

Entailment Relation Aware Paraphrase Generation

Abhilasha Sancheti,^{1,2} Balaji Vasan Srinivasan,² Rachel Rudinger¹

¹University of Maryland, College Park

²Adobe Research

sancheti@umd.edu, balsrini@adobe.com, rudinger@umd.edu

Abstract

We introduce a new task of *entailment relation aware paraphrase generation* which aims at generating a paraphrase conforming to a given entailment relation (e.g., equivalent, forward entailing, or reverse entailing) with respect to a given input. We propose a reinforcement learning-based weakly-supervised paraphrasing system, **ERAP**, that can be trained using existing paraphrase and natural language inference (NLI) corpora without an explicit task-specific corpus. A combination of automated and human evaluations show that ERAP generates paraphrases conforming to the specified entailment relation and are of good quality as compared to the baselines and uncontrolled paraphrasing systems. Using ERAP for augmenting training data for downstream textual entailment task improves performance over an uncontrolled paraphrasing system, and introduces fewer training artifacts, indicating the benefit of explicit control during paraphrasing.

1 Introduction

Paraphrase is “an alternative surface form in the same language expressing the same semantic content as the original form” (Madnani and Dorr 2010). Although the logical definition of paraphrase requires strict semantic equivalence (or bi-directional entailment (Androustopoulos and Malakasiotis 2010)) between a sequence and its paraphrase, data-driven paraphrasing accepts a broader definition of *approximate* semantic equivalence (Bhagat and Hovy 2013). Moreover, existing automatically curated paraphrase resources do not align with this logical definition. For instance, pivot-based paraphrasing rules extracted by Ganitkevitch, Van Durme, and Callison-Burch (2013) contain hypernym or hyponym pairs, e.g., due to variation in the discourse structure of translations, and unrelated pairs, e.g., due to misalignments or polysemy in the foreign language.

While this flexibility of *approximate* semantic equivalence allows for greater diversity in expressing a sequence, it comes at the cost of the ability to precisely control the semantic entailment relationship (henceforth “entailment relation”) between a sequence and its paraphrase. This trade-off severely limits the applicability of paraphrasing systems or resources to a variety of downstream natural language understanding (NLU) tasks (e.g., machine translation, question

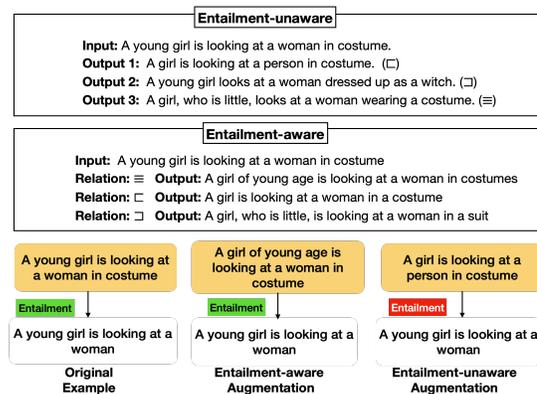


Figure 1: Entailment-unaware system might output *approximately* equivalent paraphrases. Label preserving augmentations generated using such system for textual entailment task can result in incorrect labels (red). Explicit entailment relation control in entailment-aware system helps in reducing such incorrectly labeled augmentations (green).

answering, information retrieval, and natural language inferencing (Pavlick et al. 2015)) (Figure 1). For instance, semantic divergences in machine translation have been shown to degrade the translation performance (Carpuat, Vyas, and Niu 2017; Pham et al. 2018).

Existing works identify directionality (forward, reverse, bi-directional, or no implication) of paraphrase and inference rules (Bhagat, Pantel, and Hovy 2007), and add semantics (natural logic entailment relationships such as equivalence, forward or reverse entailment, etc.) to data-driven paraphrasing resources (Pavlick et al. 2015) leading to improvements in lexical expansion and proof-based RTE systems, respectively. However, entailment relation control in paraphrase generation is, to our knowledge, a relatively unexplored topic, despite its potential benefit to downstream applications (Madnani and Dorr 2010) such as Multi-Document Summarization (MDS) (or Information Retrieval (IR)) wherein having such a control could allow the MDS (or IR) system to choose either the more specific (reverse entailing) or general (forward entailing) sentence (or query) depending on the purpose of the summary (or user needs).

