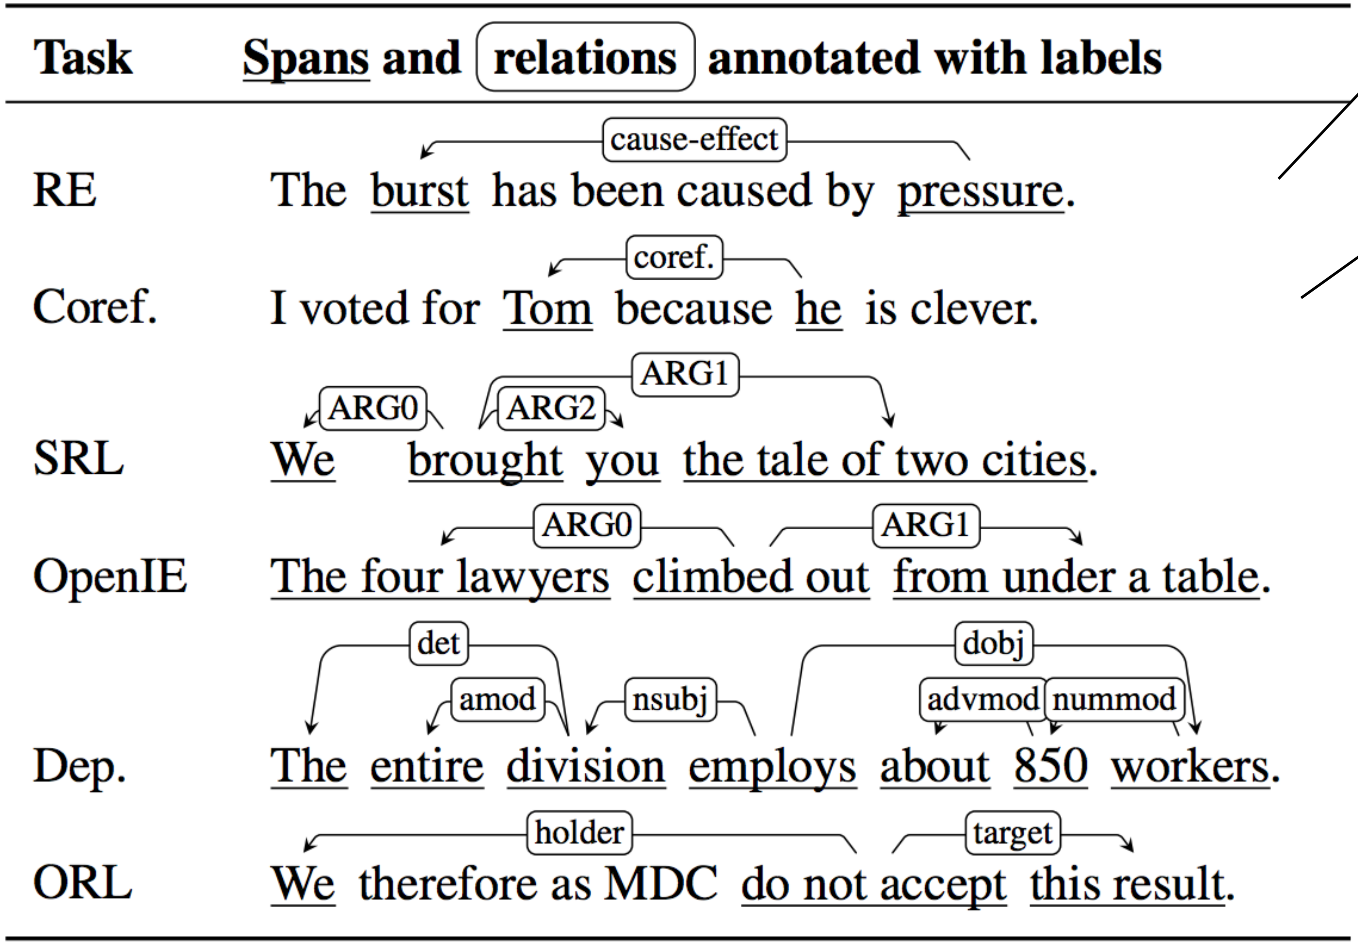

5. Relation extraction

Spans are entities.

Relations are their relationships.

