Paraphrase Identification via Textual Inference

Ning Shi, Bradley Hauer, Jai Riley, Greg Kondrak

Alberta Machine Intelligence Institute (Amii)
Department of Computing Science
University of Alberta, Edmonton, Canada

Natural Language Inference

Natural **L**anguage **I**nference (NLI) involves three labels that describe the relationship between two sentences.

Entailment, Contradiction, Neutral

For example:

S_1: "This man is <u>surfing</u>."
S_2: "A man is <u>on water</u>."
Surfing: an aquatic activity or website browsing?

Paraphrase Identification

Paraphrase **I**dentification (PI) is the task of deciding whether two sentences convey the same meaning.

Hypothesis:

Paraphrasing corresponds to **bidirectional textual entailment**.

Prior work:

- A blend of modules complicates the analysis
- Bias to traditional PI methods
- Lacks any theoretical formalization

Our Contributions

We present the **first theoretical formalization** implying a **practical reduction of PI to NLI** (PI2NLI), validated by fine-tuning an NLI model for PI.

- A theoretical task reduction showing how PI can be reduced to NLI.
- Extensive evaluation across zero-shot and fine-tuning.
- We found fine-tuned NLI models can **outperform** dedicated PI models on PI datasets.

A **P-to-Q** reduction solves an instance of a problem **P** by combining the solutions of one or more instances of **Q**.

(Hauer and Kondrak, 2023)