To address the lack of entailment relation control in para-

phrasing systems, we introduce a new task of **entailment relation aware paraphrase generation**: given a sequence and an entailment relation, generate a paraphrase which conforms to the given entailment relation. We consider three entailment relations (controls) in the spirit of monotonicity calculus (Valencia 1991): (1) **Equivalence** (\equiv) refers to semantically equivalent paraphrases (e.g., synonyms) where input sequence entails its paraphrase and vice-versa; (2) **Forward Entailment** (\sqsubset) refers to paraphrases that lose information from the input or generalize it (e.g., hypernyms) i.e. input sequence entails its paraphrase; (3) **Reverse Entailment** (\supset) refers to paraphrases that add information to the input or makes it specific (e.g., hyponyms) i.e. input sequence is entailed by its paraphrase.

The unavailability of paraphrase pairs annotated with such a relation makes it infeasible to directly train a sequence-to-sequence model for this task. Collecting such annotations for existing large paraphrase corpora such as ParaBank (Hu et al. 2019b) or ParaNMT (Wieting and Gimpel 2018) is expensive due to scale. We address this challenge in 3 ways: (1) by building a novel *entailment relation oracle* based on natural language inference task (NLI) (Bowman et al. 2015a; Williams, Nangia, and Bowman 2018) to obtain weak-supervision for entailment relation for existing paraphrase corpora; (2) by recasting an existing NLI dataset, SICK (Marelli et al. 2014), into a small supervised dataset for this task, and (3) by proposing **Entailment Relation Aware Paraphraser (ERAP)** which is a reinforcement learning based (RL-based) weakly-supervised system that can be trained only using existing paraphrase and NLI corpora, with or without weak-supervision for entailment relation.

Intrinsic and extrinsic evaluations show advantage of entailment relation aware (henceforth “entailment-aware”) paraphrasing systems over entailment-unaware (standard uncontrolled paraphrase generation) counterparts. Intrinsic evaluation of ERAP (via a combination of automatic and human measures) on recasted SICK (§3) dataset shows that generated paraphrases conform to the given entailment relation with high accuracy while maintaining good or improved paraphrase quality when compared against entailment-unaware baselines. Extrinsic data-augmentation experiments (§5) on textual entailment task show that augmenting training sets using entailment-aware paraphrasing system leads to improved performance over entailment-unaware paraphrasing system, and makes it less susceptible to making incorrect predictions on adversarial examples.

2 Entailment Relation Aware Paraphraser

Task Definition. Given a sequence of tokens $X = [x_1, \dots, x_n]$, and an entailment relation $\mathcal{R} \in \{\text{Equivalence } (\equiv), \text{Forward Entailment } (\sqsubset), \text{Reverse Entailment } (\supset)\}$, we generate a paraphrase $Y = [y_1, \dots, y_m]$ such that the entailment relationship between X and Y is \mathcal{R} . \hat{Y} is the generated paraphrase and Y is the reference paraphrase.

Neural paraphrasing systems (Prakash et al. 2016; Li et al. 2018) employ a supervised sequence-to-sequence model to generate paraphrases. However, building a supervised model for this task requires paraphrase pairs with entailment rela-

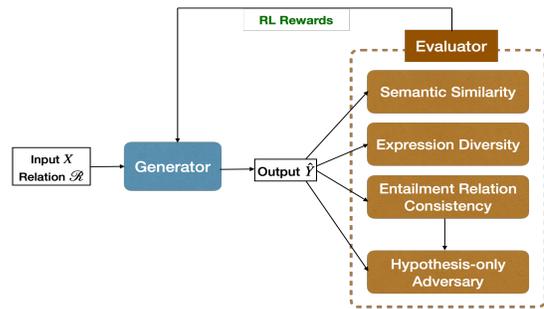


Figure 2: ERAP: Generator takes in a sequence X , an entailment relation \mathcal{R} , and outputs a paraphrase \hat{Y} . \hat{Y} is scored by various scorers in the evaluator and a combined score (known as reward) is sent back to train the generator. Hypothesis-only adversary is adversarially trained on \hat{Y} and predictions from the entailment relation consistency scorer.

tion annotations. To address this, we propose an RL-based paraphrasing system ERAP which can be trained with existing paraphrase and NLI corpora without any additional annotations. **ERAP** (Figure 2) consists of a paraphrase generator (§2.1) and an evaluator (§2.2) comprising of various scorers to assess the quality of generated paraphrases for different aspects. Scores from the evaluator are combined (§2.3) to provide feedback to the generator in the form of rewards. Employing RL allows us to explicitly optimize the generator over measures accounting for the quality of generated paraphrases, including the non-differentiable ones.