6. Coreference resolution

Spans are mentions

Relations are references

Relation-oriented Tasks

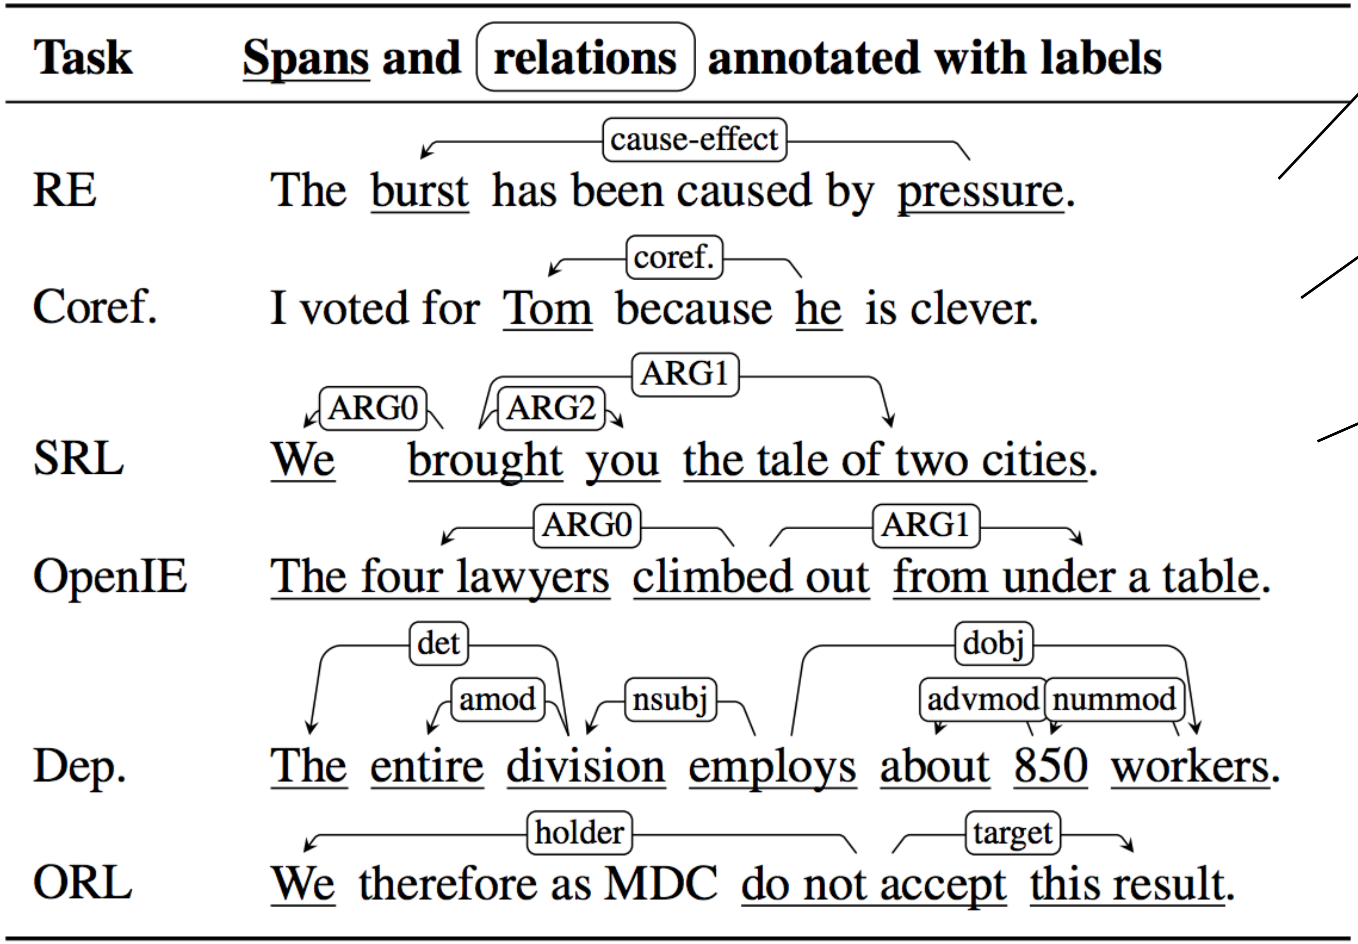

5. Relation extraction
Spans are entities.
Relations are their relationships.

6. Coreference resolution
Spans are mentions
Relations are references

7. Semantic role labeling
Spans are predicates/arguments
Relations link predicates with arguments

Relation-oriented Tasks

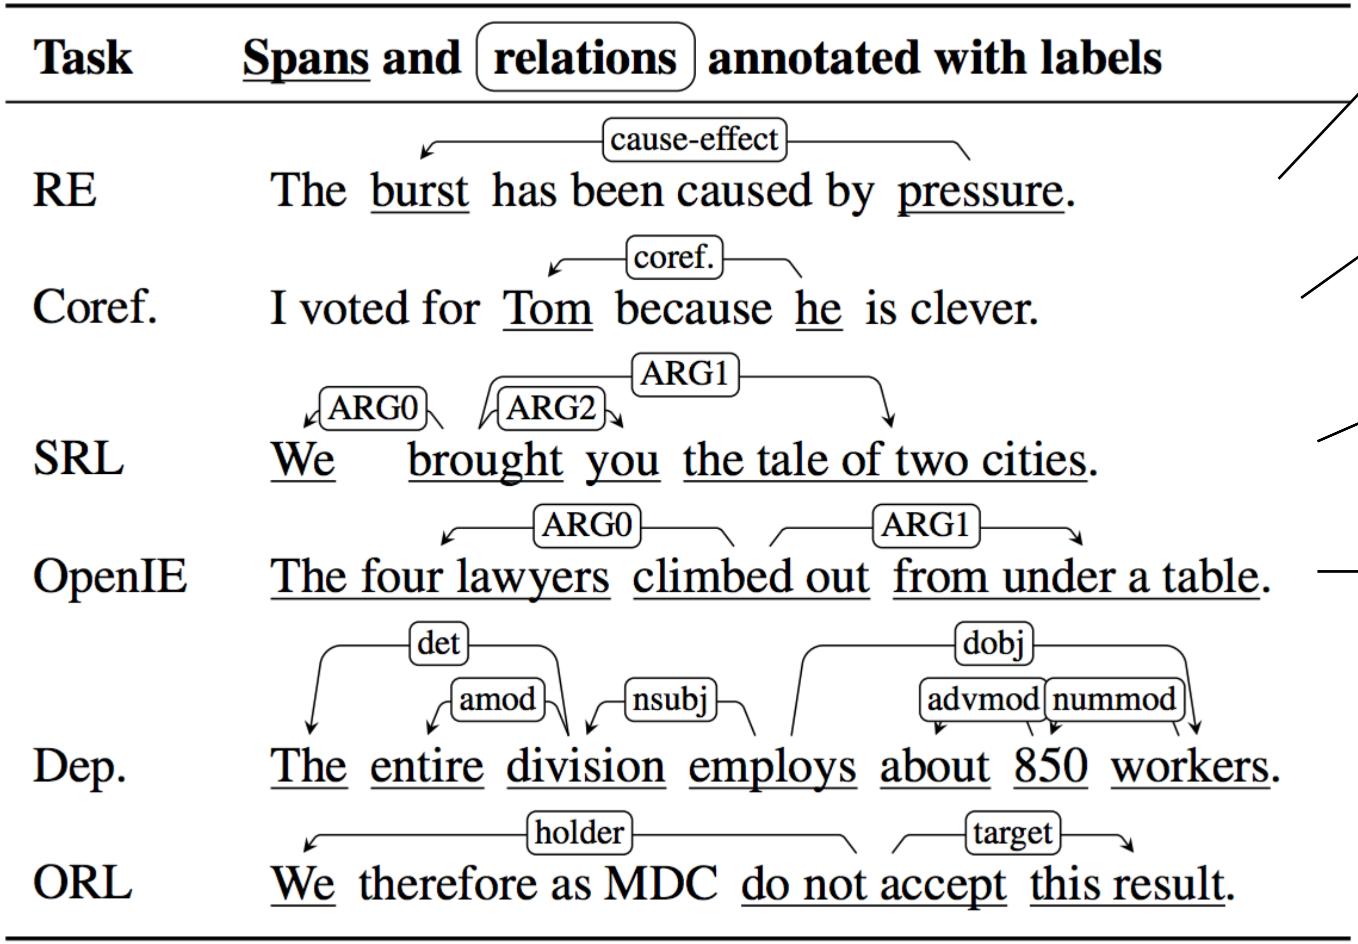

5. Relation extraction
Spans are entities.
Relations are their relationships.

6. Coreference resolution
Spans are mentions
Relations are references

7. Semantic role labeling
Spans are predicates/arguments
Relations link predicates with arguments

8. Open information extraction
