Sentence Simplification with Deep Reinforcement Learning

Xingxing Zhang and Mirella Lapata

Institute for Language, Cognition and Computation
School of Informatics, University of Edinburgh

9 September, 2017

Sentence Simplification: Motivation

Mathematical secrets of ancient tablet unlocked after nearly a century of study

By Maev Kennedy, The Guardian
09/01/2017

Word Count **800**

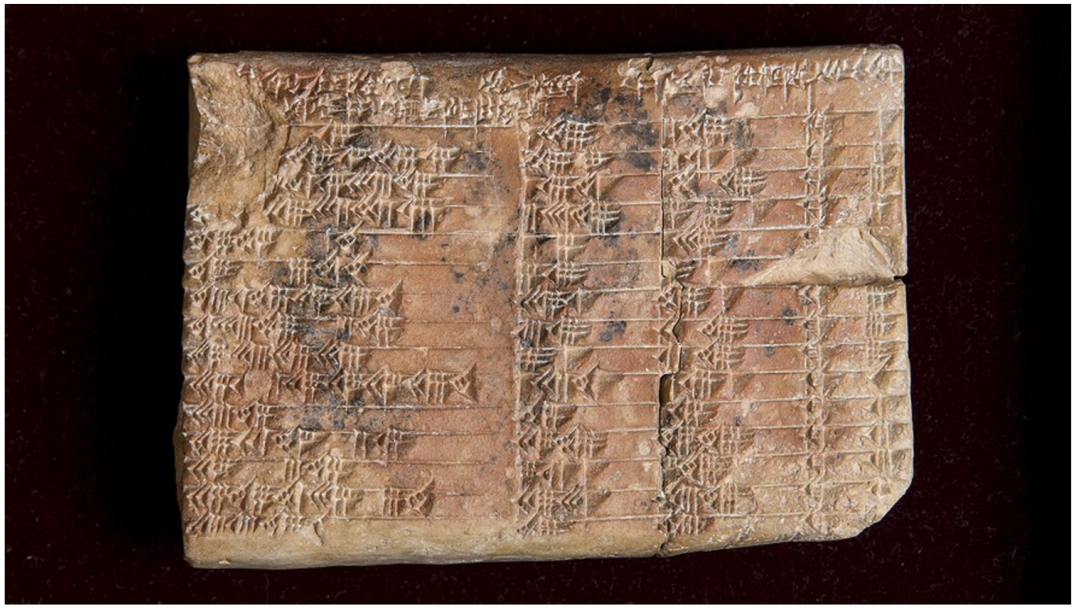

Scientists have discovered the purpose of a famous 3,700-year-old Babylonian clay tablet, revealing it is the world's oldest and most accurate trigonometric table, possibly used to calculate how to construct palaces and temples and build canals. The new research shows the Babylonians, not the Greeks, were the first to study trigonometry – the study of triangles – and reveals an ancient mathematical sophistication that had been hidden until now. Photo by: UNSW/Andrew Kelly

At least 1,000 years before the Greek mathematician Pythagoras looked at a right-angled triangle and worked out that the square of the longest side is always equal to the sum of the

Sentence Simplification: Motivation

Ancient clay tablet rewrites history of math

By Maev Kennedy, The Guardian, adapted by Newsela staff
09/01/2017

Word Count **374**

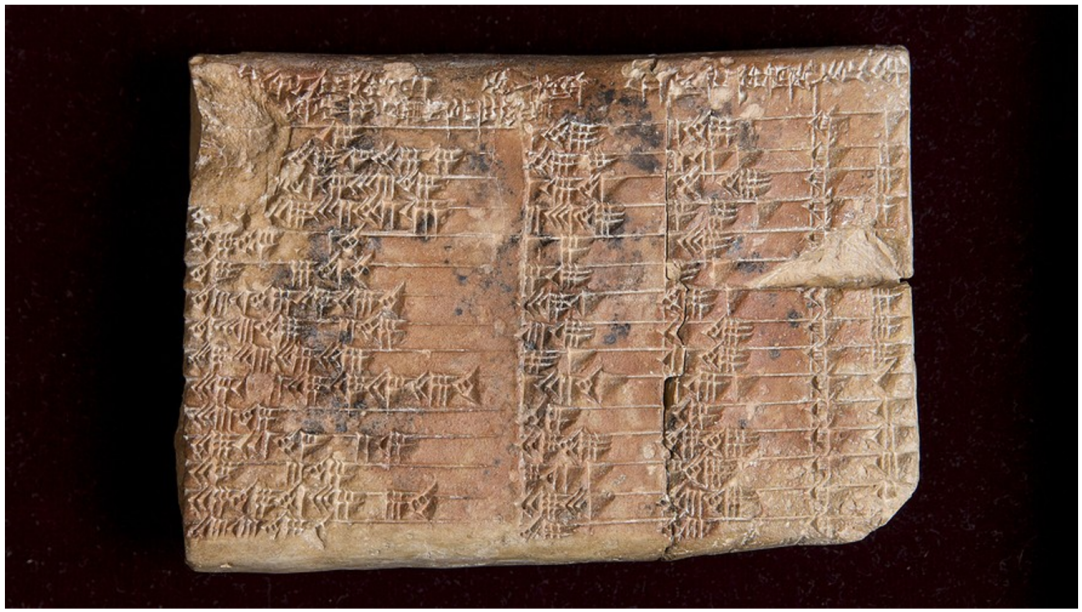

Scientists have discovered the purpose of a famous 3,700-year-old Babylonian clay tablet, revealing it is the world's oldest and most accurate trigonometric table, possibly used to calculate how to construct palaces and temples and build canals. The new research shows the Babylonians, not the Greeks, were the first to study trigonometry – the study of triangles – and reveals an ancient mathematical sophistication that had been hidden until now. Photo by: UNSW/Andrew Kelly

One of the most famous people to study math was Pythagoras. Pythagoras lived about 2,500 years ago. Pythagoras is famous for his idea about triangles.

Sentence Simplification: The Task

- Rewrite a complex sentence into a simpler one.
- Word/phrase replacement, word/phrase deletion

  Most Americans **favor** the *Defense Department 's* decision *to lift a ban on women in combat jobs* and do not believe it will hurt military *effectiveness* , a new poll shows .

  Most Americans **support** the decision and do not believe it will hurt the military , a new poll shows .

- Sentence splitting

  In 1883 , Faur married Marie Fremiet , **with whom he** had two sons .

  In 1883 Faur married Marie Fremiet . **They** had two sons .

Related Work

Rule-based Methods (to name a few)

- Carroll et al., (1999)
- Siddharthan et al., (2004)

Syntax-based Machine Translation (SBMT)

- Zhu et al., (2010)
- Woodsend and Lapata, (2011)
- Xu et al., (2016)

Phrase-based Machine Translation (PBMT)

- Wubben et al., (2012)
- Coster and Kauchak, (2011)
- Kauchak, (2013)

Simplification == Sequence to Sequence Learning?

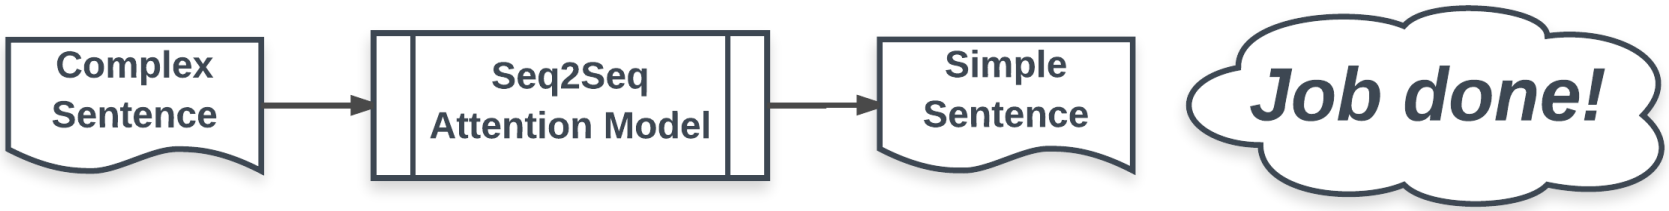

Simplification == Sequence to Sequence Learning?

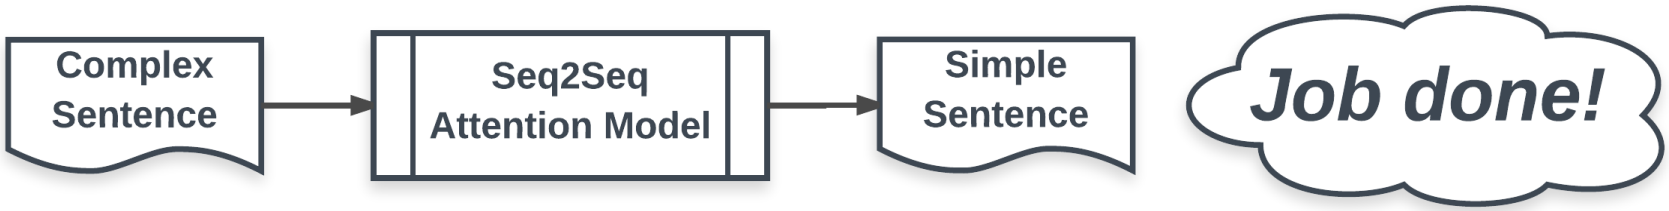

- Problem: Seq2SeqA **Repeats** the Complex Sentence
  - Why?

Dataset	**Copy Operation (%)**
Newsela	73%
WikiSmall	82%
WikiLarge	72%

Task Specific Constraints

What does a good simplification look like?

- Target Sentences are **Simpler**!
- They **Preserve the Meaning** of the Original Sentence!
- They must be **Fluent**!

Task Specific Constraints

What does a good simplification look like?

- Target Sentences are **Simpler**!
- They **Preserve the Meaning** of the Original Sentence!
- They must be **Fluent**!

How can we instill this knowledge in a Seq2Seq model?

- Change the training algorithm?

Task Specific Constraints

What does a good simplification look like?

- Target Sentences are **Simpler**!
- They **Preserve the Meaning** of the Original Sentence!
- They must be **Fluent**!

How can we instill this knowledge in a Seq2Seq model?