2.1 Paraphrase Generator

The generator is a transformer-based (Vaswani et al. 2017) sequence-to-sequence model which takes (X, \mathcal{R}) and generates \hat{Y} . We denote the generator as $G(\hat{Y}|X, \mathcal{R}; \theta_g)$, where θ_g refers to parameters of the generator. We incorporate the entailment relation as a special token prepended to the input sequence. This way, entailment relation receives special treatment (Kobus, Crego, and Senellart 2017) and the generator learns to generate paraphrases for a given X , and \mathcal{R} .

2.2 Paraphrase Evaluator

The evaluator comprises of several scorers to assess the quality of the generated paraphrase for three aspects: *semantic similarity* with the input, *expression diversity* from the input, and *entailment relation consistency*. It provides rewards to the paraphrases generated by the generator as feedback which is used to update the parameters of the generator. We describe the various scorers below.

Semantic Similarity Scorer provides reward which encourages the generated paraphrase \hat{Y} to have similar meaning as the input sequence X . We use MoverScore (Zhao et al. 2019) to measure the semantic similarity between the generated paraphrase and the input, denoted as $r_s(X, \hat{Y})$. MoverScore combines contextualized representations with word mover’s distance (Kusner et al. 2015) and has shown high correlation with human judgment of text quality.

Expression Diversity Scorer rewards the generated paraphrase to ensure that it uses different tokens or surface form to express the input. We measure this aspect by computing n-grams dissimilarity (inverse BLEU (Papineni et al. 2002)),

$$r_d(X, \hat{Y}) = 1 - \text{BLEU}(\hat{Y}, X) \quad (1)$$

Following Hu et al. (2019b), we use modified BLEU without length penalty to avoid generating short paraphrases which can result in high inverse BLEU scores.

Entailment Relation Consistency Scorer is a novel scorer designed to reward the generated paraphrase in such a way that encourages it to adhere to the given entailment relation \mathcal{R} . To compute the reward, we build an oracle $O(X, Y)$ (details in §3) based on natural language inferencing (NLI) and use likelihood of the given entailment relation from the Oracle as the score. $r_l(X, \hat{Y}, \mathcal{R}) = O(l = \mathcal{R} | X, \hat{Y})$. As will be discussed further in §4.3, we found that entailment relation consistency scorer can result in generator learning simple heuristics (e.g., adding same adjective such as ‘desert’, or trailing tokens like ‘and says’ or ‘with mexico’ for \sqsupset or short outputs for \sqsubset) leading to degenerate paraphrases having high consistency score.

Inspired by the idea of *hypothesis-only* baselines (Poliak et al. 2018) for NLI task, we build a novel RoBERTa-based **Hypothesis-only Adversary**, $A(l | \hat{Y})$, to penalize the generated paraphrases resorting to such heuristics. The adversary is a 3-class classifier trained on the paraphrases generated during the training phase with the oracle prediction for (X, \hat{Y}) pair as the ground-truth. The Adversary loss $L(A)$ is

$$L(A) = - \sum_{c=1}^{|C|} O(X, \hat{Y}) \log(A(l = c | \hat{Y})), \quad (2)$$

where $|C| = 3$ is the number of entailment relations. Training the adversary in this way helps in adapting to the heuristics taken by the generator during the course of training. The generator and the adversary are trained alternatively, similar to a GAN (Goodfellow et al. 2014) setup. The penalty is computed as the likelihood of entailment relation being \mathcal{R} using the adversary, $p_l(\hat{Y}, \mathcal{R}) = A(l = \mathcal{R} | \hat{Y})$. We only penalize those generated paraphrases for which predicted relation is same as the input relation because incorrect prediction denotes no heuristic is taken by the generator.

2.3 Reinforcement Learning Setup

The output paraphrases from the generator are sent to the scorers for evaluation. The various scores from the scorers are combined to give feedback (in the form of reward) to the generator to update its parameters and to improve the quality of the generated paraphrases conforming to the given relation. We emphasize that although the scores from our scorers are not differentiable with respect to θ_g , we can still use them by employing RL (the REINFORCE algorithm) to update the parameters of the generator (Williams 1992).