Spans are predicates/arguments
Relations link predicates with arguments

Relation-oriented Tasks

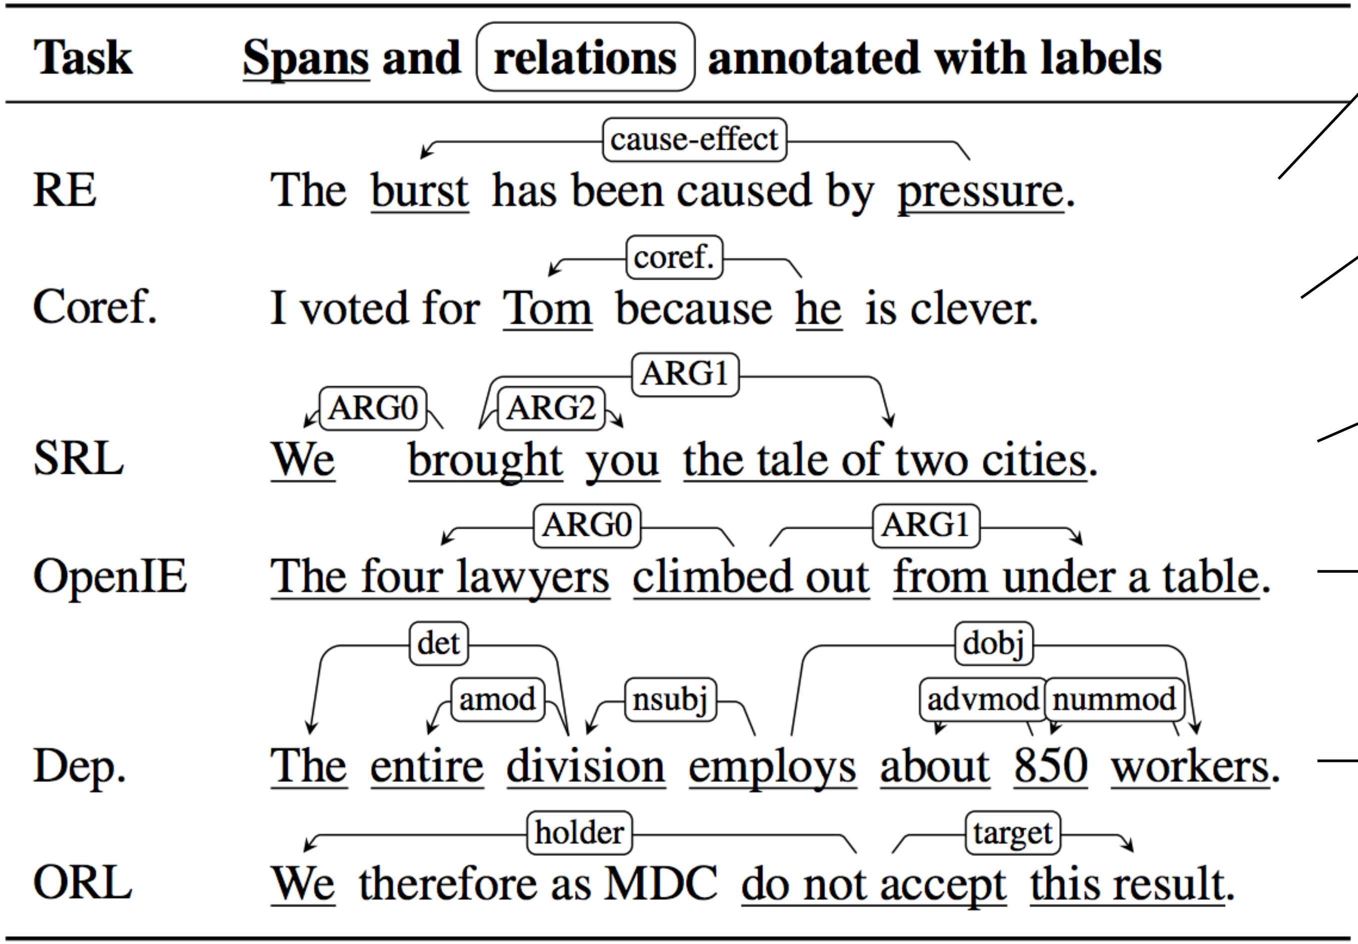

5. Relation extraction
Spans are entities.
Relations are their relationships.

6. Coreference resolution
Spans are mentions
Relations are references

7. Semantic role labeling
Spans are predicates/arguments
Relations link predicates with arguments

8. Open information extraction
Spans are predicates/arguments
Relations link predicates with arguments

9. Dependency parsing
Spans are words
Relations are their dependencies

Relation-oriented Tasks

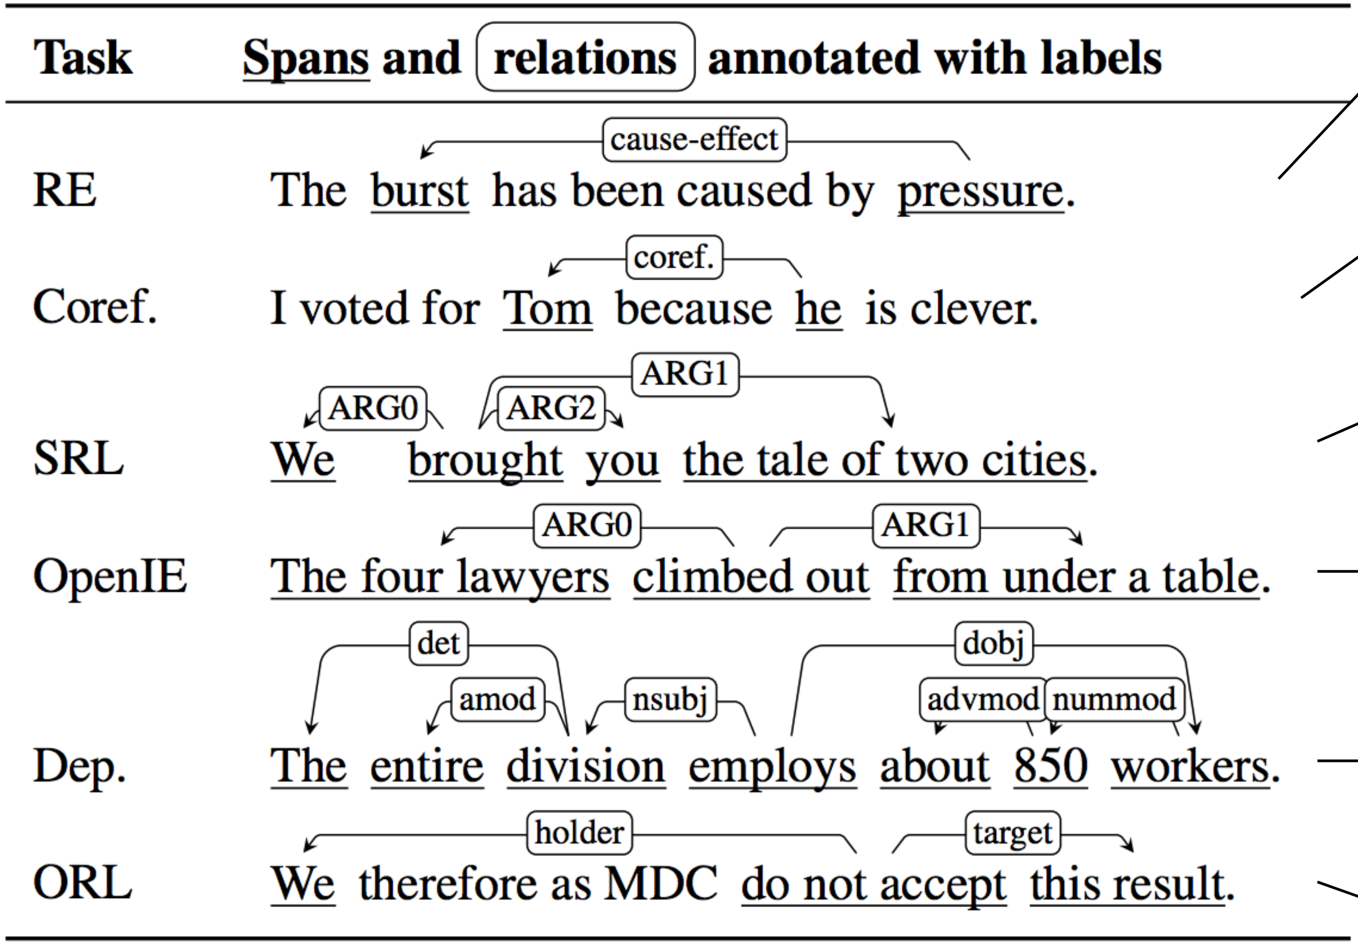

5. Relation extraction
Spans are entities.
Relations are their relationships.