- Change the training algorithm?
- Reinforcement Learning (Policy Gradient)
  Williams, (1992); Ranzato et al., (2015)

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

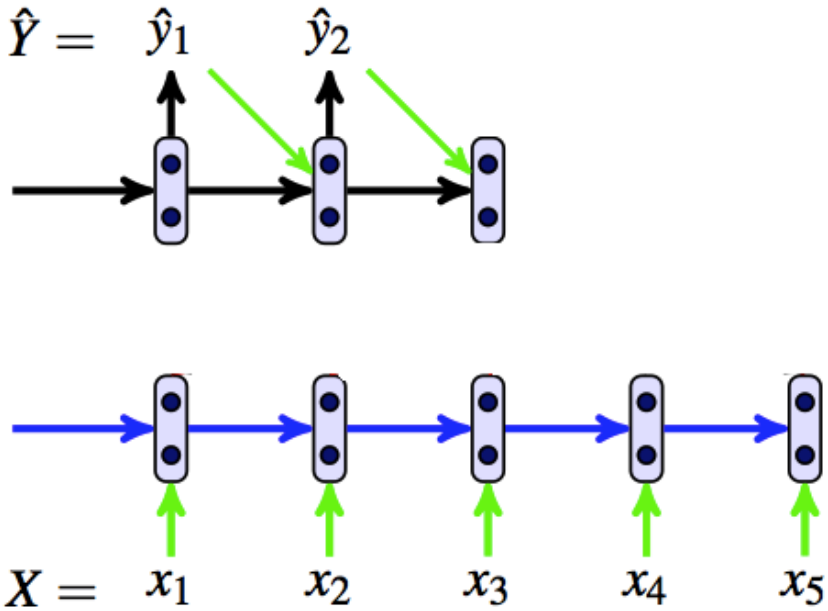

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

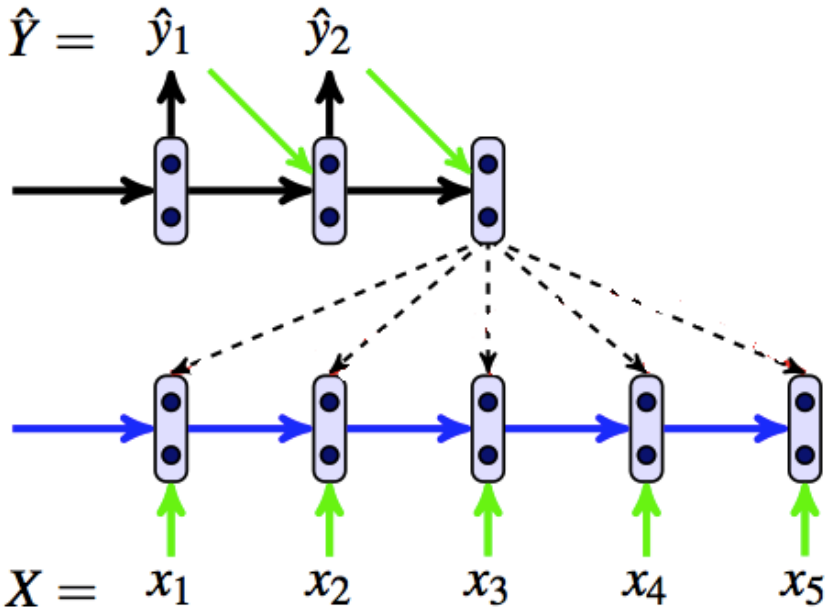

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

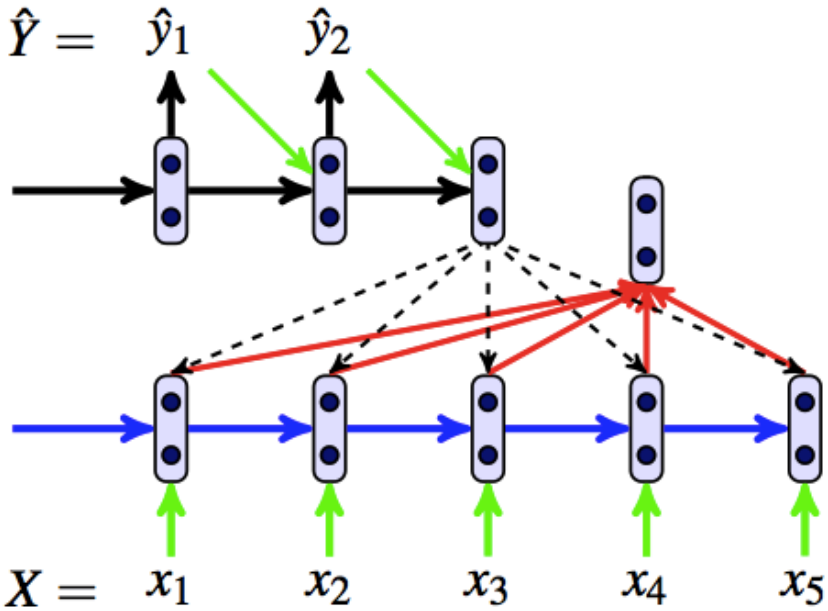

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

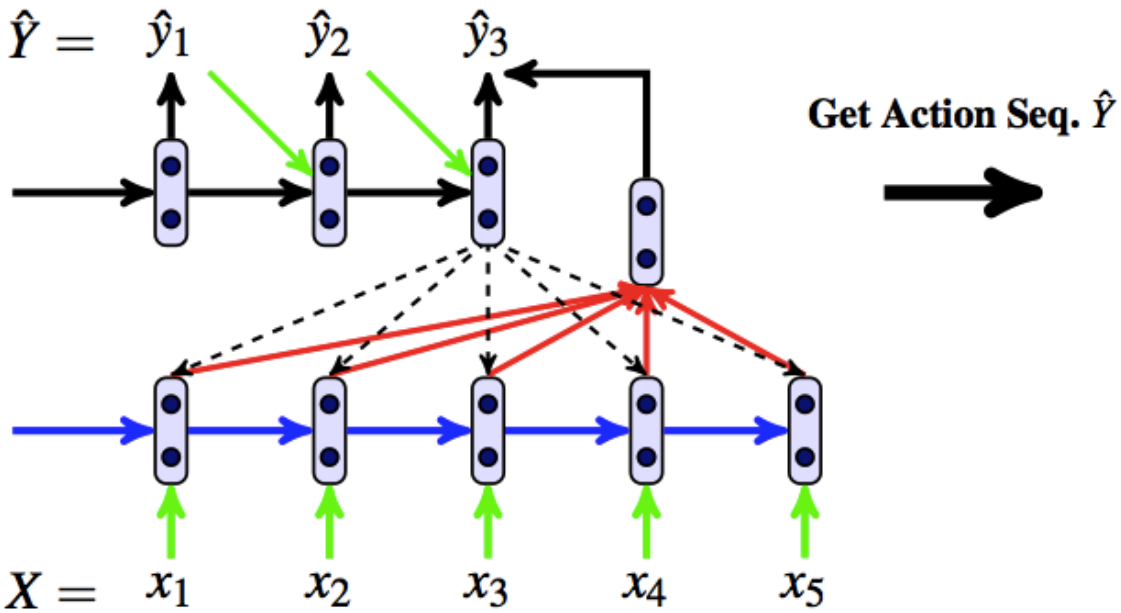

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

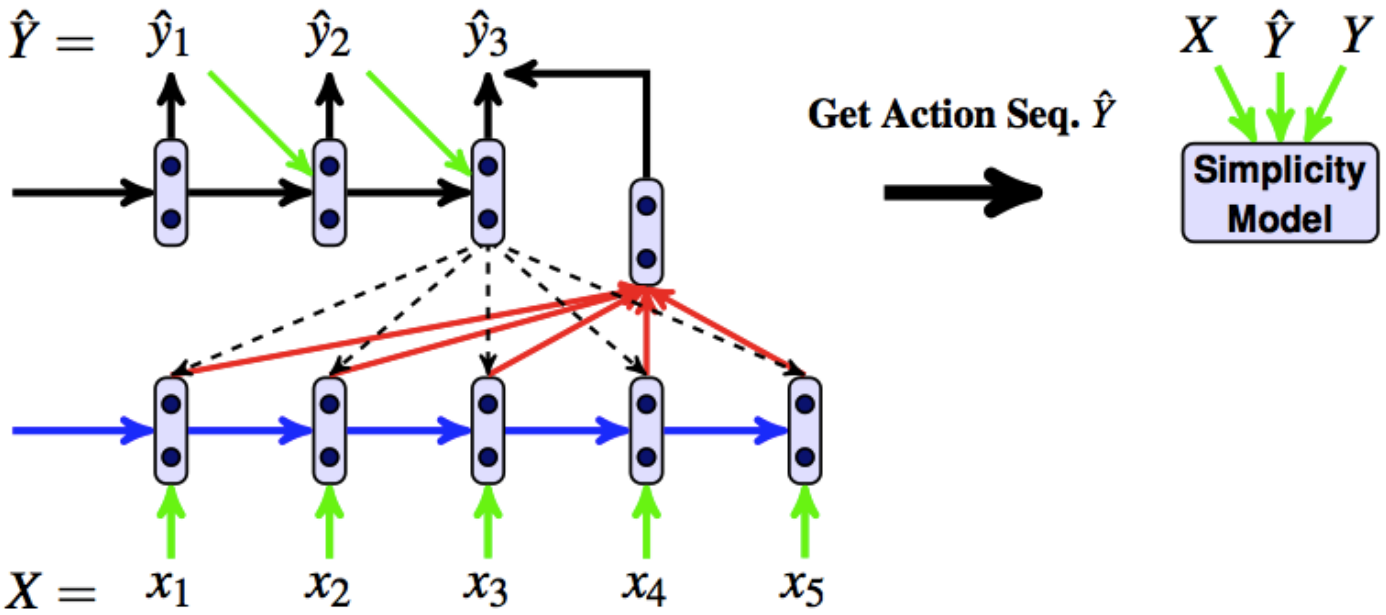

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