In RL paradigm, state at time t is defined as $s_t = (X, \mathcal{R}, \hat{Y}_{1:t-1})$ where $\hat{Y}_{1:t-1}$ refers to the first $t - 1$ tokens that are already generated in the paraphrase. The action at time t is the t^{th} token to be generated. Let V be the vocabulary, and T be the maximum output length. The total

expected reward of the current generator is then given by $J(G) = \sum_{t=1}^T \mathbb{E}_{\hat{Y}_{1:t-1} \sim G} [\sum_{y_t \in V} \mathbb{P}(y_t | s_t) Q(s_t, y_t)]$, where $\mathbb{P}(y_t | s_t)$ is the likelihood of token y_t given the current state s_t , and $Q(s_t, y_t)$ is the cumulative discounted reward for a paraphrase extended from $\hat{Y}_{1:t-1}$. The total reward, Q , is defined as the sum of the token level rewards.

$$Q(s_t, y_t) = \sum_{\tau=t}^T \gamma^{\tau-t} r(s_\tau, y_\tau), \quad (3)$$

where $r(s_\tau, y_\tau)$ is the reward of token y_τ at state s_τ , and $\gamma \in (0, 1)$ is a discounting factor so that the future rewards have decreasing weights, since their estimates are less accurate. If we consider that $\hat{Y}_{1:t-1}$ has been given then for every y_t , the total expected reward becomes

$$J(G) = \sum_{t=1}^T \sum_{y_t \in V} \mathbb{P}(y_t | s_t) Q(s_t, y_t). \quad (4)$$

Sequence Sampling. To obtain $r(s_t, y_t)$ at each time step t , we need scores for each token. However, by design these scorers only evaluate complete sequences instead of single token or partial sequences. We therefore use the technique of rolling out (Yu et al. 2017), where the generator ‘‘rolls out’’ a given sub-sequence $\hat{Y}_{1:t}$ to generate complete sequence by sampling the remaining part of the sequence $\hat{Y}_{t+1:T}$. Following Gong et al. (2019), we use a combination of beam search and multinomial sampling to balance reward estimation accuracy at each time step and diversity of the generated sequence. We first generate a reference paraphrase $\hat{Y}_{1:T}^{ref}$ using beam search and draw n samples of complete sequences $\hat{Y}_{1:T}$ by rolling out the sub-sequence $\hat{Y}_{1:t}^{ref}$ using multinomial sampling to estimate reward at each time step t .

Reward Estimation. We send n samples of complete sequences drawn from the sub-sequence $\hat{Y}_{1:t}^{ref}$ to the scorers. The combined score $f(s_t, y_t)$ for an action y_t at state s_t is computed by averaging the score of the complete sequences rolled out from $\hat{Y}_{1:t}^{ref}$ defined as

$$f(s_t, y_t) = \frac{1}{n} \sum_{i=1}^n \alpha \cdot (r_l(X, \hat{Y}^i, \mathcal{R}) - p_l(\hat{Y}^i, \mathcal{R})) + \beta \cdot r_s(X, \hat{Y}^i) + \delta \cdot r_d(X, \hat{Y}^i), \quad (5)$$

where α , β , δ , and n are hyperparameters empirically set to 0.4, 0.4, 0.2, and 2, respectively. These parameters control the trade-off between different aspects for this multi-objective task. Following Siddique, Oymak, and Hristidis (2020), we threshold¹ the scorers’ scores so that the final reward maintains a good balance across various scores. For example, generating diverse tokens at the expense of losing too much on the semantic similarity is not desirable. Similarly, copying the input sequence as-is to the generation is clearly not a paraphrase (i.e., $r_s(X, \hat{Y}) = 1$). We define reward $r(s_t, y_t)$ for action y_t at state s_t as:

$$r(s_t, y_t) = \begin{cases} f(s_t, y_t) - f(s_{t-1}, y_{t-1}), & t > 1, \\ f(s_1, y_1), & t = 1 \end{cases} \quad (6)$$

¹If $0.3 \leq r_s(X, \hat{Y}) \leq 0.98$ then the score is used as is otherwise, 0. Similarly, if $r_s(X, \hat{Y}) > 0$ after thresholding then r_d , r_l , and p_l are computed as defined, otherwise 0.