6. Coreference resolution
Spans are mentions
Relations are references

7. Semantic role labeling
Spans are predicates/arguments
Relations link predicates with arguments

8. Open information extraction
Spans are predicates/arguments
Relations link predicates with arguments

9. Dependency parsing
Spans are words
Relations are their dependencies

10. Opinion role labeling
Spans are opinions/holders/targets
Relations link opinions to holders/targets

SpanRel Model

Barack Obama was born in Hawaii

SpanRel Model

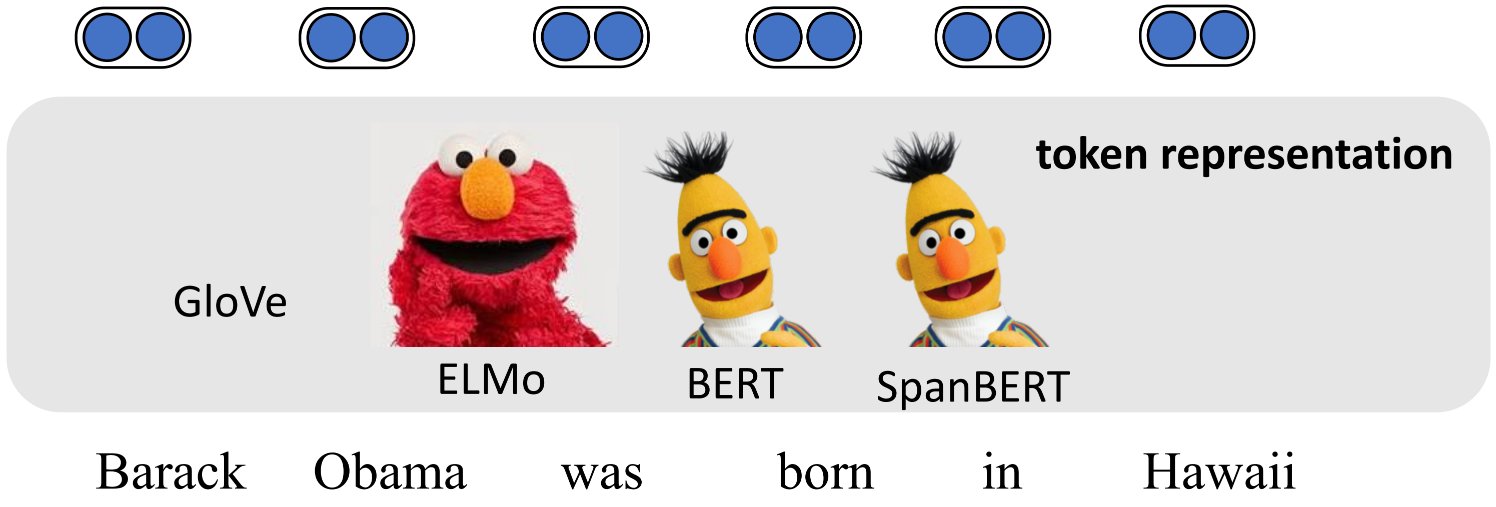

SpanRel Model

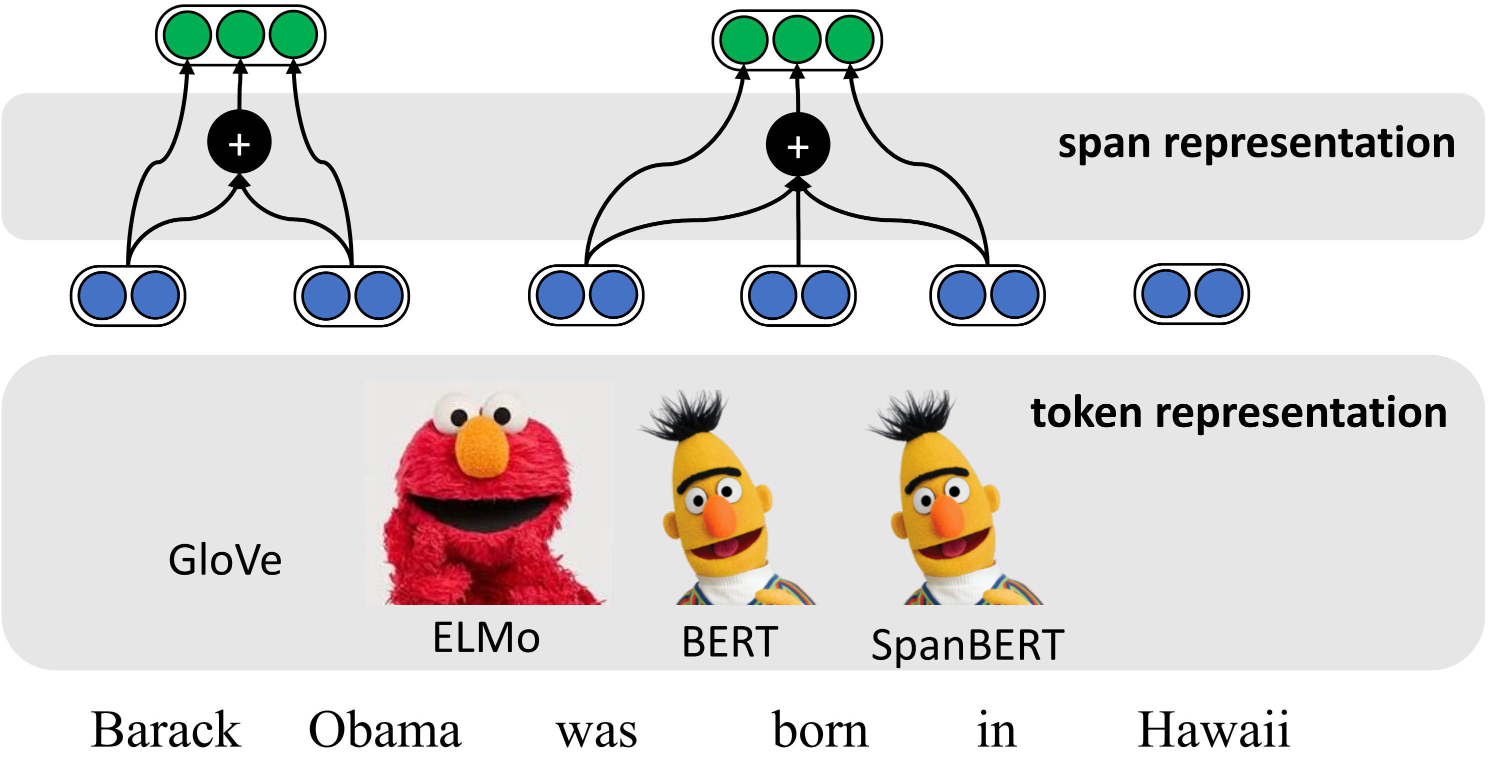

SpanRel Model

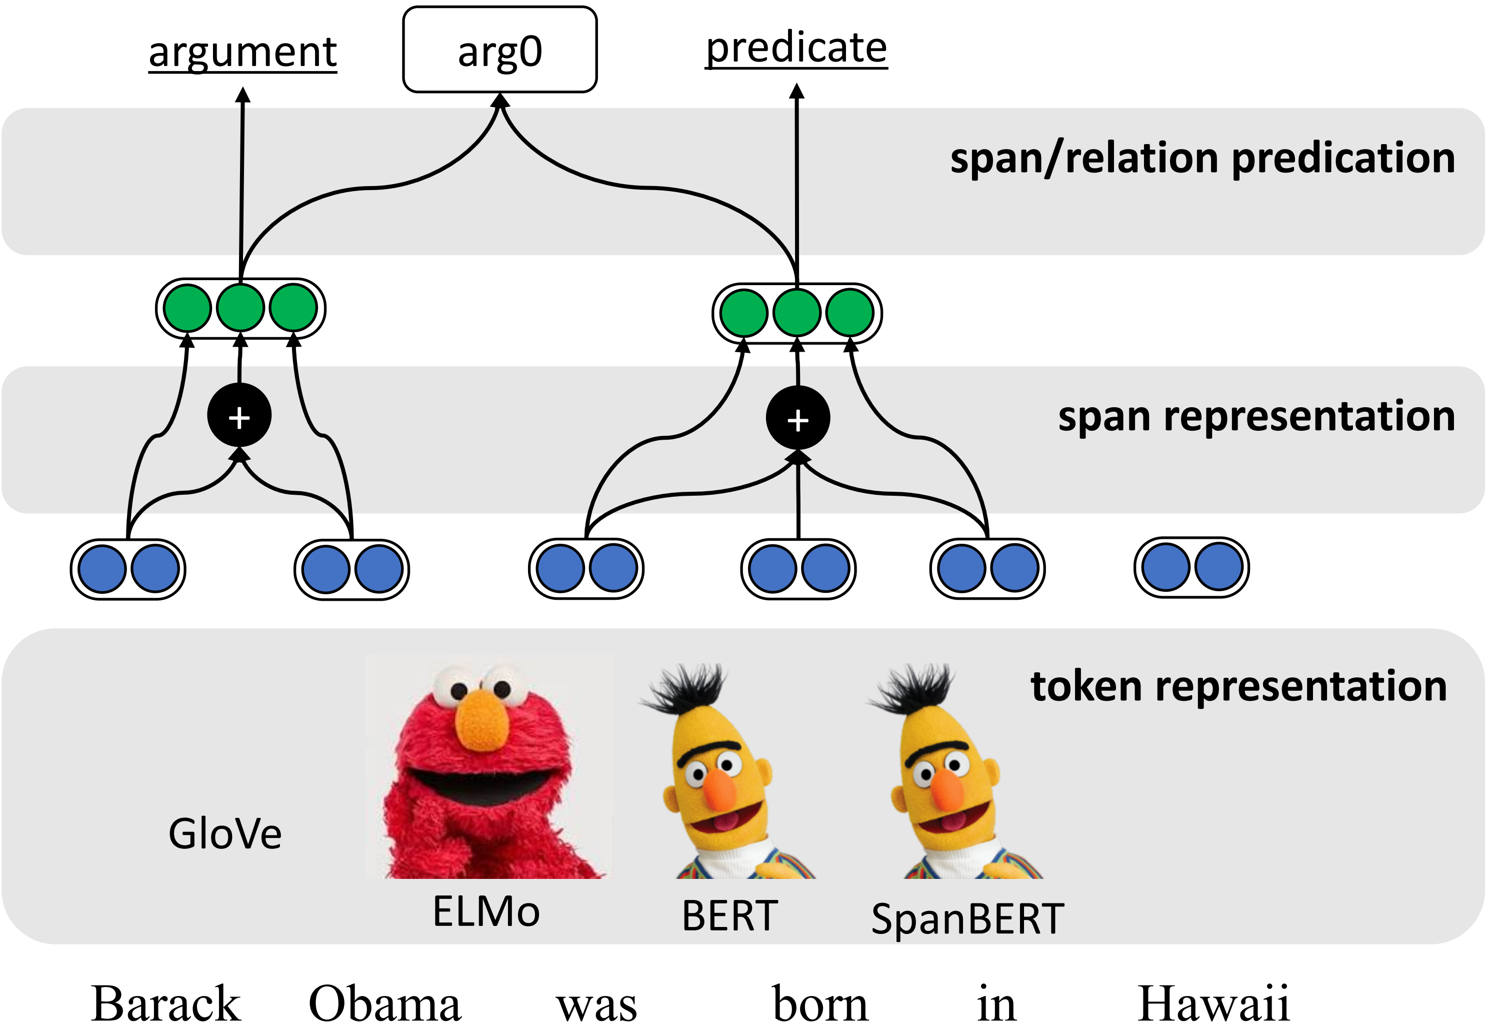

SpanRel Model

Shared across tasks