Sentence-level Relations

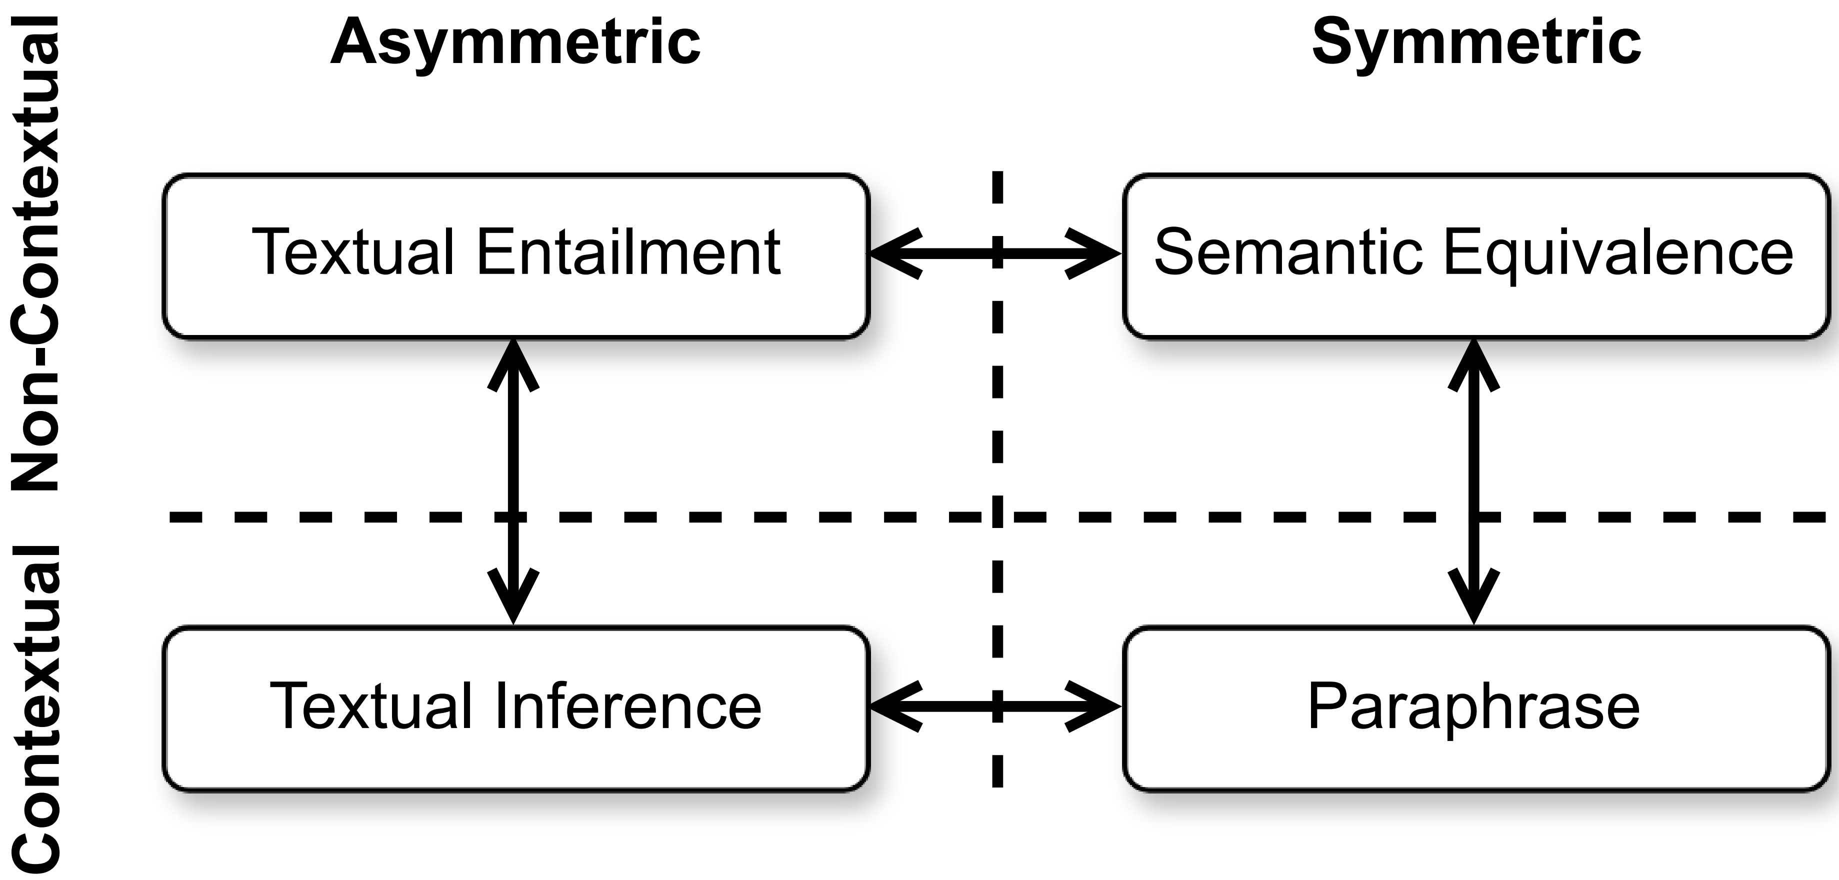

Equivalence and Paraphrasing

Semantic **Eq**uivalence relation, **SEQ(S_1, S_2)** :=
"the sentences S_1 and S_2 convey the same meaning"

Paraphrase **R**elation, **PR(C, S_1, S_2)** := "the sentences S_1 and S_2 convey the same meaning
<u>given the context C</u>"

The relationship in between:

$$\mathbf{SEQ(S_1, S_2) \Leftrightarrow \forall C : PR(C, S_1, S_2)}$$

Example:
S_1: "<u>We</u> must work hard to win this election."
S_2: "<u>The Democrats</u> must work hard to win this election."

Entailment and Inference

Textual **E**ntailment, $TE(S_1, S_2) :=$
"the sentence S_2 can be inferred from the sentence S_1"

Textual **I**nference, $TI(C, S_1, S_2) :=$ "the sentence S_2 can be inferred from the sentence S_1 given the context C"

The relationship in between:

$$\mathbf{TE(S_1, S_2) \Leftrightarrow \forall C : TI(C, S_1, S_2)}$$

Example:

S_1: "This man is surfing."
S_2: "A man is on water."

Proposition

Given context C, sentences S_1 and S_2 are paraphrases if and only if they can be mutually inferred from each other.

Formally:

$$\mathbf{PR}(C, S_1, S_2) \Leftrightarrow \mathbf{TI}(C, S_1, S_2) \wedge \mathbf{TI}(C, S_2, S_1)$$

Context:

- Context includes common sense and world knowledge.
- In practice, context is embedded in the data distribution.