The discounted cumulative reward $Q(s_t, y_t)$ is then computed from the rewards $r(s_\tau, y_\tau)$ at each time step using Eq. 3 and the total expected reward is derived using Eq. 4. The generator loss $L(G)$ is defined as $-J(G)$.

2.4 Training Details

Pre-training has been shown to be critical for RL to work in unsupervised settings (Siddique, Oymak, and Hristidis 2020; Gong et al. 2019) therefore, we pre-train the generator on existing large paraphrase corpora *e.g.*, ParaBank (Hu et al. 2019b) or ParaNMT (Wieting and Gimpel 2018) in two ways: (1) **Entailment-aware** uses Oracle (§3) to obtain entailment relation for paraphrase pairs in the train set of paraphrase corpora, filter the semantically-divergent (§3) pairs, upsample or downsample to have balanced data across relations, and train the generator with weak-supervision for entailment relation and gold-paraphrases, and (2) **Entailment-unaware** trains the generator on paraphrase pairs as-is without any entailment relation. Pre-training is done in a supervised manner with the cross-entropy loss and offers immediate benefits for generator to learn paraphrasing transformations and have warm-start leading to faster model training.

RL-based Fine-tuning. We fine-tune the generator using feedback from the evaluator on recasted SICK dataset (details in §3). For any practical purposes, our RL-fine-tuning approach only requires input sequences without any annotations for entailment relation or ground-truth paraphrases. However, for a fair comparison against supervised or weakly-supervised baselines (§4.1), we use the gold-entailment relation for recasted SICK during RL fine-tuning.

3 Collecting Labeled Paraphrase Data

Entailment-aware paraphrasing requires paraphrase pairs annotated with entailment relation. However, collecting such annotations for large paraphrase corpora such as ParaBank² (Hu et al. 2019b) is too costly. To obtain entailment relations automatically, we train a NLI classifier and use it to derive the entailment relations as described below.

Entailment Relation Oracle. NLI is a standard natural language understanding task of determining whether a *hypothesis* h is true (entailment³ E), false (contradiction C), or undetermined (neutral N) given a *premise* p (MacCartney 2009). To build an entailment relation oracle, $O(X, Y)$, we first train a RoBERTa-based (Liu et al. 2019) 3-class classifier, $o(l|\langle p, h \rangle)$, to predict the uni-directional (E, N, C) labels given a $\langle p, h \rangle$ pair. This classifier is then run forwards ($\langle X, Y \rangle$) and backwards ($\langle Y, X \rangle$) on the paraphrase pairs to get the uni-directional predictions which are further used to derive the entailment relations as in Eq. 7. The Oracle is used to generate weak-supervision for entailment relations for existing paraphrase corpora, and assess the generated

²It consists of 50 million high quality English paraphrases obtained training a Czech-English neural machine translation (NMT) system and adding lexical-constraints to NMT decoding procedure.

³Entailment in NLI is a uni-directional relation while Equivalence is a bi-directional entailment relation.

Split	Recasted SICK				SICK NLI		
	\equiv	\sqsubset	\sqsupset	Others	E	N	C
Train	1344	684	684	420	1274	2524	641
Dev	196	63	63	43	143	281	71
Test	1386	814	814	494	1404	2790	712

Table 1: E, N, C denote entailment, neutral, and contradiction, respectively. Others refers to neutral or invalid relation.

paraphrases for relation consistency. We only focus on \equiv , \sqsubset , and \sqsupset relations as contradictory, neutral or invalid pairs are considered as semantically-divergent sentence pairs.

$$O(X, Y) = \begin{cases} \equiv & \text{if } o(l|\langle X, Y \rangle) = E \ \& \ o(l|\langle Y, X \rangle) = E \\ \sqsubset & \text{if } o(l|\langle X, Y \rangle) = E \ \& \ o(l|\langle Y, X \rangle) = N \\ \sqsupset & \text{if } o(l|\langle X, Y \rangle) = N \ \& \ o(l|\langle Y, X \rangle) = E \\ C & \text{if } o(l|\langle X, Y \rangle) = C \ \& \ o(l|\langle Y, X \rangle) = C \\ N & \text{if } o(l|\langle X, Y \rangle) = N \ \& \ o(l|\langle Y, X \rangle) = N \\ \text{Invalid} & \text{,otherwise} \end{cases} \quad (7)$$