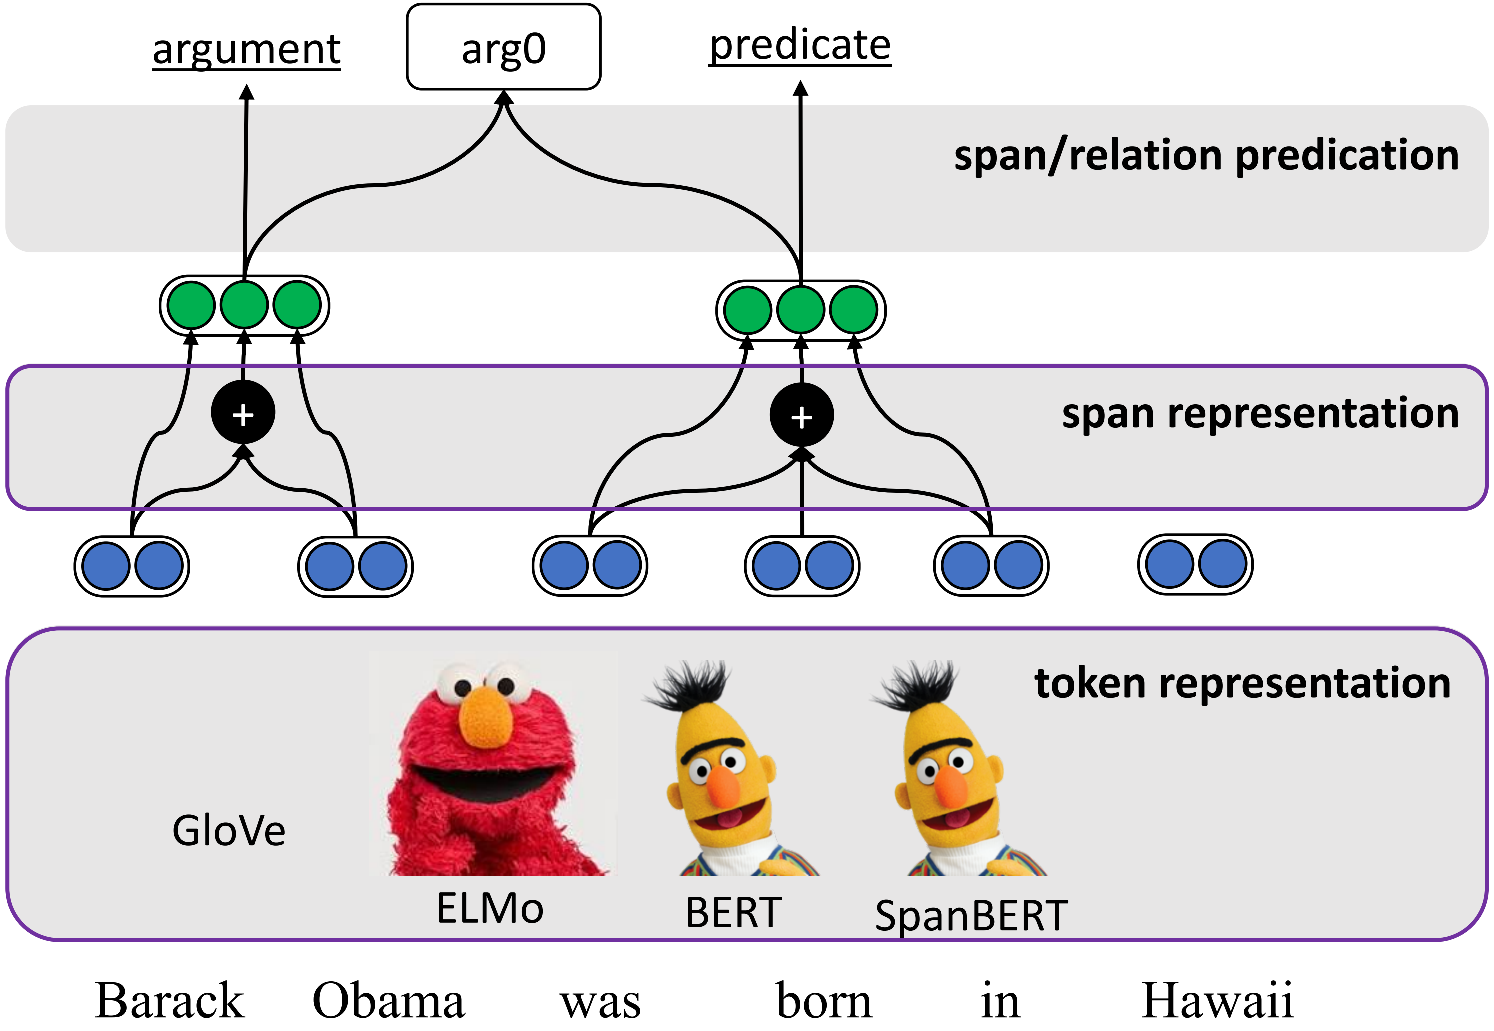

SpanRel Model

Task-specific

Shared across tasks

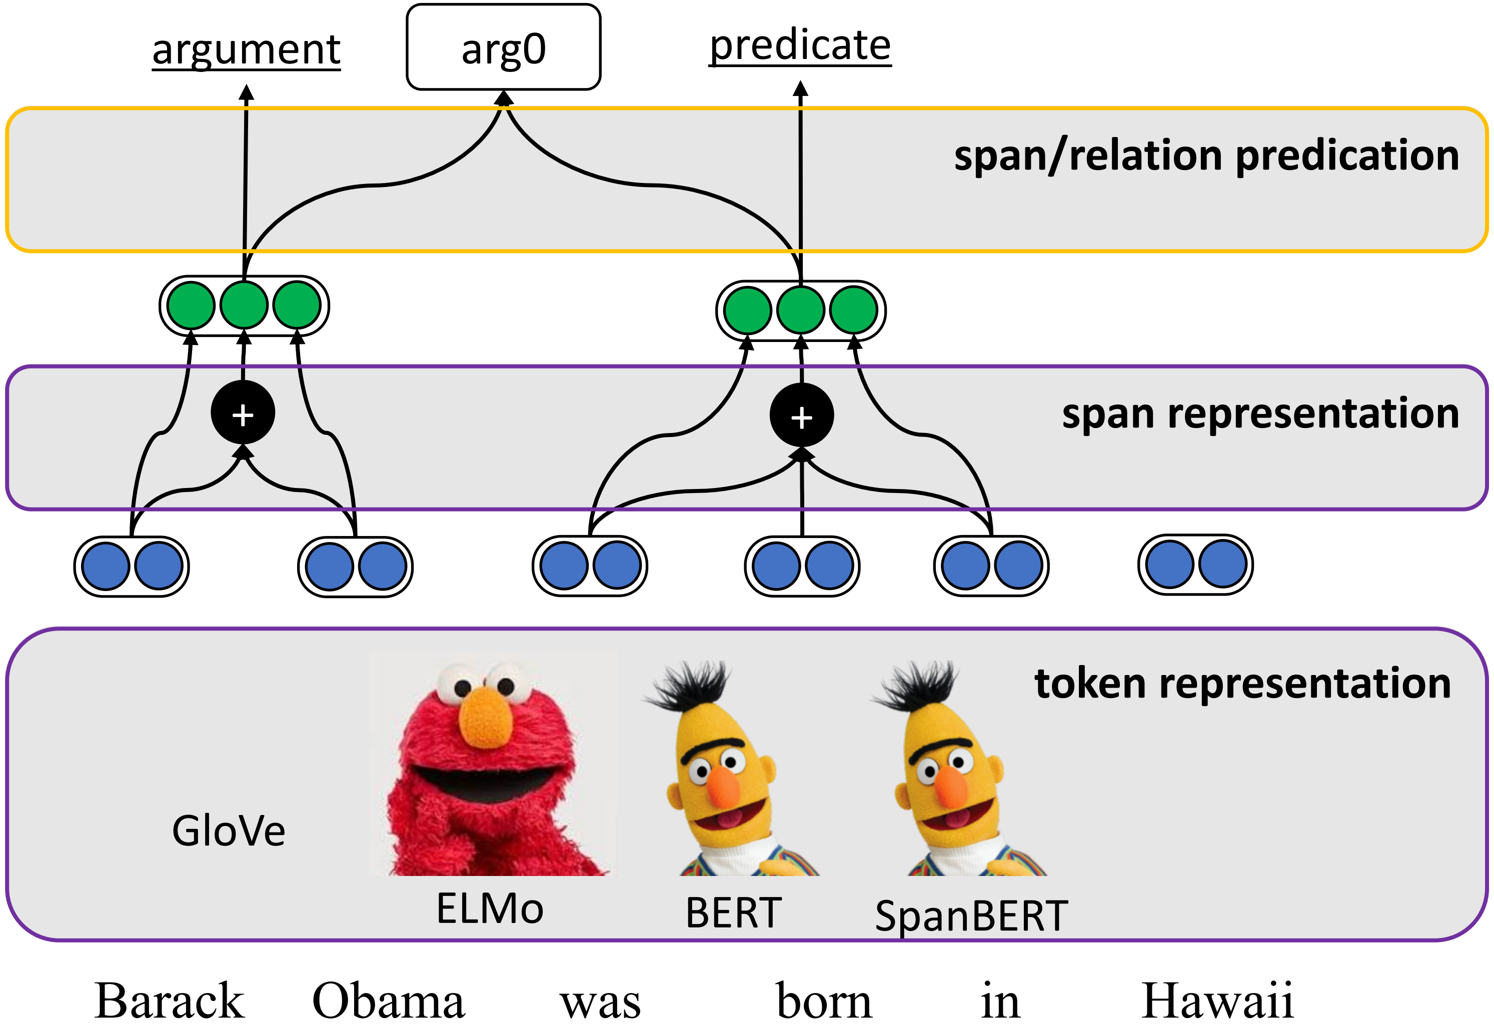

Experimental Settings

- General Language Analysis Datasets (GLAD) benchmark.
  - 8 datasets: CoNLL03, WLP, SemEval10, OntoNotes, OIE2016, PTB, SemEval14, MPQA 3.0.
- Evaluation metrics.
  - Major metric: span-based F1.
  - Task-specific metrics.
- Implementation details.
  - Token representation: GloVe, ELMo, BERT, SpanBERT.
  - Different pruning ratio/max span length for different tasks.

Comparison with SOTA

- Achieves comparable performances as task-specific SOTA methods

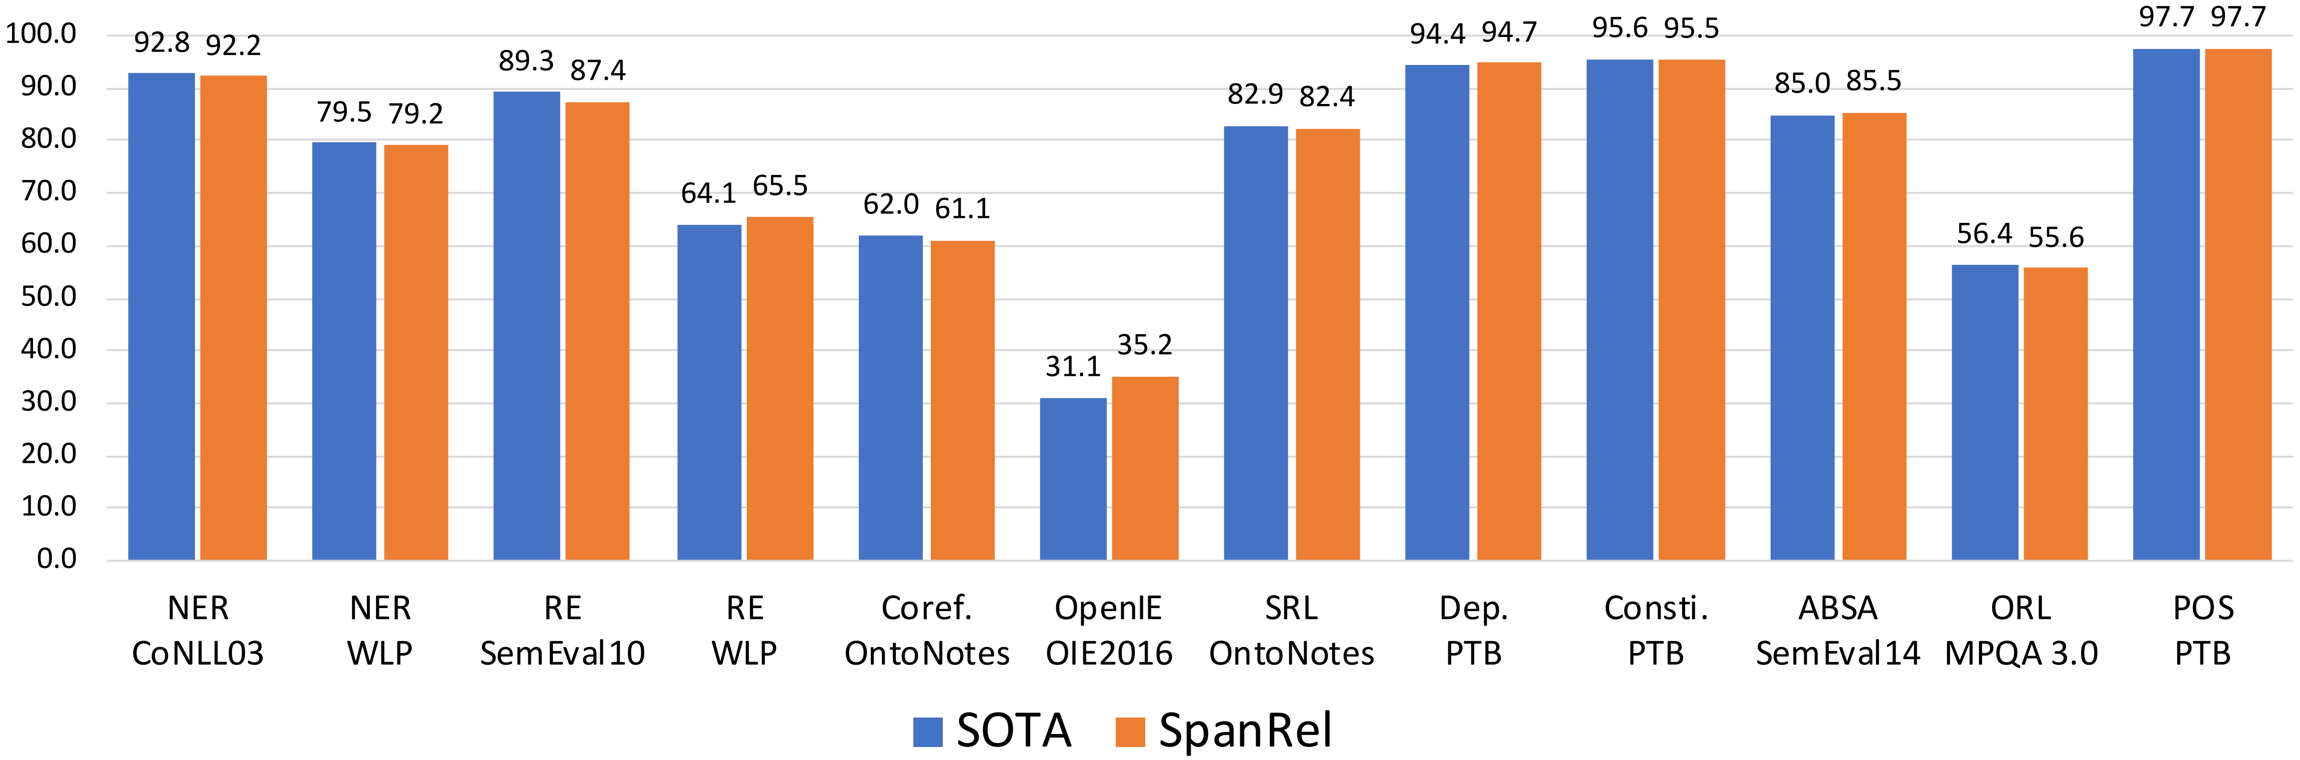

Multi-task Learning

- Significant improvements on 5/15 tasks with SpanBERT.