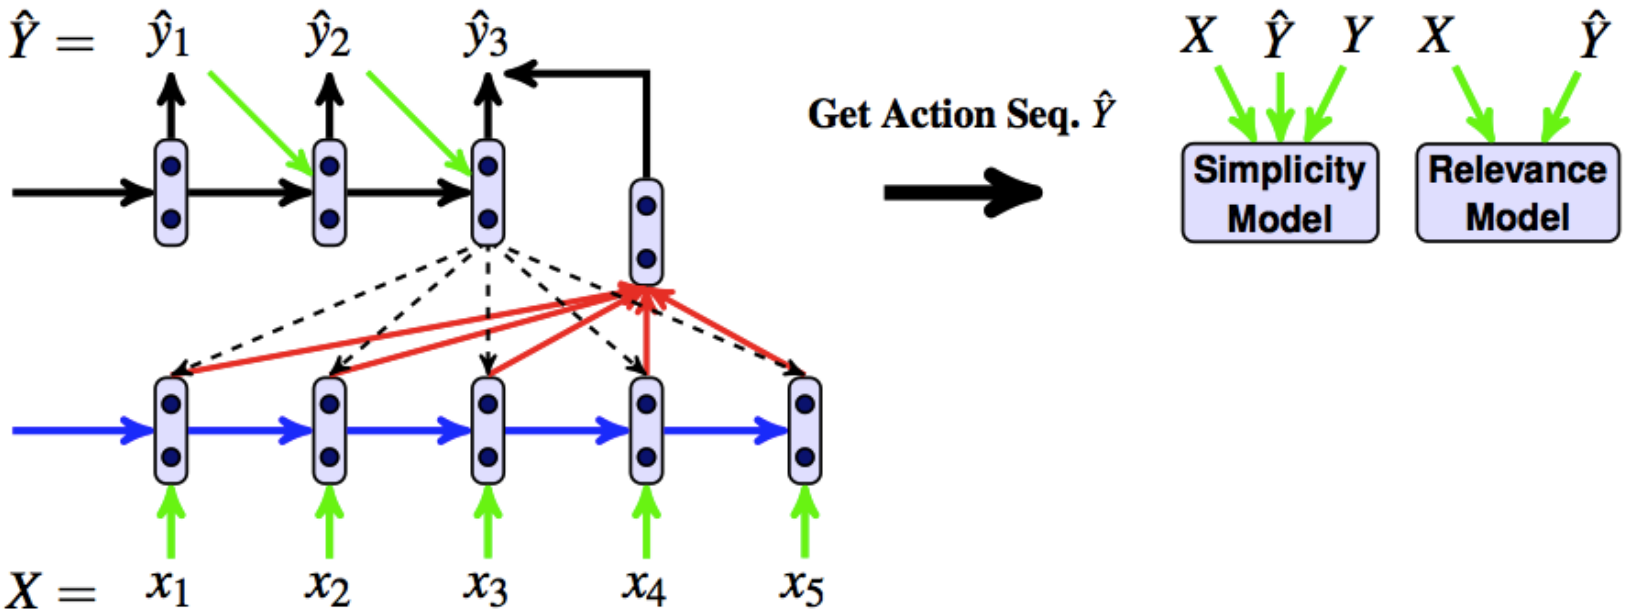

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

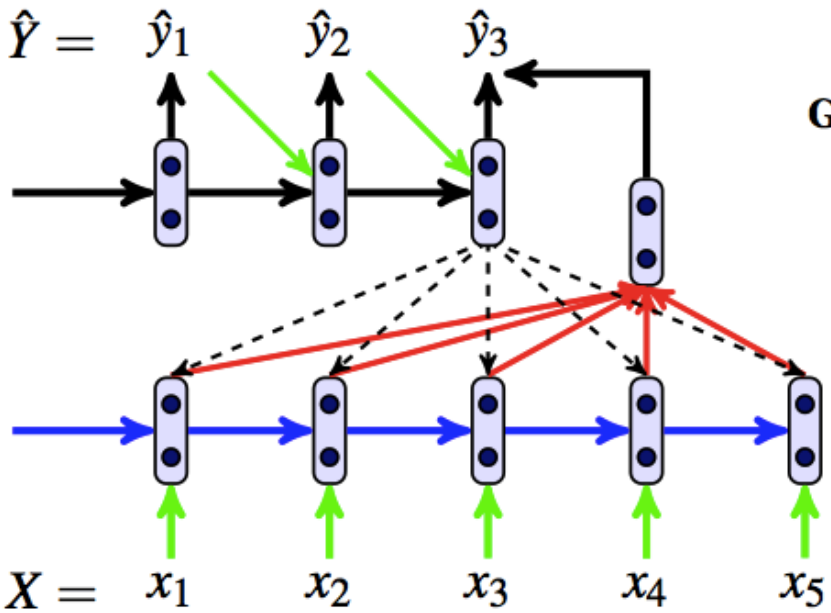

Get Action Seq. $\hat{Y}$

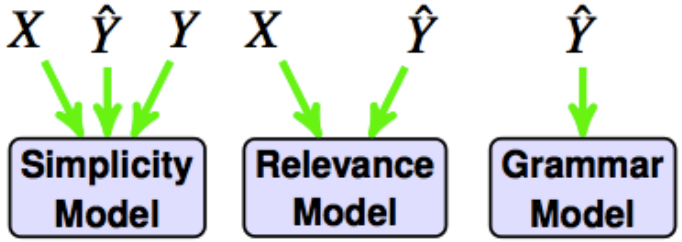

DRESS: **D**eep **RE**inforcement **S**entence **S**implification

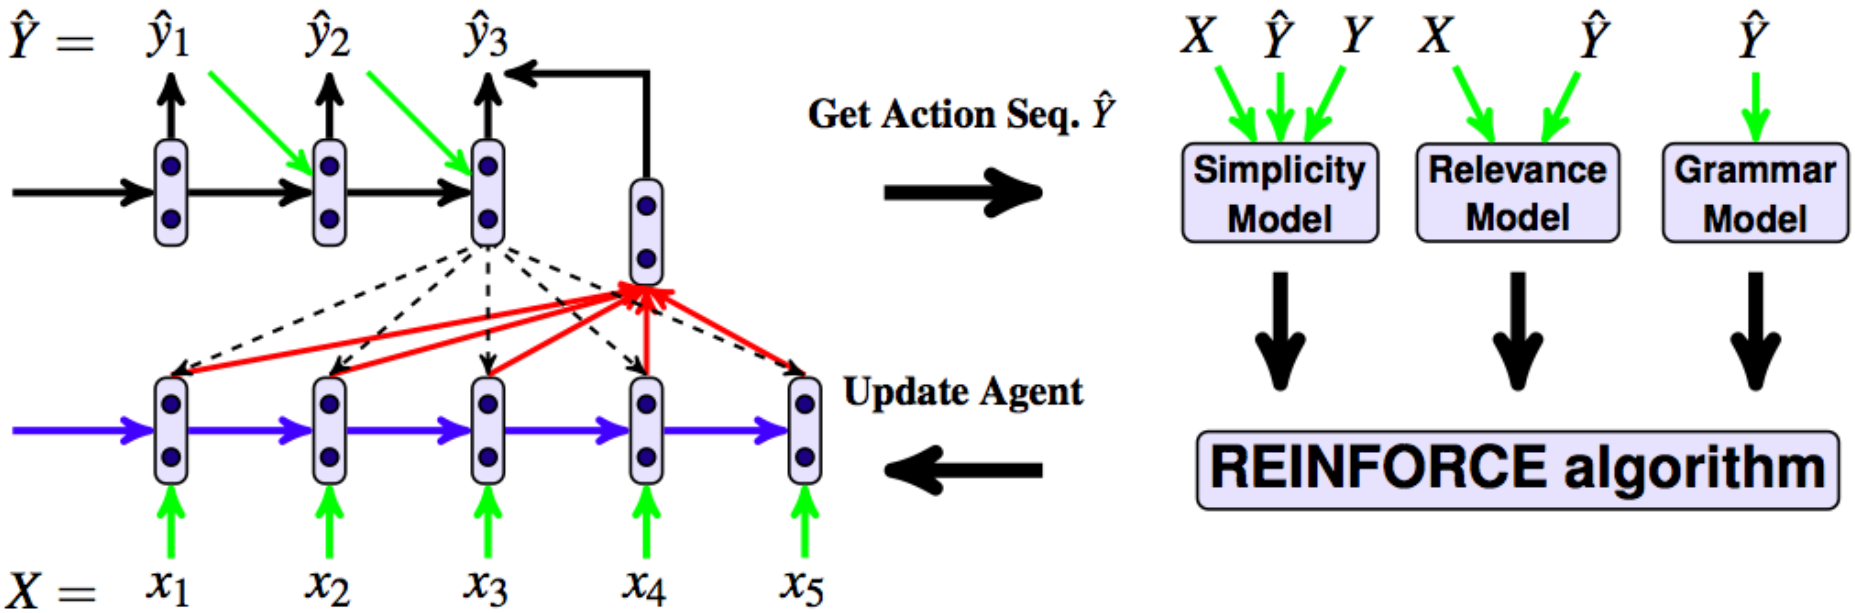

Rewards

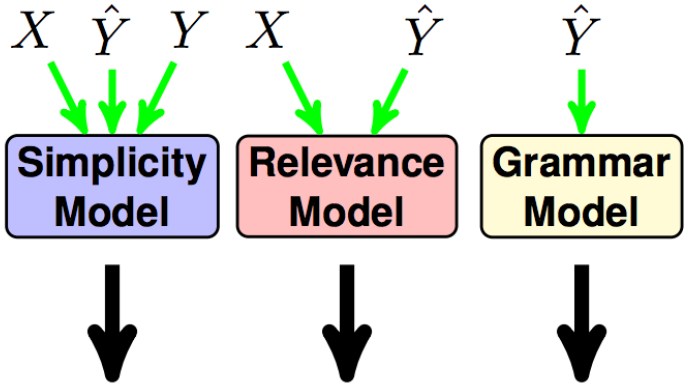

$$r(\hat{Y}) = \lambda^S r^S + \lambda^R r^R + \lambda^F r^F$$

Simplicity Reward

- SARI (Xu et al., 2016)
  - **S**ystem output **A**gainst **R**eferences and against the **I**nput sentence

$$r^S = \beta \, \textsc{Sari}(X, \hat{Y}, Y) + (1 - \beta) \, \textsc{Sari}(X, Y, \hat{Y})$$