Recasting SICK Dataset. SICK (Marelli et al. 2014) is a NLI dataset created from sentences describing the same picture or video which are near paraphrases. It consists of sentence pairs (p, h) with human-annotated NLI labels for both directions $\langle p, h \rangle$ and $\langle h, p \rangle$. We recast this dataset to obtain paraphrase pairs with entailment relation annotations derived using the gold bi-directional labels in the same way as O . We only consider the sentence pairs which were created by combining meaning-preserving transformations (details in appendix). We augment this data by adding valid samples obtained by reversing sentence pairs ($\forall p \sqsubset h$, we add $h \sqsupset p$ and $\forall p \equiv h$, we add $h \equiv p$). Data statistics in Table 1.

Oracle Evaluation. We train the NLI classifier o on existing NLI datasets namely, MNLI (Williams, Nangia, and Bowman 2018), SNLI (Bowman et al. 2015a), SICK (Marelli et al. 2014) as well as diagnostic datasets such as, HANS (McCoy, Pavlick, and Linzen 2019), others introduced in Glockner, Shwartz, and Goldberg (2018); Min et al. (2020), using cross-entropy loss. Combining diagnostic datasets during training has shown to improve robustness of NLI systems which can resort to simple lexical or syntactic heuristics (Glockner, Shwartz, and Goldberg 2018; Poliak et al. 2018) to perform well on the task. The accuracy of $o(l|\langle p, h \rangle)$ on the combined test sets of the datasets used for training is 92.32% and the accuracy of Entailment Relation Oracle $O(X, Y)$ on the test set of recasted SICK dataset is 81.55%. Before using the Oracle to obtain weak-supervision for entailment relation for training purposes, we validate it by manually annotating 50 random samples from ParaBank. 78% of the annotated relations were same as the Oracle predictions when C, N, and Invalid labels were combined.

4 Intrinsic Evaluation

Here we provide details on the entailment-aware and unaware comparison models, and the evaluation measures.

4.1 Comparison Models

To contextualize ERAP’s performance, we train several related models including supervised and weakly-supervised,

Aware	BLEU↑	Div↑	iBLEU↑	\mathcal{R} -Con↑
✗	32.54	46.57	17.78	–
✓	33.08	58.24	19.06	72.34

Table 2: Evaluation of the generator pre-trained on ParaBank using entailment-aware (✓) and unaware (✗) settings.

entailment-aware and unaware models to obtain lower and upper bound performance on recasted SICK as follows: (1) the generator is trained on recasted SICK in an entailment-aware (**S2S-A**) and unaware (**S2S-U**) supervised setting; (2) the generator is pre-trained on ParaBank dataset in entailment-aware (**Pre-train-A**) and unaware (**Pre-train-U**) setting to directly test on the test set of recasted SICK; (3) the pre-trained generators are fine-tuned on recasted SICK in entailment-aware (**Fine-tune-A**) and unaware (**Fine-tune-U**) supervised setting; (4) multiple outputs ($k \in \{1, 5, 10, 20\}$) are sampled using nucleus sampling (Holtzman et al. 2019) from S2S-U (**RR-S2S-U**) or Fine-tune-U (**RR-FT-U**) and re-ranked based on the combined score $f(s_t, y_t)$. The highest scoring output is considered as the final output for RR-S2S-U and RR-FT-U.

4.2 Evaluation Measures

Automatic evaluation to evaluate the quality of paraphrases is primarily done using *i*BLEU (Sun and Zhou 2012) which penalizes for copying from the input. Following Liu et al. (2020), we also report BLEU (Papineni et al. 2002) (up to 4 n-grams) and Diversity (Div, measured identical to Eq. 1) scores to understand the trade-off between these measures. We also compute, \mathcal{R} -Con, defined as the percentage of test examples for which the entailment relation predicted using oracle is same as the given entailment relation.