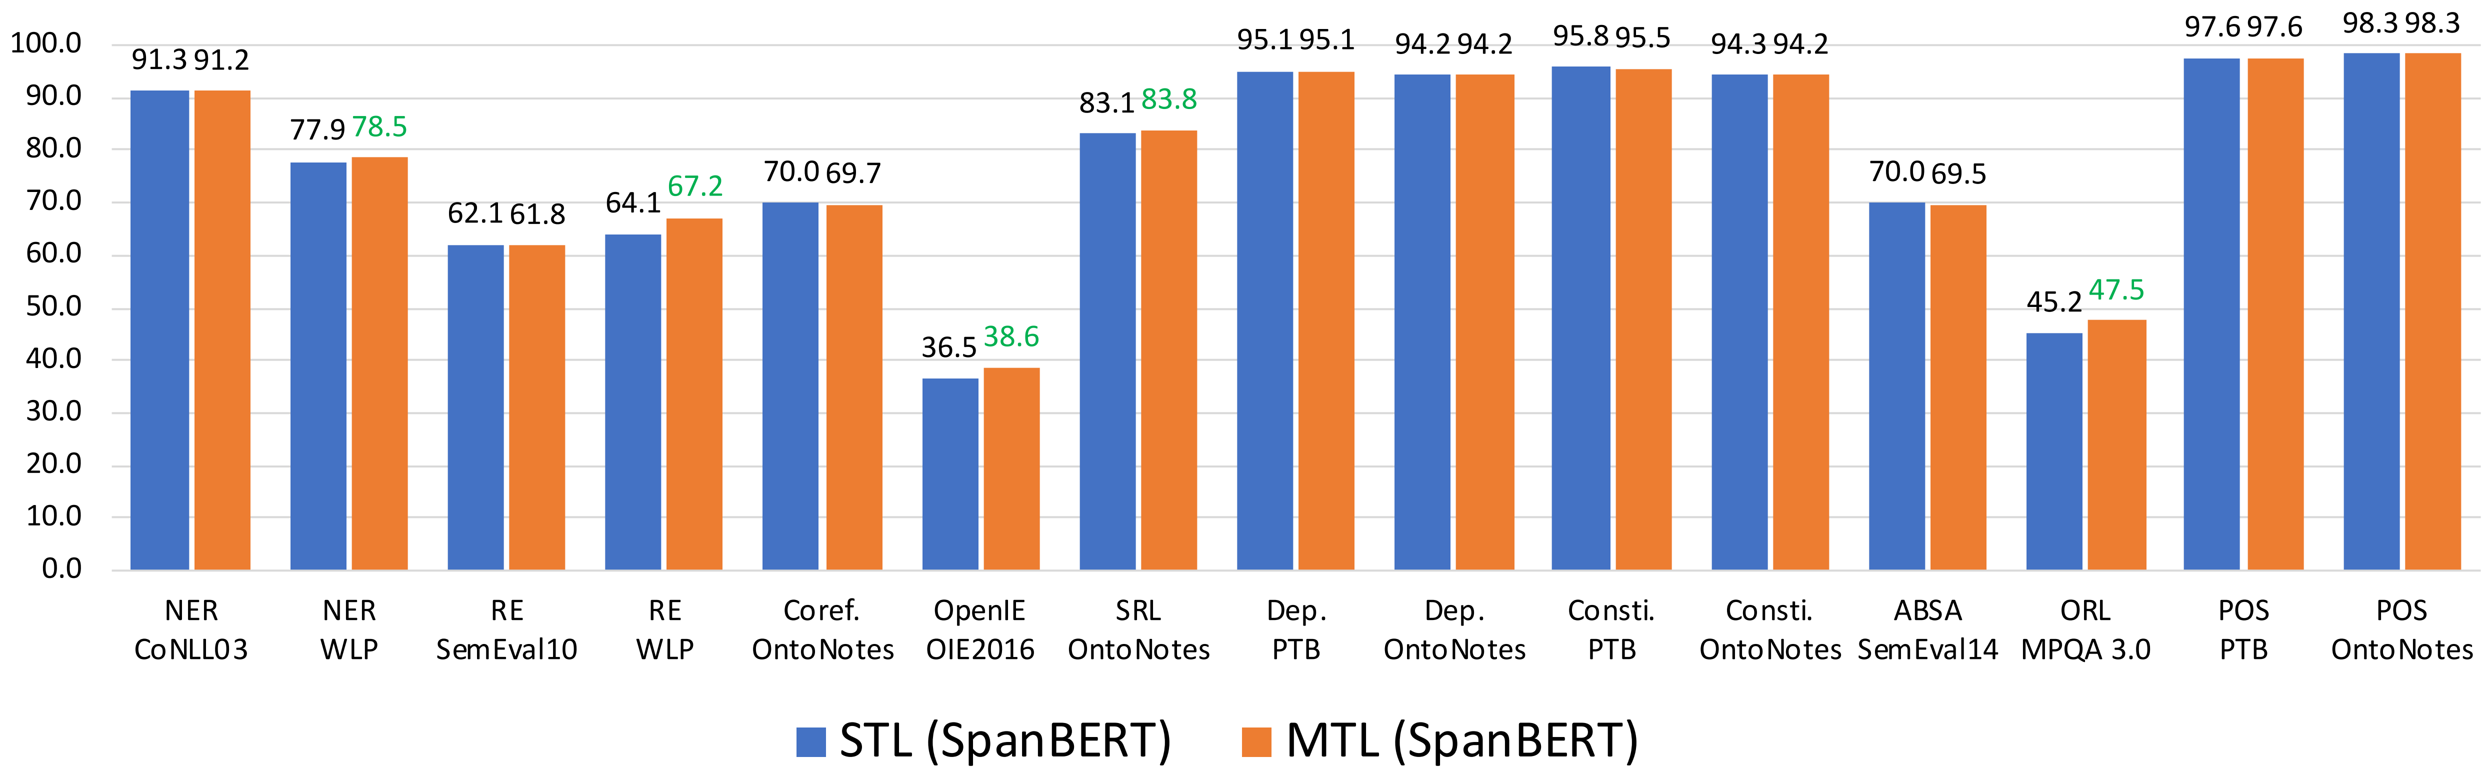

SpanBERT: 5 tasks↑, 0 tasks ↓

Multi-task Learning

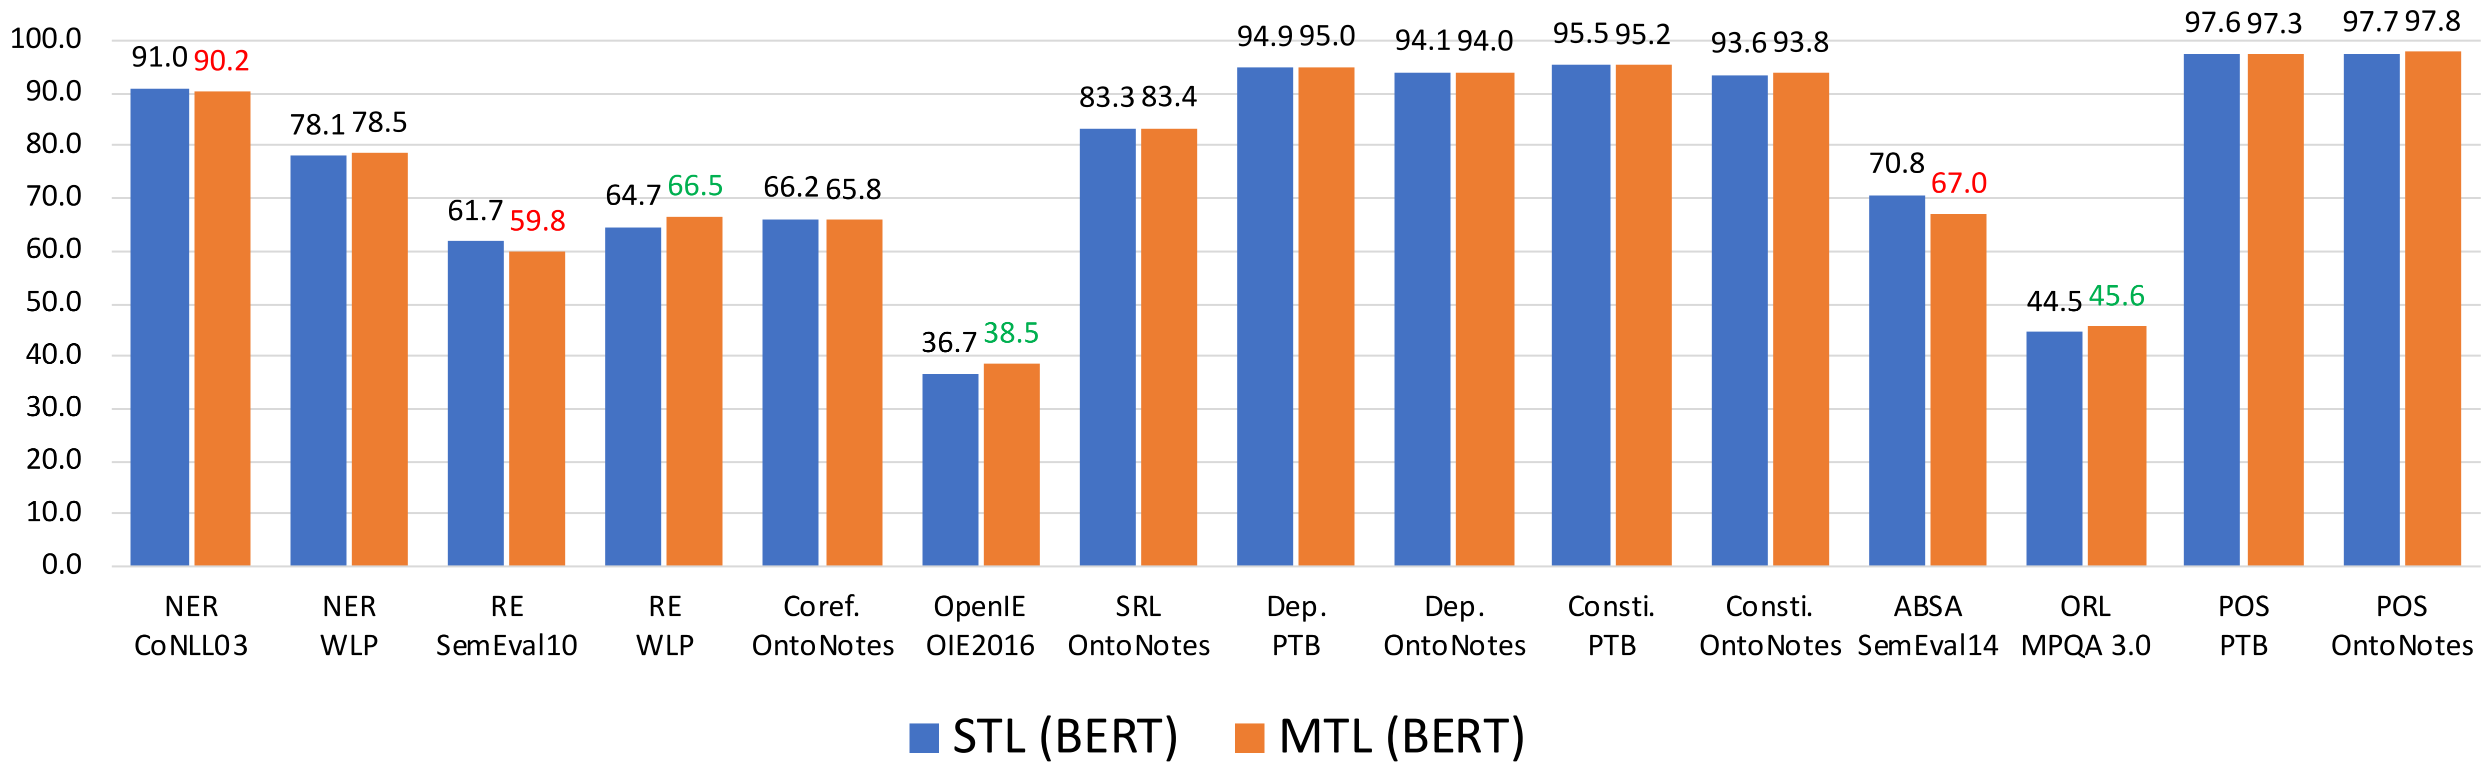

BERT: 3 tasks↑, 3 tasks ↓

Multi-task Learning