Data Adaptation

Positive PI instances:

We convert each positive PI instance into two distinct NLI positive instances, one in each direction.

$PI(S_1, S_2, True)$ -> $NLI(S_1, S_2, Entailment)$ **AND** $NLI(S_2, S_1, Entailment)$

Negative PI instances:

We generate a negative NLI instance (randomly selected as either contradiction or neutral) in one randomly selected direction.

$PI(S_1, S_2, False)$ -> $NLI(S_1, S_2, Contradiction)$ **OR** $NLI(S_2, S_1, Contradiction)$
$NLI(S_1, S_2, Neutral)$ **OR** $NLI(S_2, S_1, Neutral)$

Results

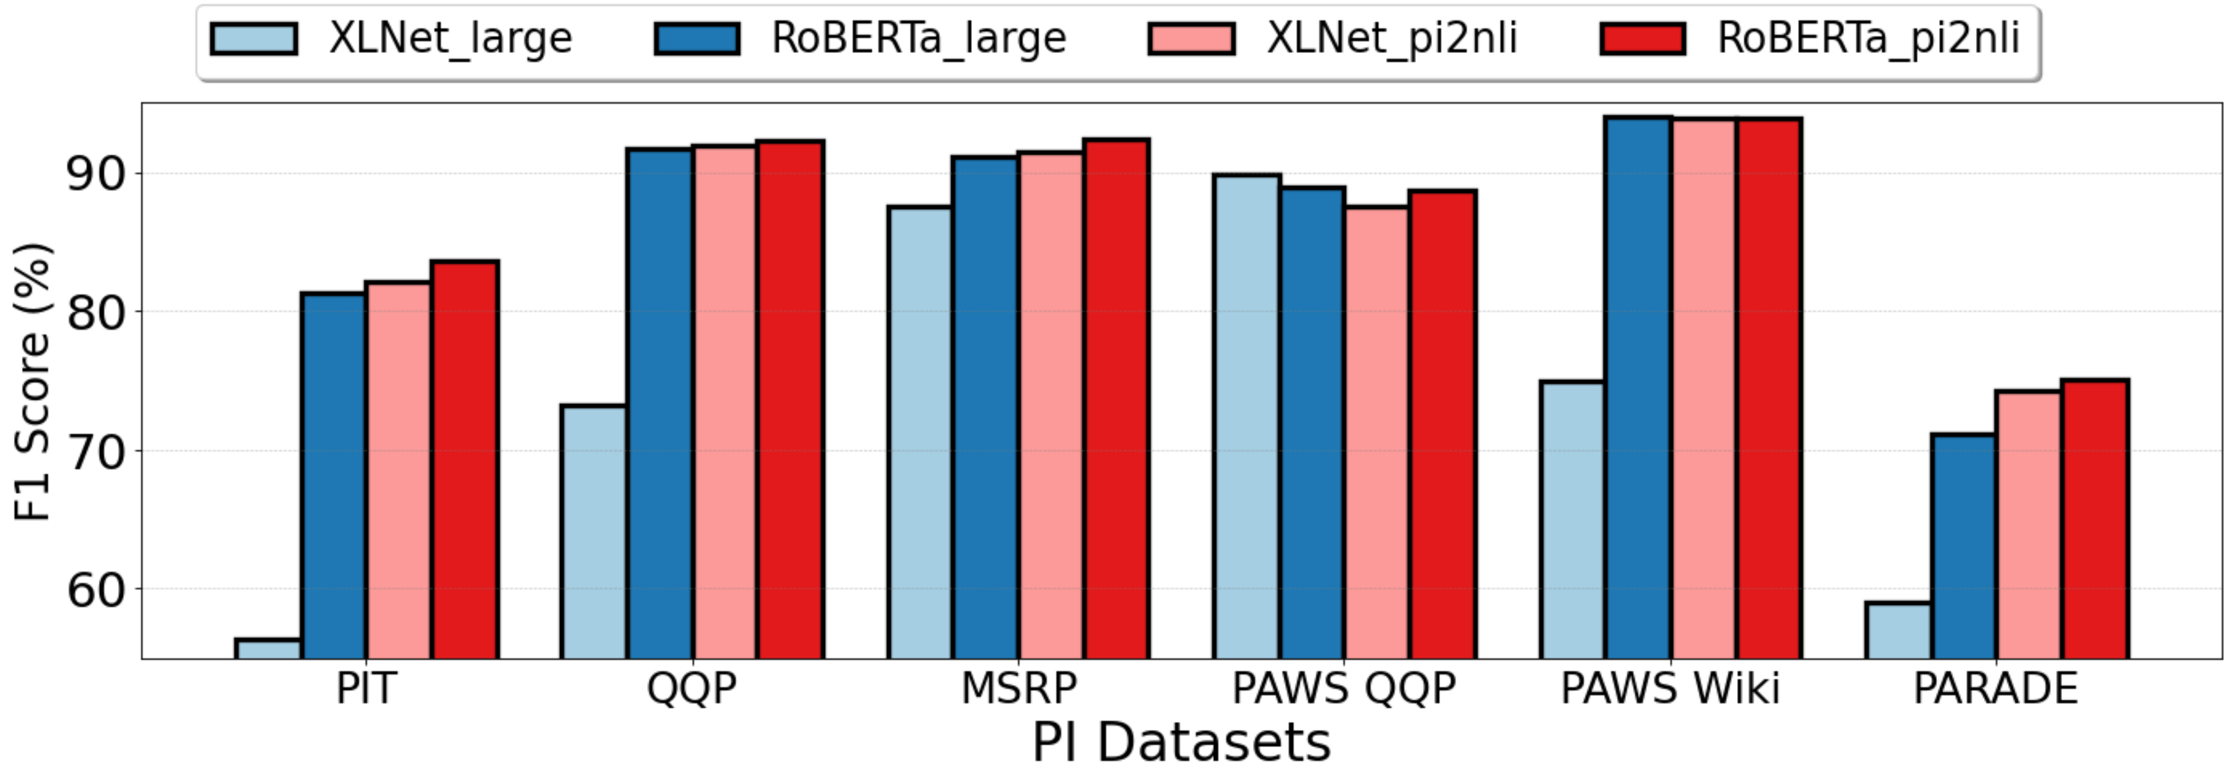

XLNet_large (Yang et al., 2019); RoBERTa_large (Liu et al., 2019)

PAWS QQP and PAWS Wiki include **adversarial** example created by word scrambling and back-translation.