Human evaluation is conducted on 4 aspects: (1) **semantic similarity** which measures the closeness in meaning between paraphrase and input on a scale of 5 (Li et al. 2018); (2) **diversity in expression** which measures if different tokens or surface-forms are used in the paraphrase with respect to the input on a scale of 5 (Siddique, Oymak, and Hristidis 2020); (3) **grammaticality** which measures if paraphrase is well-formed and comprehensible on a scale of 5 (Li et al. 2018); (4) **relation consistency** which measures the % of examples for which the annotated entailment relation is same as the input relation. Three annotations per sample are collected for similarity, diversity, and grammaticality using Amazon Mechanical Turk (AMT), and the authors (blinded to the identity of the model and following proper guidelines) manually annotate for relation consistency as it is more technical and AMT annotators were unable to get the qualification questions correct. More details in Appendix.

4.3 Results and Analysis

To use paraphrasing models for downstream tasks, we need to ensure that the generated paraphrases conform to the specified entailment relation and are of good quality.

Automatic evaluation. We first evaluate the pre-trained generators on a held-out set from ParaBank containing

Model	\mathcal{R} -T	BLEU↑	Div↑	iBLEU↑	\mathcal{R} -Con↑
Pre-train-U	✗	14.92	76.73	7.53	–
Pre-train-A	✓	17.20	74.25	8.75	65.53
S2S-U	✗	30.93	59.88	17.62	–
S2S-A	✓	31.44	63.90	18.77	38.42
RR-S2S-U	✓	30.06	64.51	17.26	51.86
RR-FT-U	✓	41.44	53.67	23.96	66.85
ERAP-U*	✓	19.37	69.70	9.43	66.89
ERAP-A	✓	28.20	59.35	14.43	68.61
Fine-tune-U	✗	41.62	51.42	23.79	–
Fine-tune-A	✓	45.21	51.60	26.73*	70.24*
Copy-input	–	51.42	0.00	21.14	45.98

Table 3: Automatic evaluation of ERAP against comparison models described in §4.1. \mathcal{R} -Con is measured only for models conditioned (\mathcal{R} -T) on \mathcal{R} at test time. Fine-tune models are upper- and Pre-train are lower-bound. ★ denotes only pre-training is done in entailment-unaware setting. Bold-face denotes best in each block and * denotes best overall.

500 examples for each relation. Table 2 shows that the entailment-aware generator outperforms its unaware counterpart across all the measures. This boost is observed with weak-supervision for entailment relation demonstrating the good quality of weak-supervision.

Next, We evaluate ERAP variants against the comparison models (§4.1) on the recasted SICK test samples belonging to $\equiv, \sqsubset, \sqsupset$ relation and report the results in Table 3⁴. *Entailment-aware (-A) variants outperform corresponding unaware (-U) variants* on *i*BLEU score, while outperforming the majority-class (i.e., \equiv) copy-input baseline (except for S2S-A). *Weakly-supervised pre-training helps in boosting the performance* in terms of *i*BLEU and \mathcal{R} -Con score as evident from higher scores for Fine-tune-A(U) model over S2S-A(U). Poor performance of S2S variants is because of the small dataset size and much harder multi-objective task. Re-ranking outputs from Fine-tune-U achieve higher *i*BLEU and consistency score than Pre-train-A which is explicitly trained with weak-supervision for relation. However, this comes at the computational cost of sampling multiple⁵ ($k=20$) outputs. Improved performance of Fine-tuned models over S2S indicates the importance of pre-training. Both the ERAP variants achieve higher *i*BLEU and consistency than its lower-bounding (Pre-trained) models but the outputs show less diversity in expression and make conservative lexical or syntactic changes. These results look encouraging until we notice Copy-input (last row) which achieves high BLEU and *i*BLEU, indicating that these metrics fail to punish against copying through the input (an observation consistent with Niu et al. (2020)).

Ablation analysis of each scorer. We demonstrate the effectiveness of each scorer in ERAP via an ablation study in Table 4. Using only consistency scorer for rewarding the generated paraphrases, a significant improvement in con-

⁴We report analogous results for ParaNMT in the Appendix available at <https://arxiv.org/pdf/2203.10483.pdf>.

⁵We report results for $k \in \{1, 5, 10\}$ in the Appendix.

Model	BLEU↑	Div↑	iBLEU↑	\mathcal{R} -Con↑
Gold-reference	–	48.58	–	81.55
Pre-train-A	17.20	74.25	8.75	65.53
+Con	24.82	58.55	12.29	<u>96.75</u>
+Con+Sim	39.78	42.05	20.24	<u>94.72</u>
+Con+Sim+Div	21.68	68.41	11.29	<u>93.60</u>
ERAP-A	28.20	40.65	14.43	68.61

Table 4: Ablation of scorers in ERAP. Con, Sim, Div refers to relation consistency, semantic similarity, and expression diversity scorers. Underline denote more copying of input for Diversity (Div) score and presence of heuristics in outputs for \mathcal{R} -Con score as compared to gold-references.