Relevance Reward

$$r^R = \cos(\mathbf{q}_X, \mathbf{q}_{\hat{Y}}) = \frac{\mathbf{q}_X \cdot \mathbf{q}_{\hat{Y}}}{||\mathbf{q}_X|| \, ||\mathbf{q}_{\hat{Y}}||}$$

- $\mathbf{q}_X$ and $\mathbf{q}_{\hat{Y}}$ are learned by an LSTM sentence encoder.
- the LSTM encoder is trained within a Sequence Auto-Encoder Model (SAE; Dai and Le, 2015)

Fluency Reward

$$r^F = \exp\left(\frac{1}{|\hat{Y}|}\sum_{i=1}^{|\hat{Y}|}\log P_{LM}(\hat{y}_i|\hat{y}_{0:i-1})\right)$$

- Normalized sentence prob assigned by an LSTM LM
- Make sure $r^F \in [0, 1]$ as r^S and r^R

Lexical Simplification Model

- Word substitution is important for simplification (Specia et al., 2012); but Seq2SeqA is not perfect at this.
- Attention score α_t as soft word alignment (Luong et al., 2015)
- Lexical Simplification Model $P_{LS}(y_t|X, \alpha_t)$

Lexical Simplification Model

- Word substitution is important for simplification (Specia et al., 2012); but Seq2SeqA is not perfect at this.
- Attention score α_t as soft word alignment (Luong et al., 2015)
- Lexical Simplification Model $P_{LS}(y_t|X, \alpha_t)$

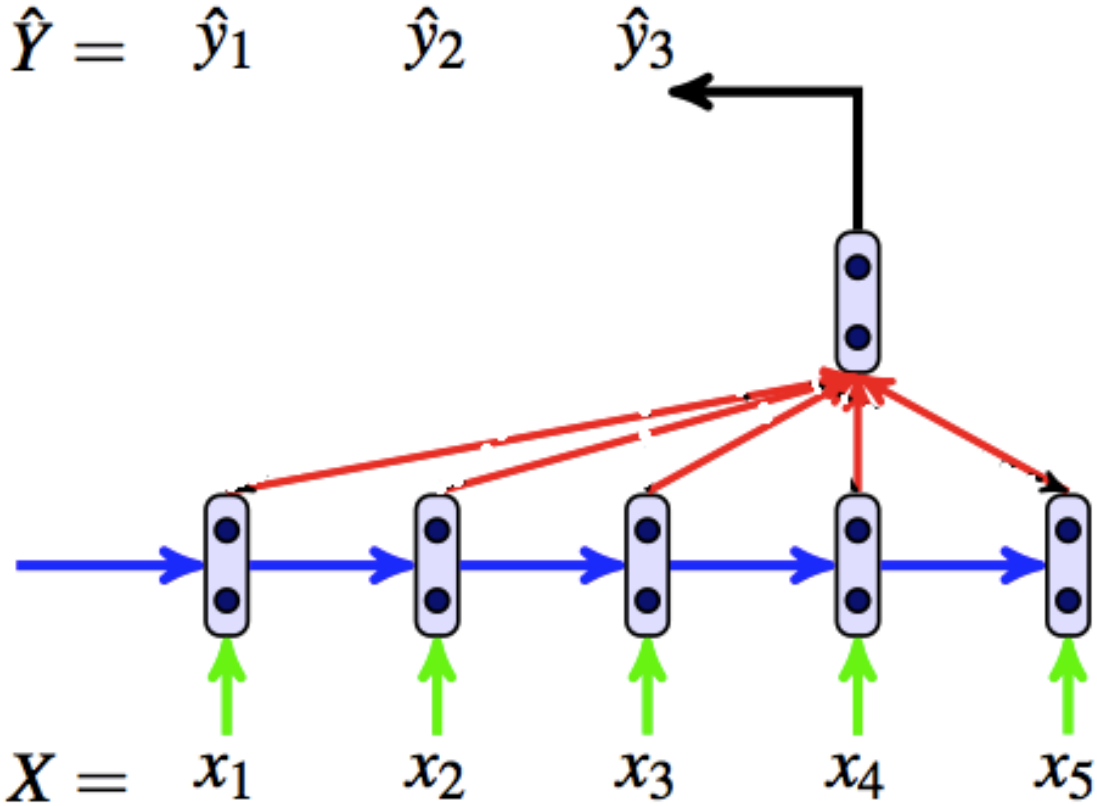

DRESS-LS: Put Things Together

$$P(y_t|y_{1:t-1}, X) = (1-\eta)\, P_{RL}(y_t|y_{1:t-1}, X) + \eta\, P_{LS}(y_t|X, \alpha_t)$$

Experiments

Datasets

Dataset	Train	Dev	Test
Newsela	94,208	1,129	1,076
WikiSmall	89,042	205	100
WikiLarge	296,402	2,000	359

- **Newsela**: Xu et al., 2015
- **WikiSmall**: Zhu et al., 2010
- **WikiLarge**:
  - Train: Kauchak, 2013; Woodsend and Lapata, 2011; Zhu et al., 2010
  - Dev and Test: Xu et al., 2016

Automatic Evaluation: Newsela Dataset

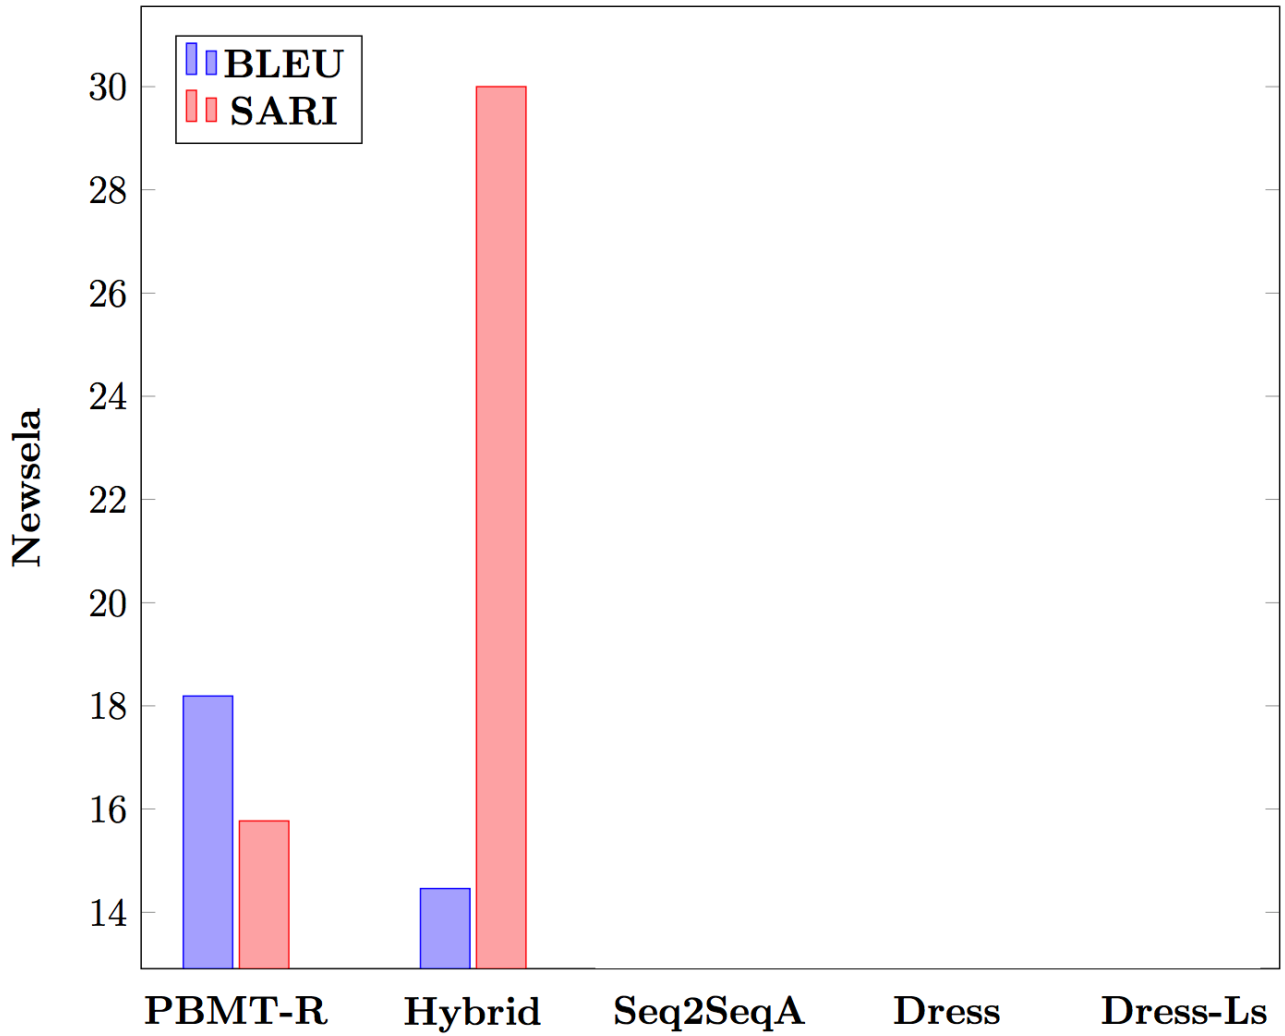

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Automatic Evaluation: Newsela Dataset