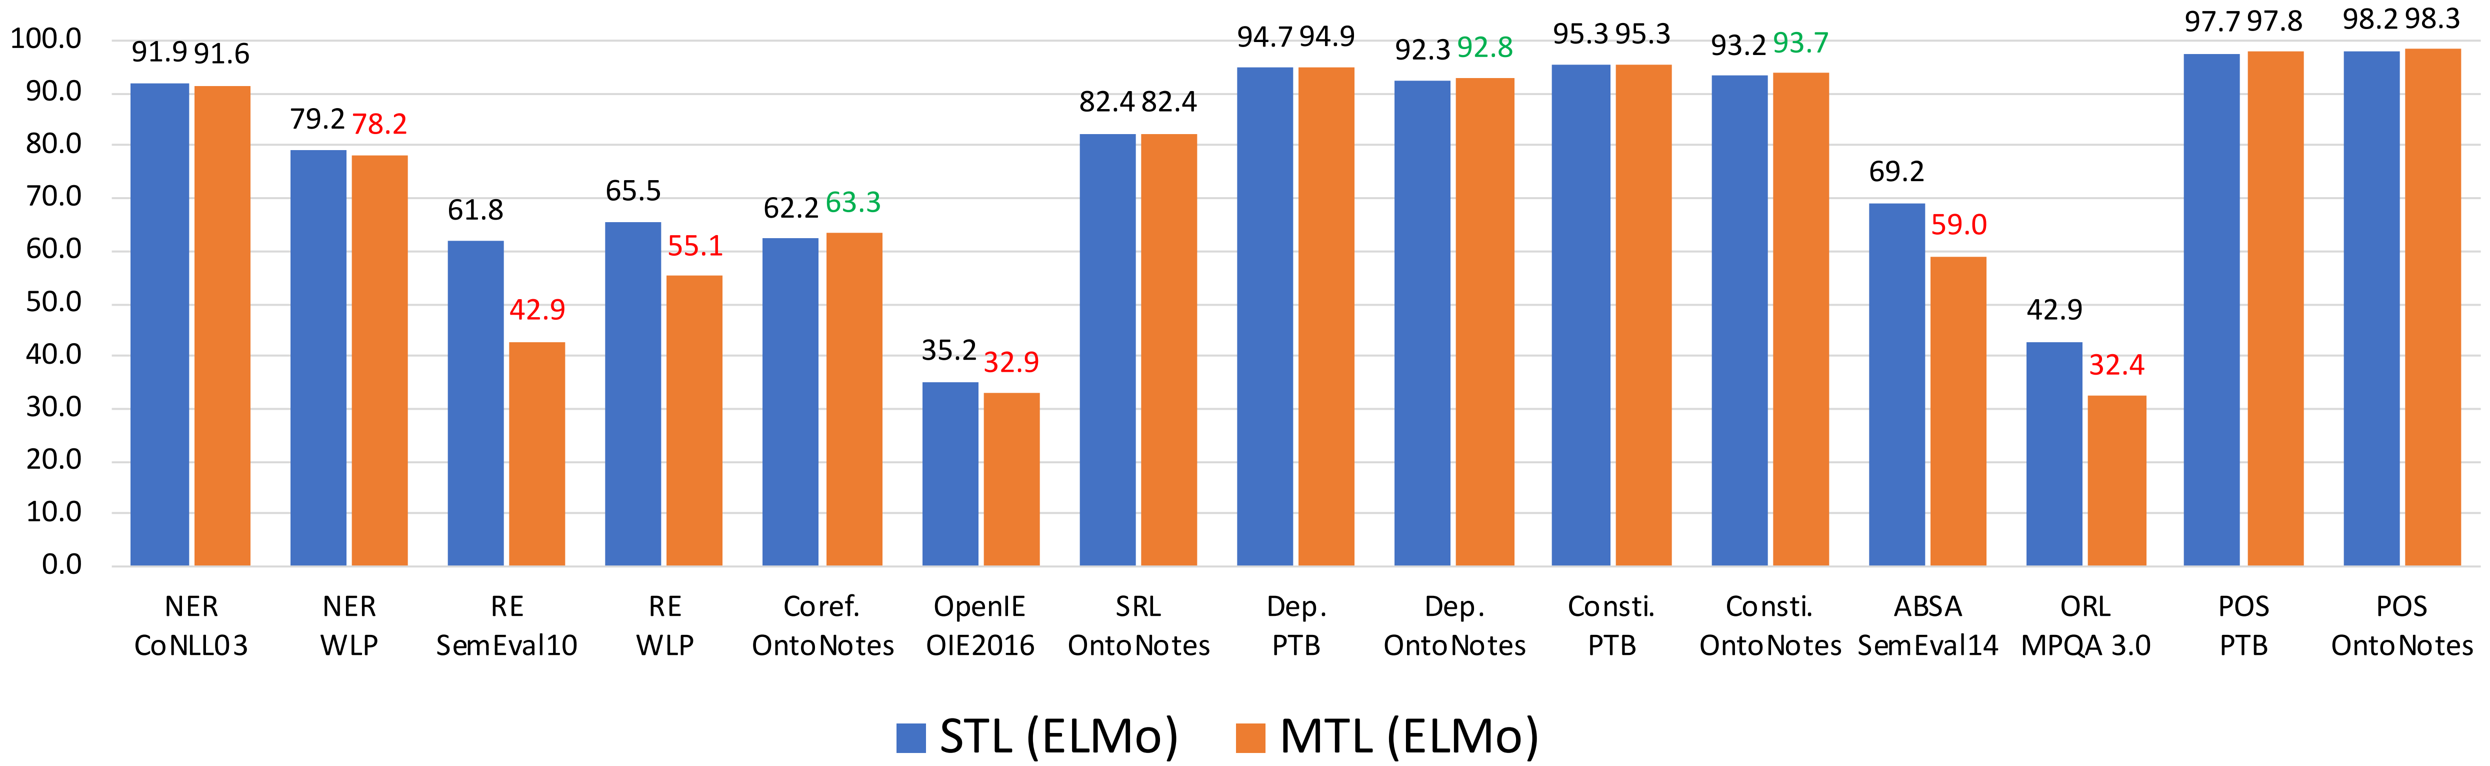

ELMo: 3 tasks ↑, 6 tasks ↓

Multi-task Learning

- Stronger models show consistent improvements from MTL, weaker models less so.