Ablation

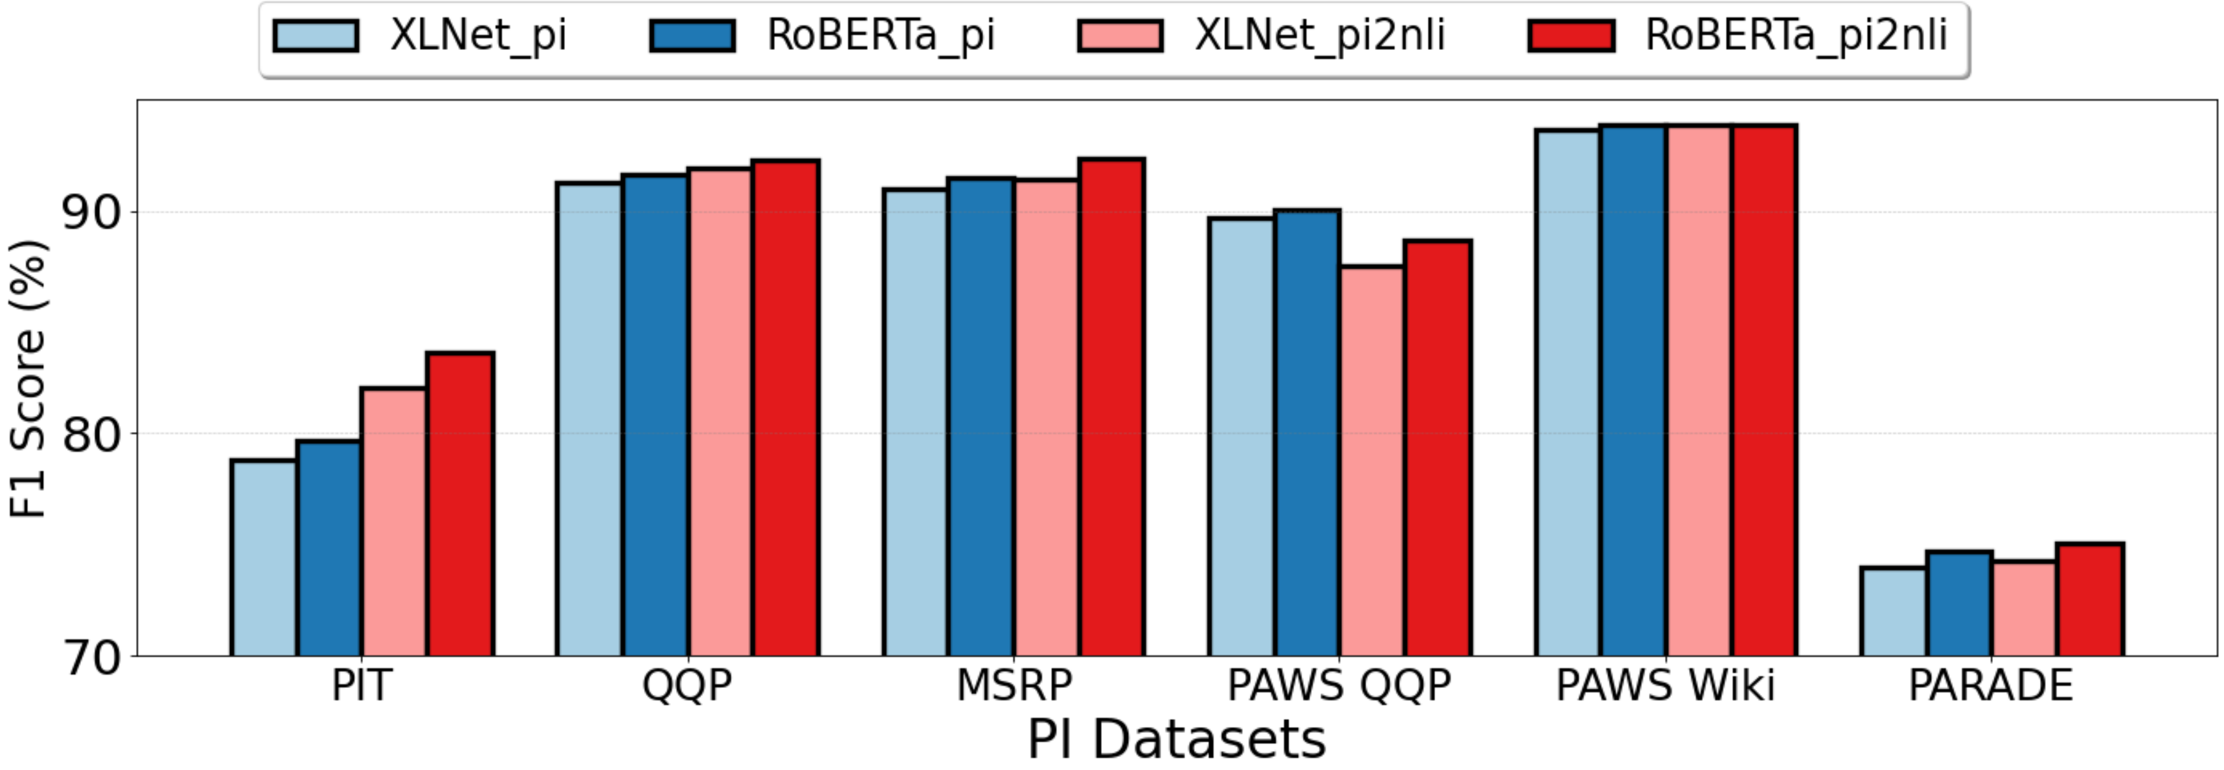

XLNet_pi, RoBERTa_pi (Nie et al., 2020)
The **same** language models with classification heads initialized from **scratch**.