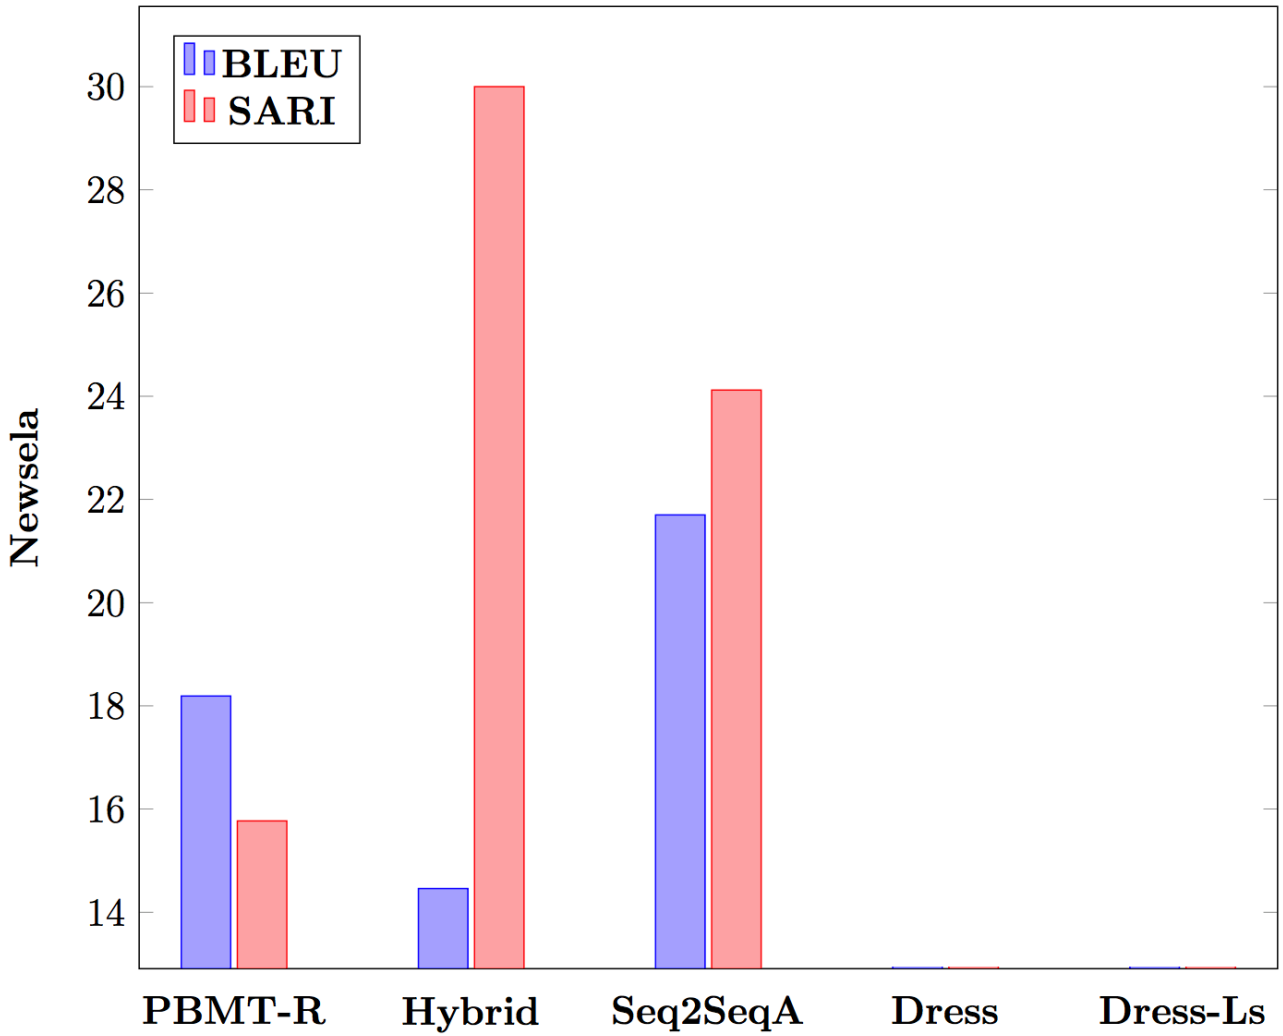

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Automatic Evaluation: Newsela Dataset

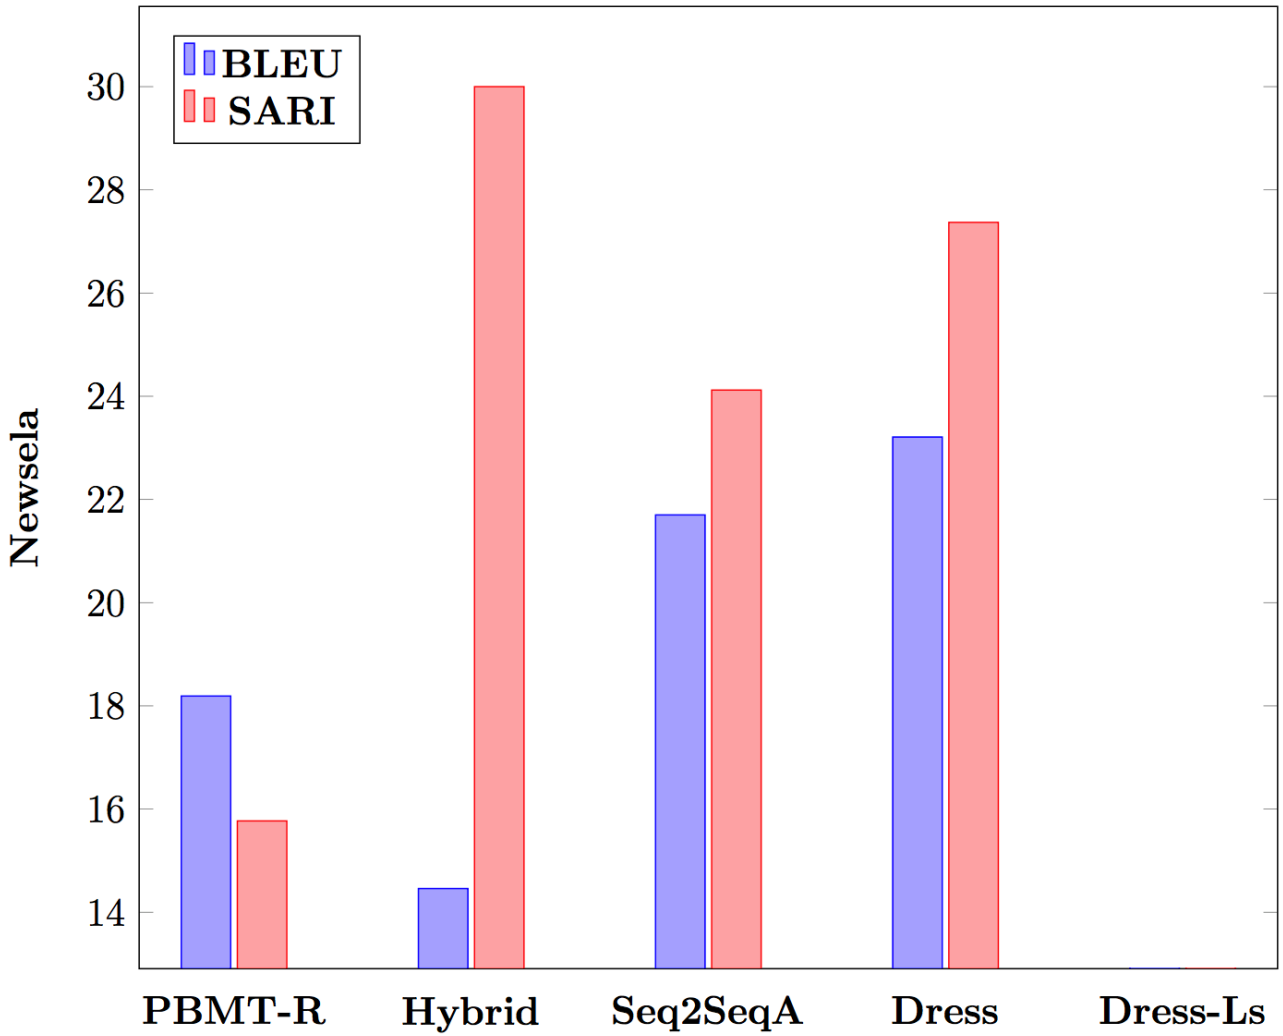

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Automatic Evaluation: Newsela Dataset