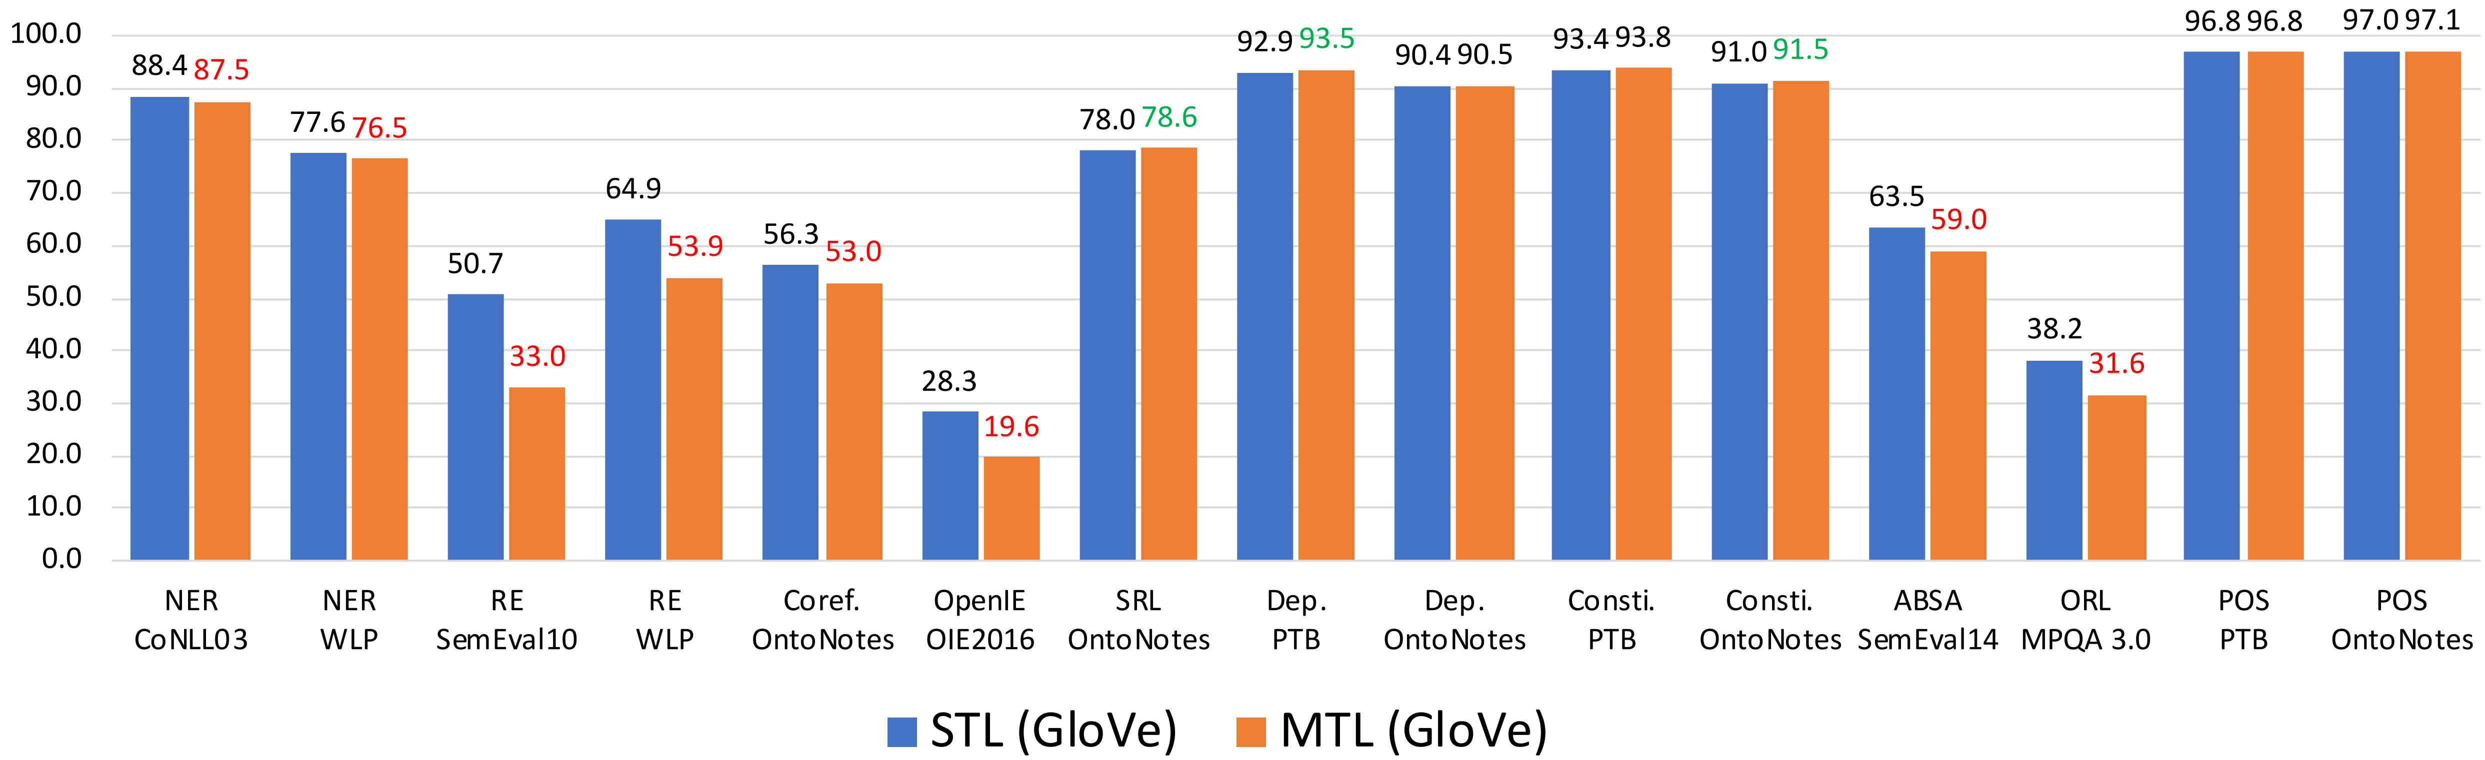

GloVe: 3 tasks↑, 8 tasks ↓

Task-relatedness Analysis

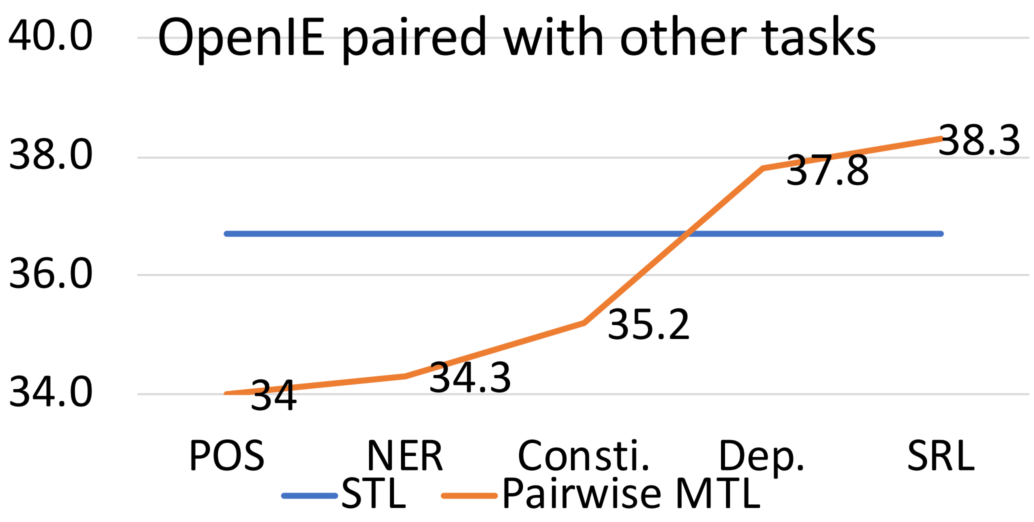

Tasks are not equally helpful to each other.

Task-relatedness Analysis

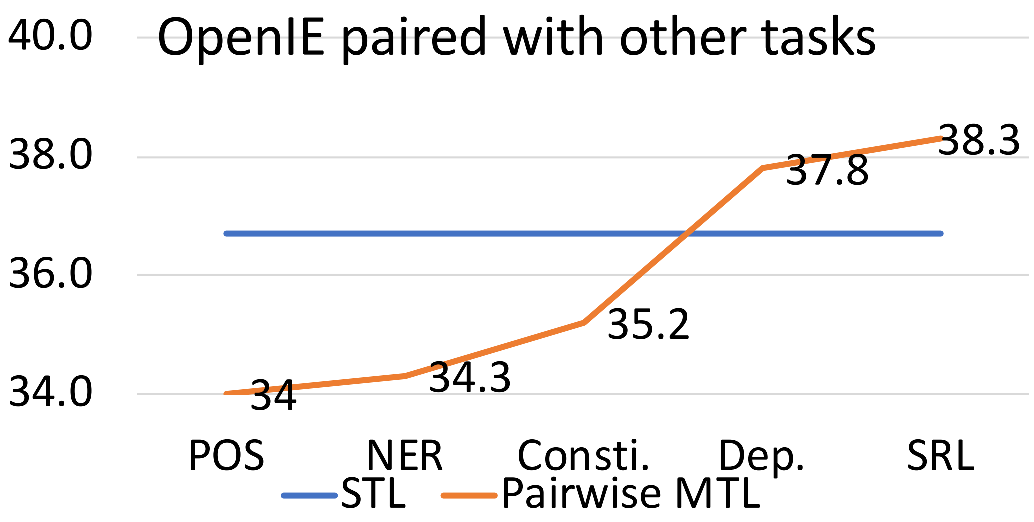

Tasks are not equally helpful to each other.

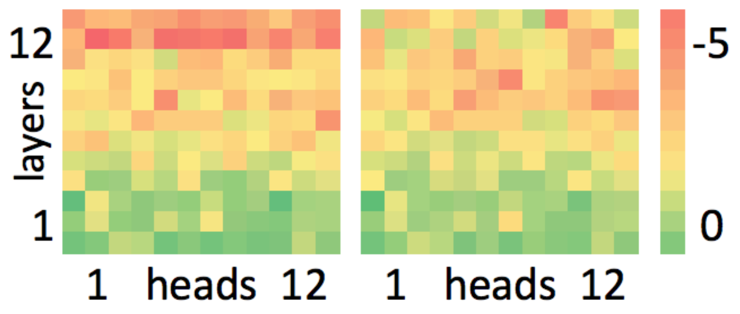

Attn similarity of OpenIE/POS and OpenIE/SRL.

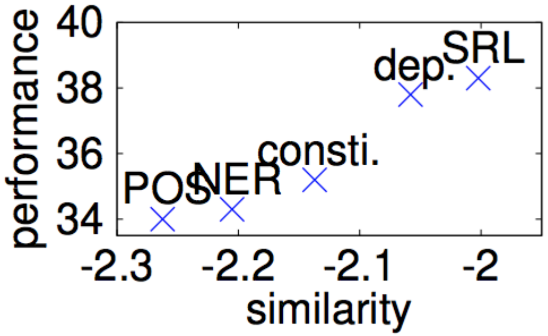

Tasks with similar attention as OpenIE help more.

Carnegie Mellon University School of Computer Science

The Ohio State University Department of Computer Science and Engineering

BOSCH Invented for life

NEULAB

Take away

1. A large variety of NLP tasks can be unified as span-relation prediction problems.

2. Multitask learning across a large number of different tasks helps, and how to better reconcile them is a challenging and rewarding future direction.

Paper: https://arxiv.org/pdf/1911.03822.pdf
Code: https://github.com/neulab/cmu-multinlp