Analysis - Context

A paraphrase identified in one dataset might **NOT** necessarily be considered a valid paraphrase in the other.

We view this **adjustment** as the process of how models learn the **context** inherent in each PI dataset.

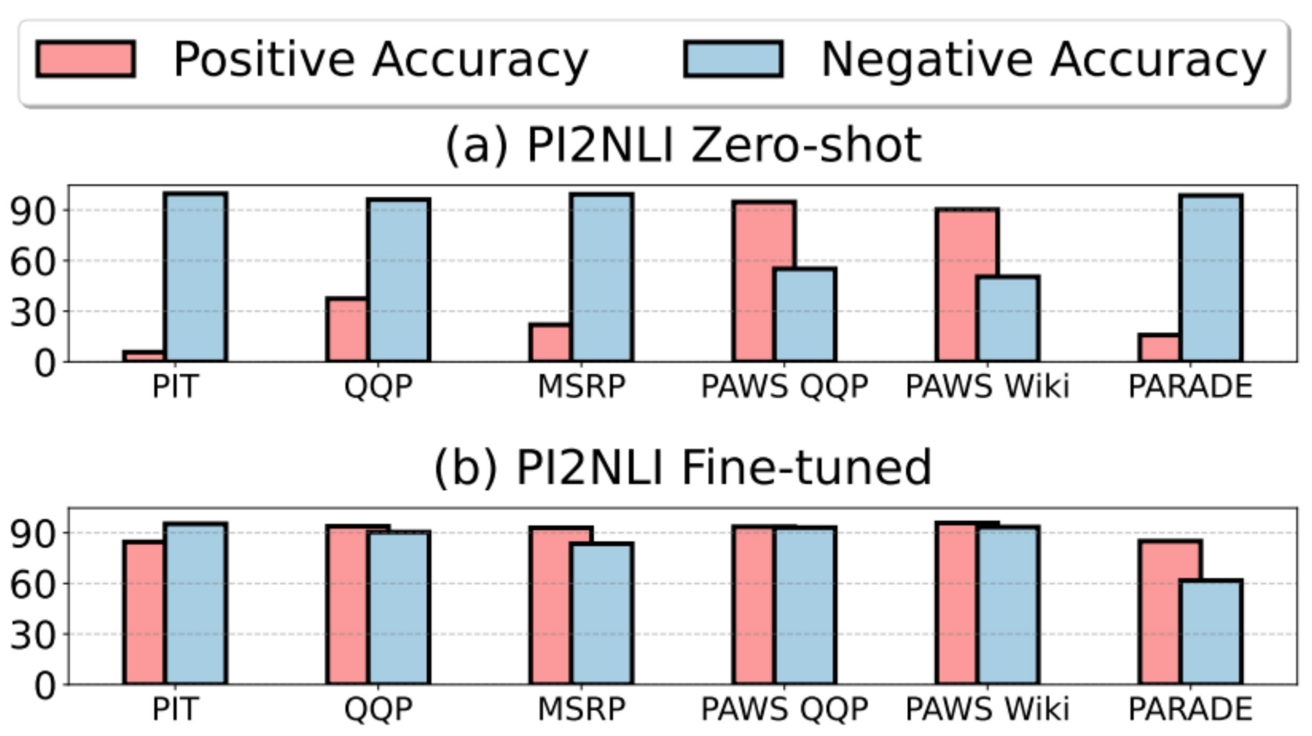

Conclusion

We have presented PI2NLI, the first attempt to reduce PI to NLI.

- PI can be reduced to NLI **theoretically** and **empirically**.
- Fine-tuned NLI models can **outperform** PI models on PI datasets.
- Applying PI2NLI in a zero-shot setting show the **limitations** in the current PI datasets.

Future:
- Using NLI to refine PI data?
- Using PI models to solve NLI?
- Using one single model to solve both PI and NLI?

github.com/ShiningLab/PI2NLI