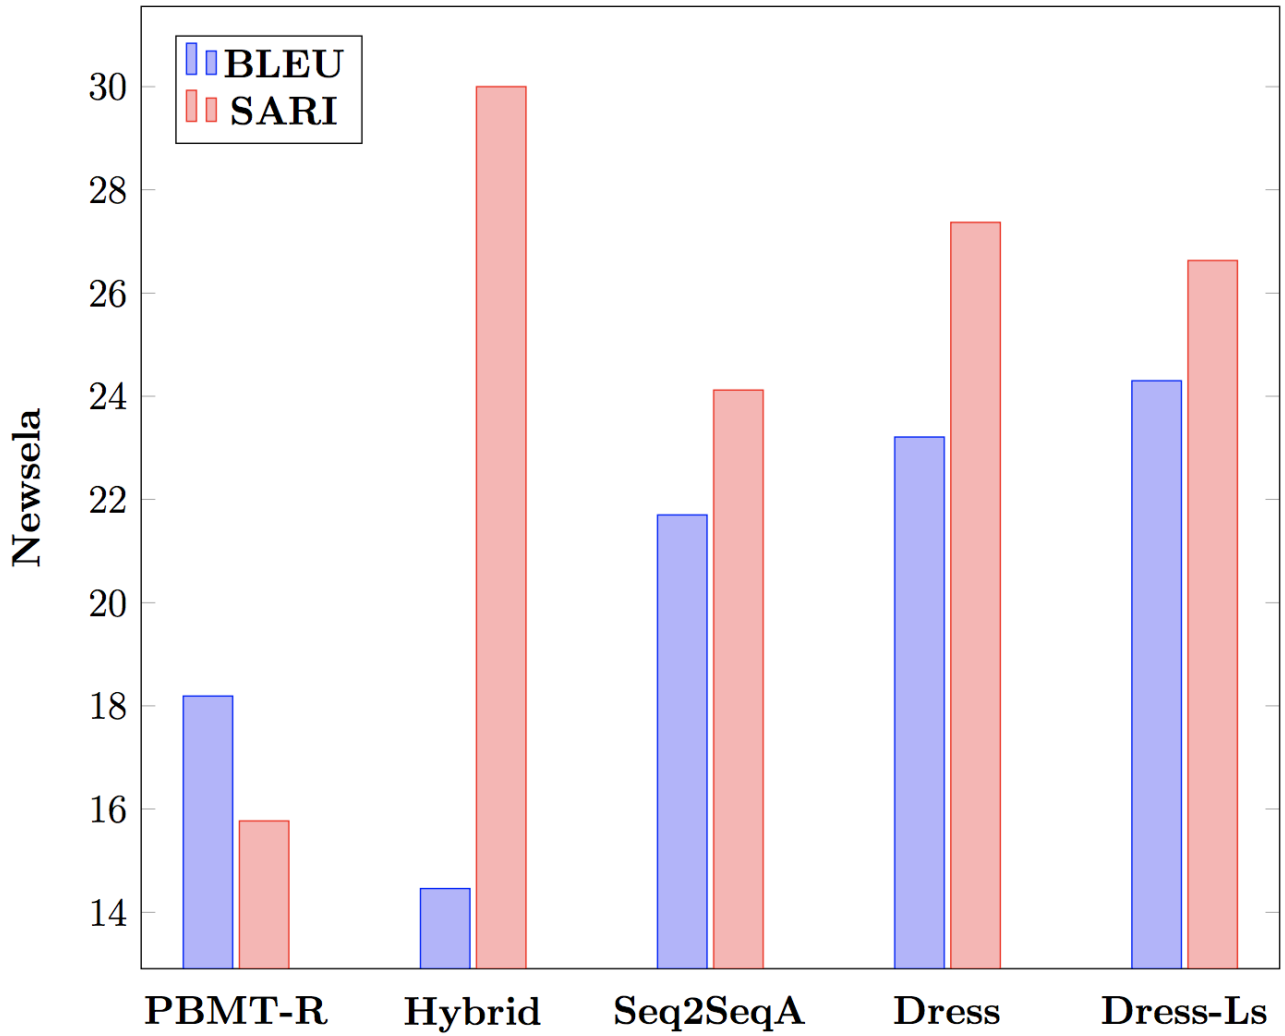

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Human Evaluation: Newsela Dataset

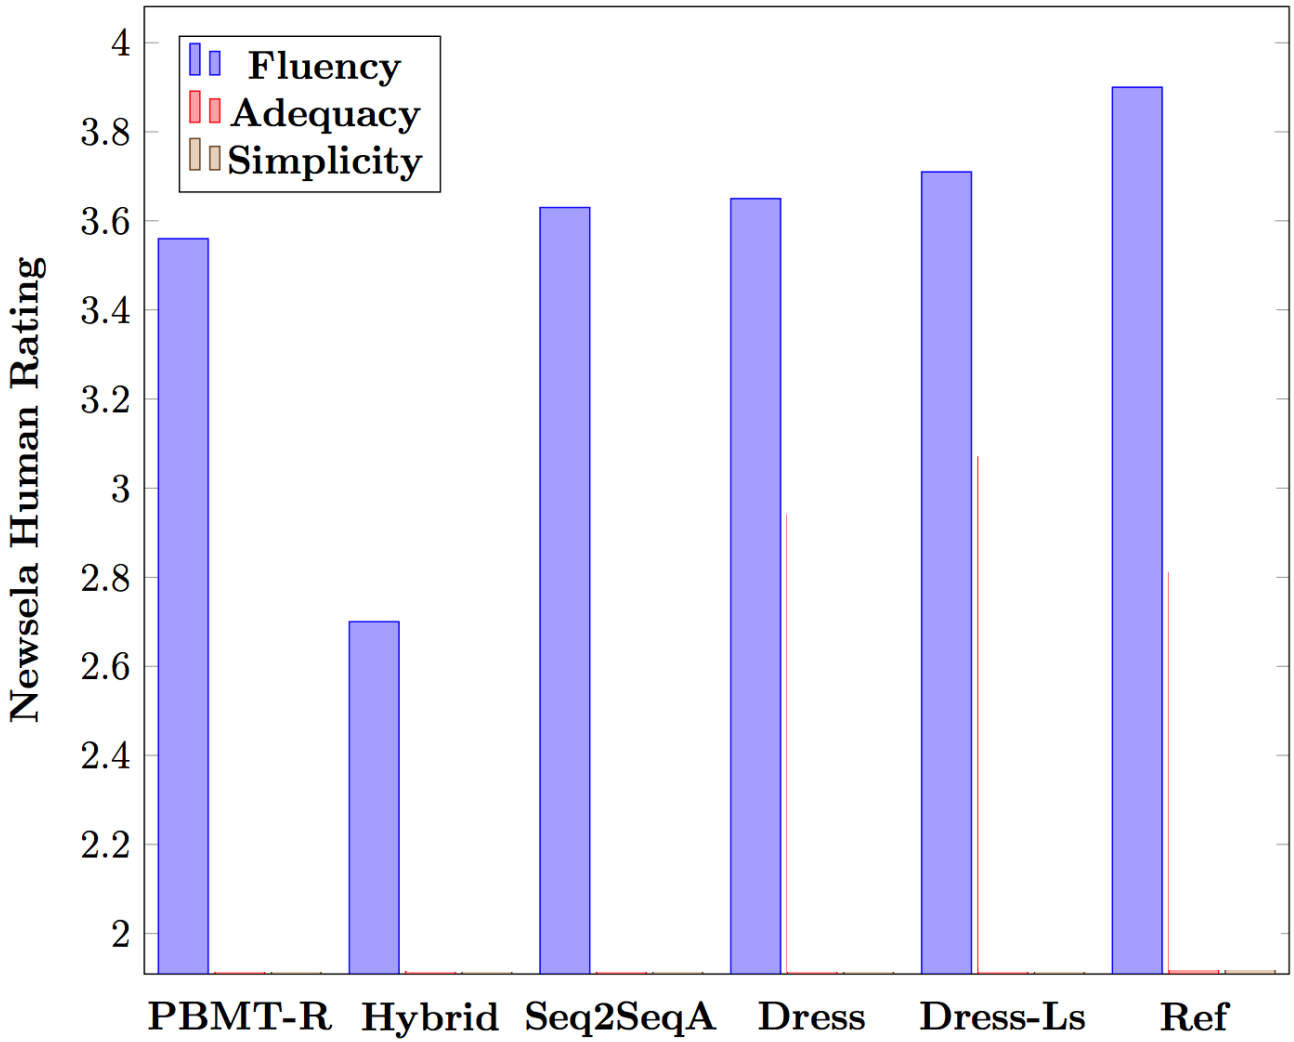

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Human Evaluation: Newsela Dataset

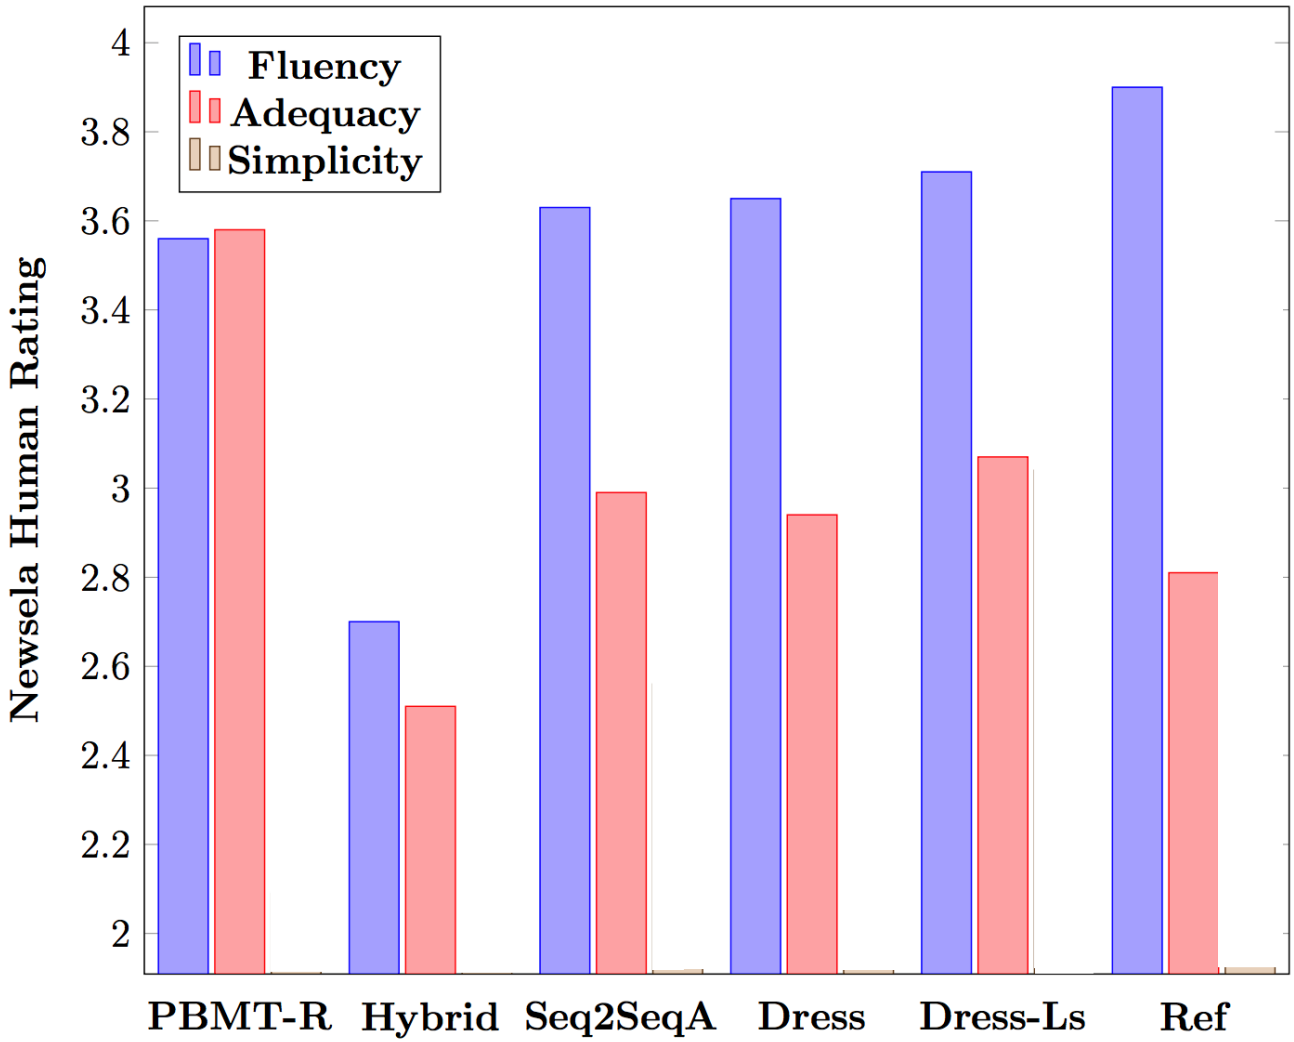

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Human Evaluation: Newsela Dataset

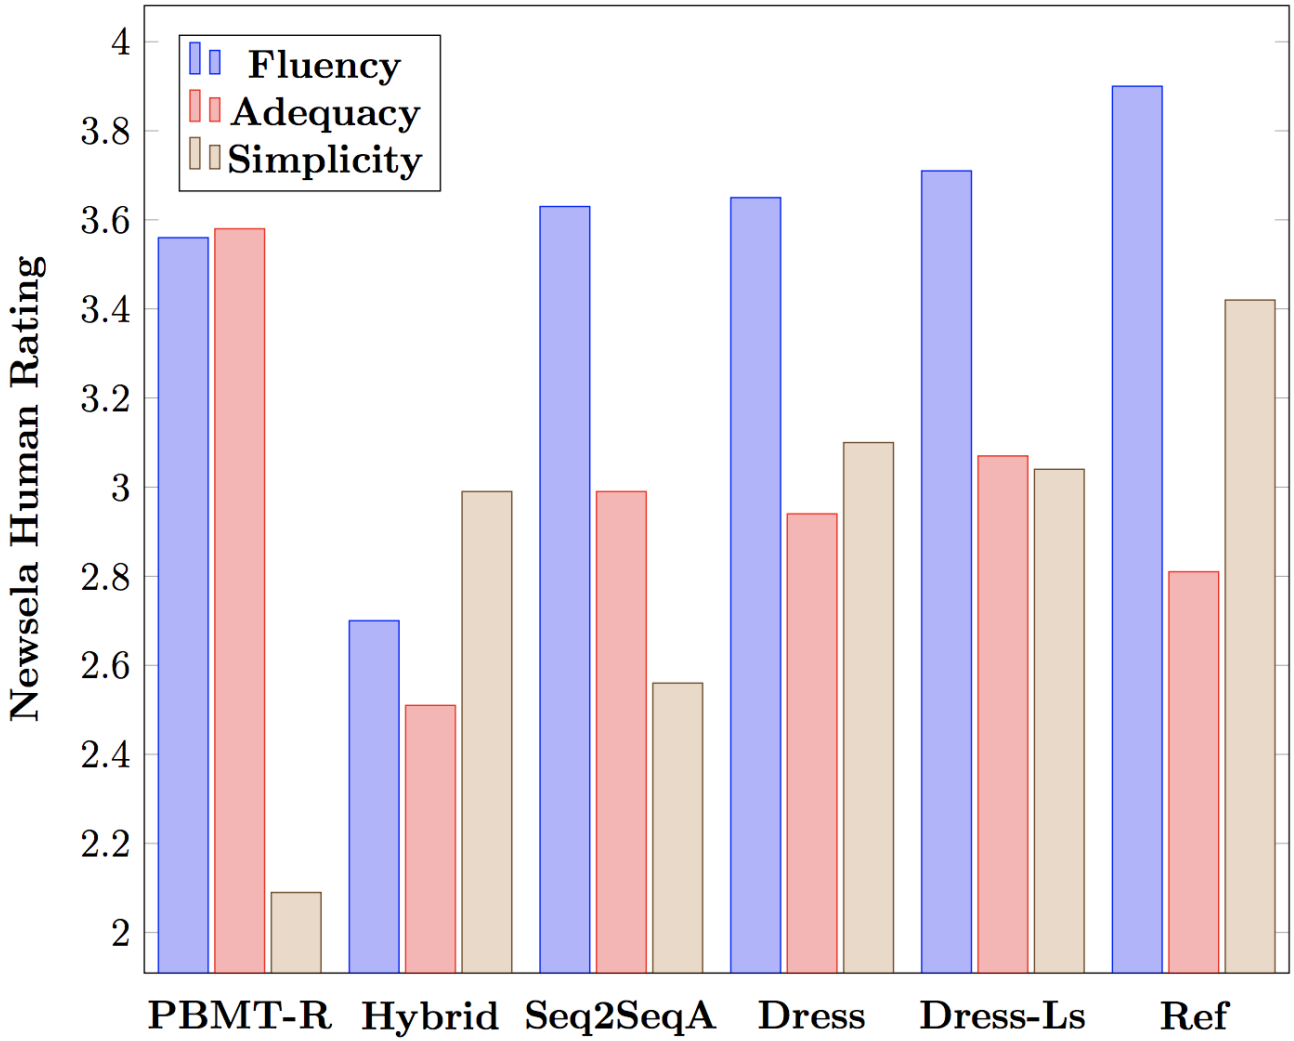

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Human Evaluation: Newsela Dataset

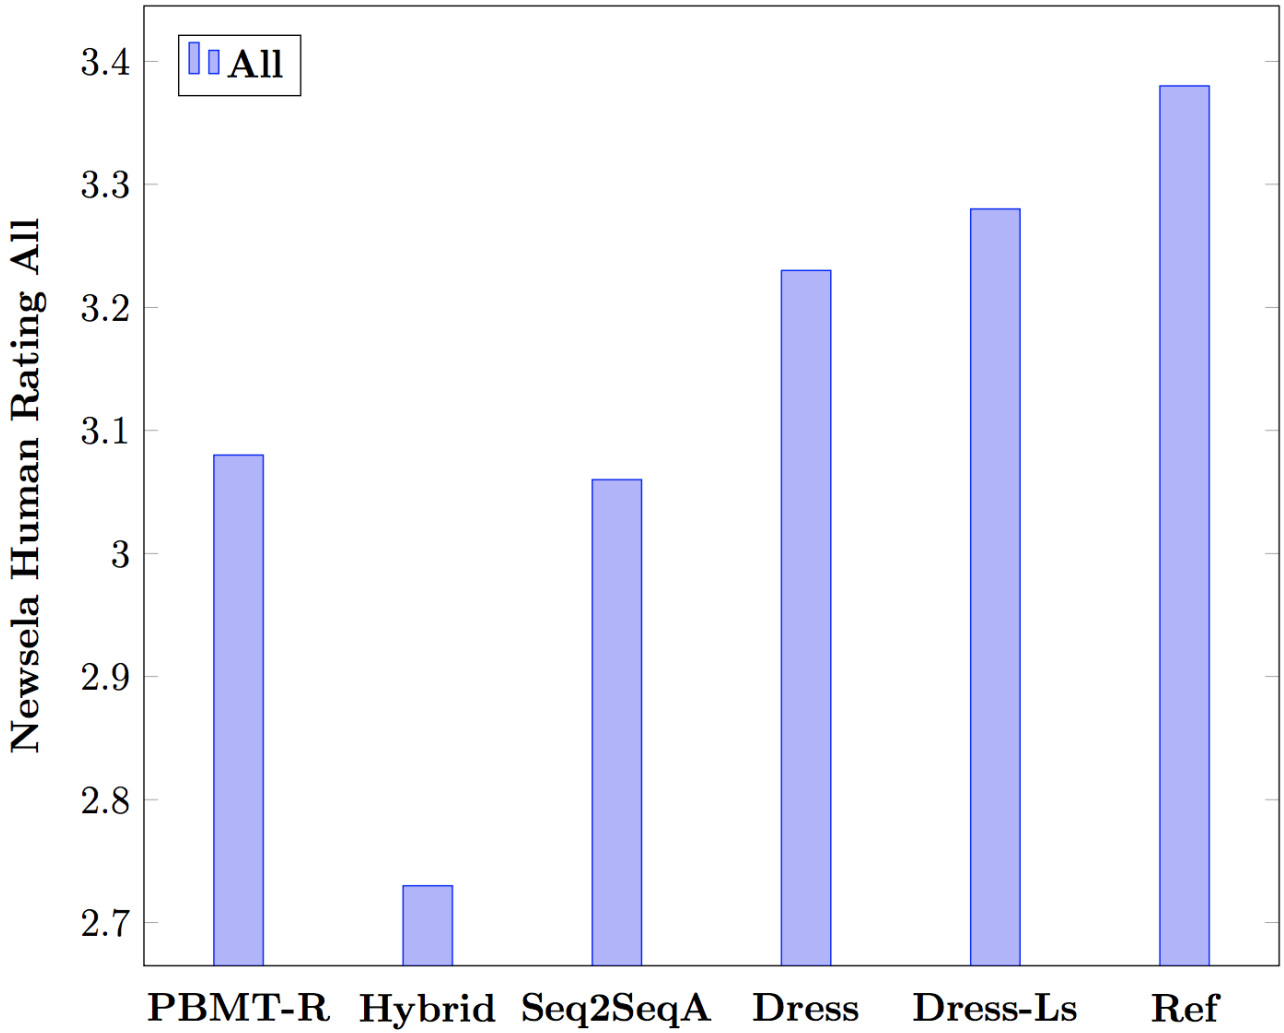

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Rewriting Operations Performed (Newsela)

Translation Edit Rate (TER; Snover et al., 2006)

Models	Len	TER	Ins	Del	Sub	Shft
PBMT-R	23.1	0.13	0.68	0.68	1.50	0.09
Hybrid	12.4	0.90	0.01	10.19	0.12	0.41
EncDecA	17.0	0.36	0.13	5.96	1.69	0.09
DRESS	14.2	0.46	0.07	8.53	1.37	0.11
DRESS-LS	14.4	0.44	0.07	8.38	1.11	0.09
Reference	12.7	0.67	0.40	10.26	3.44	0.73

EncDecA == Seq2Seq with Attention model; PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Rewriting Operations Performed (Newsela)

Translation Edit Rate (TER; Snover et al., 2006)

Models	Len	TER	Ins	Del	Sub	Shft
PBMT-R	23.1	0.13	0.68	0.68	1.50	0.09
Hybrid	12.4	0.90	0.01	10.19	0.12	0.41
EncDecA	17.0	0.36	0.13	5.96	1.69	0.09
DRESS	14.2	0.46	0.07	8.53	1.37	0.11
DRESS-LS	14.4	0.44	0.07	8.38	1.11	0.09
Reference	12.7	0.67	0.40	10.26	3.44	0.73

EncDecA == Seq2Seq with Attention model; PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Rewriting Operations Performed (Newsela)

Translation Edit Rate (TER; Snover et al., 2006)

Models	Len	TER	Ins	Del	Sub	Shft
PBMT-R	23.1	0.13	0.68	0.68	1.50	0.09
Hybrid	12.4	0.90	0.01	10.19	0.12	0.41
EncDecA	17.0	0.36	0.13	5.96	1.69	0.09
DRESS	14.2	0.46	0.07	8.53	1.37	0.11
DRESS-LS	14.4	0.44	0.07	8.38	1.11	0.09
Reference	12.7	0.67	0.40	10.26	3.44	0.73

EncDecA == Seq2Seq with Attention model; PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Human Evaluation: WikiSmall Dataset

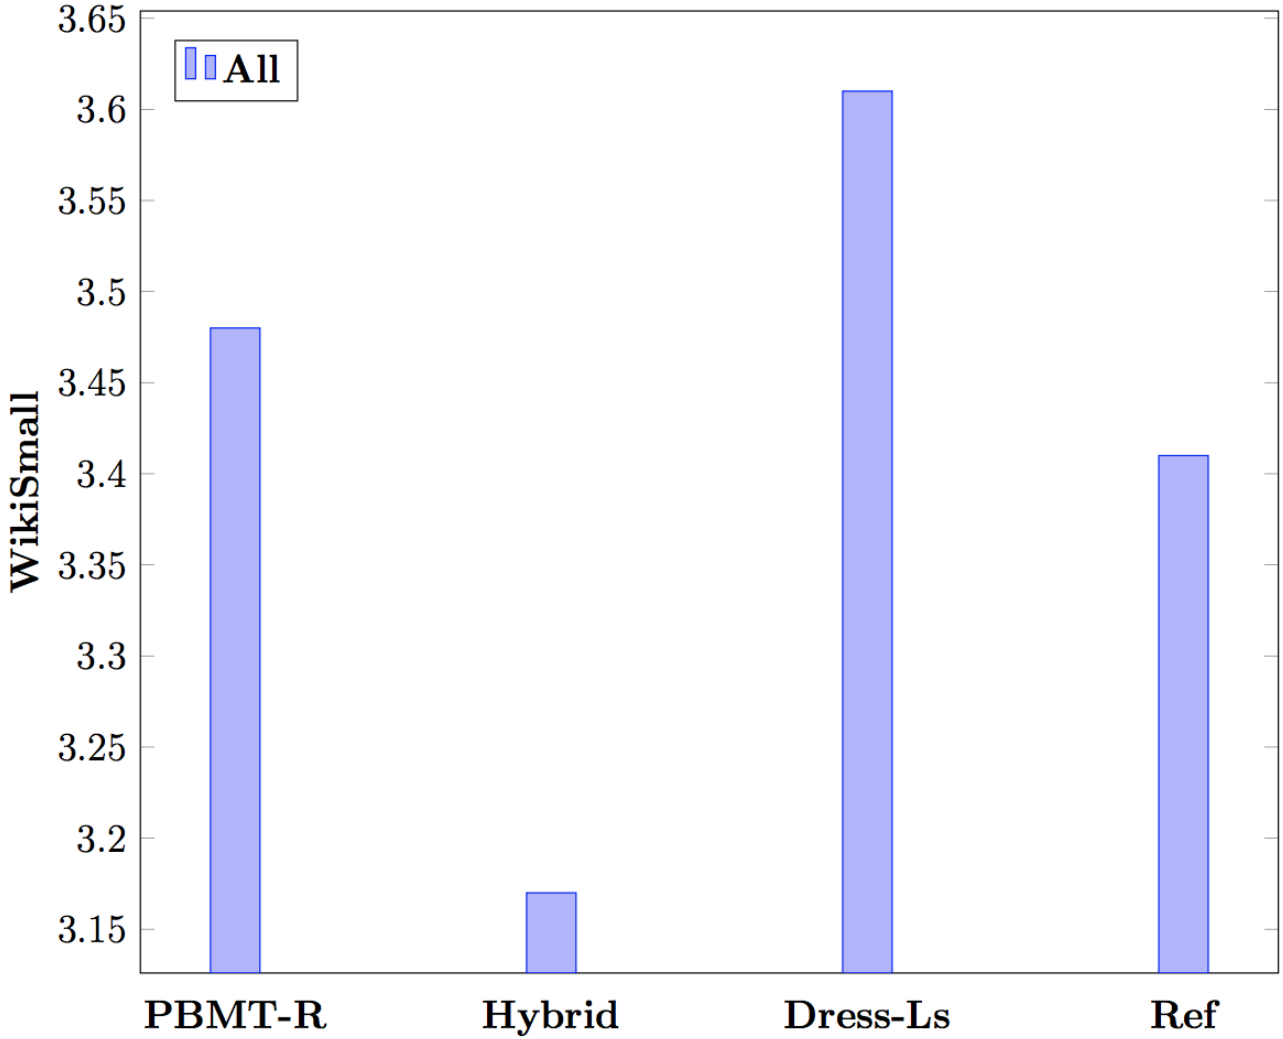

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014)

Human Evaluation: WikiLarge Dataset

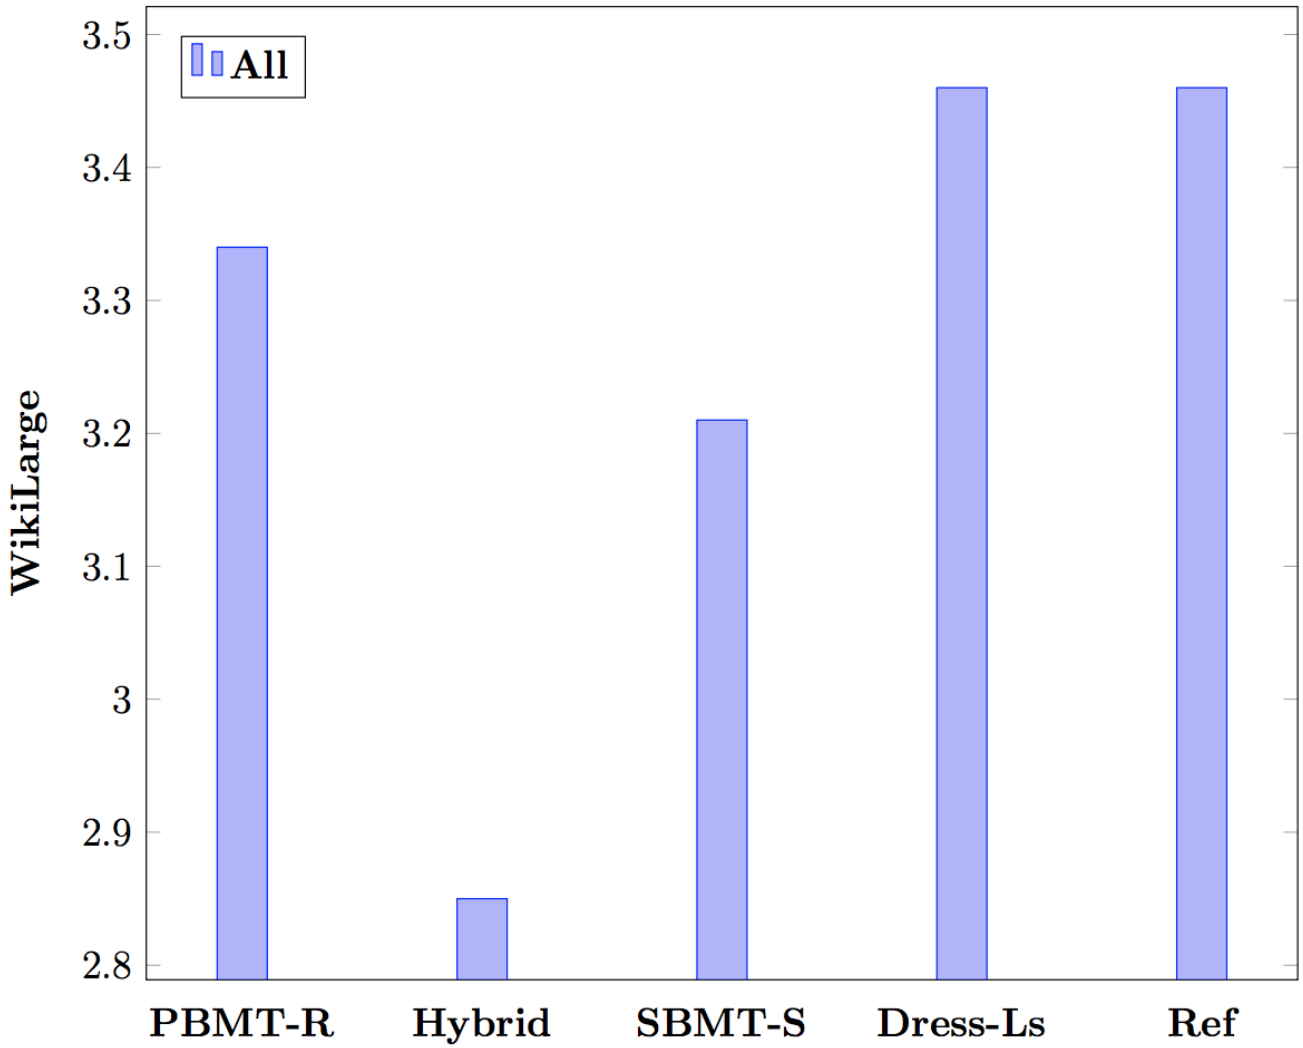

PBMT-R (Wubben et al., 2012), Hybrid (Narayan and Gardent, 2014), SBMT-S (Xu et al., 2016)

Conclusions

- We apply Reinforcement Learning to inject our prior knowledge into the sentence simplification task
- We also proposed a lexical simplification model to further improve the performance
- Future Work:
  - Model sentence splitting in our framework
  - Document level simplification
- Code Available:
  - https://github.com/XingxingZhang/dress