1. Qualitative outputs for ablation study	
Input \sqsubset	a shirtless man is escorting a horse that is pulling a carriage along a road
Reference	a shirtless man is leading a horse that is pulling a carriage
Only Con	a shirtless man escorts it.
Con+Sim	a shirtless man is escorting a horse.
Con+Sim+Div	a shirtless man escorts a horse.
ERAP-A	a shirtless person is escorting a horse who is dragging a carriage.
2. Heuristic learned w/o Hypothesis-only Adversary	
Input \sqsubset	a man and a woman are walking through a wooded area
Reference	a man and a woman are walking together through the woods
-Adversary	a <i>desert</i> man and a woman walk through a wooded area
ERAP-A	a man and a woman are walking down a path through a wooded area
Input \sqsubset	four girls are doing backbends and playing outdoors
Reference	four kids are doing backbends in the park
-Adversary	four girls do backbends and play outside <i>with mexico</i> .
ERAP-A	four girls do backbends and play games outside.
3. Output from comparison models and ERAP	
Input \equiv	a ribbon is being twirled by a girl in pink
Reference	a girl in pink is twirling a ribbon
S2S-U	a girl is talking to a girl in pink
S2S-A	a girl in blue is surfing on a pink
Fine-tune-U	a girl in pink is turning a tape
Fine-tune-A	a girl in pink is circling the ribbon.
RR-S2S-U	a girl is talking to a girl in a pink built
RR-FT-A	a girl in a pink jacket is swinging a ribbon
Pre-train-U	a girl in pink is circling the ribbon.
Pre-train-A	the ribbon is being twisted by a girl in pink
ERAP-U	the ribbon is being twisted by the girl in the pink
ERAP-A	a girl in pink is twirling a ribbon.

Figure 3: Qualitative outputs: 1 showing the effectiveness of various scorers, 2 showing heuristic learned in the absence of hypothesis-only adversary, and 3 from various models.

sistency score is observed as compared to Pre-train-A and Gold-references. However, this high score may occur at the cost of semantic similarity (e.g., 1 in Figure 3) wherein output conforms to \sqsubset relation at the cost of losing much of the content. Adding similarity scorer, helps in retaining some of the content (higher BLEU and iBLEU) but results in copying (low diversity) from the input. Addition of diversity scorer helps in introducing diversity in expression. However, model is still prone to heuristics (e.g., losing most of the content from input (1 in Figure 3), or adding irrelevant to-

Model	\mathcal{R} -T	Sim↑	Div↑	Gram↑	\mathcal{R} -Con↑
Pre-train-U	✗	4.60	2.62	4.73	–
Pre-train-A	✓	4.67	2.60	4.67	48.00
RR-S2S-U	✓	2.72	3.15	3.46	24.00
RR-FT-U	✓	3.05	2.89	4.27	28.00
ERAP-U	✓	3.98	2.85	4.10	40.00
ERAP-A	✓	3.95	2.68	4.42	64.00
Fine-tune-U	✗	3.87	3.10	4.83	–
Fine-tune-A	✓	3.80	3.04	4.68	48.00

Table 5: Average scores across 3 annotators are reported for Similarity (Sim, $\alpha=0.65$), Diversity (Div, $\alpha=0.55$), and Grammaticality (Gram, $\alpha=0.72$) and % of correct specified relation for \mathcal{R} -Con ($\alpha=0.70$). Moderate to strong inter-rater reliability is observed with Krippendorff’s α .

kens ‘with mexico’ or ‘desert’ (2 in Figure 3)) for ensuring high consistency score. Introducing Adversary reduces the heuristics learned by the generator. Together all the scorers help maintain a good balance for this multi-objective task.

Human evaluation. We report the human evaluation for 25 test outputs each from 8 models for 4 measures in Table 5. ERAP-A achieves the highest consistency while maintaining a good balance between similarity, diversity and grammaticality. RR-S2S-U has the highest diversity which comes at the cost of semantic similarity and grammaticality (e.g., 3 in Figure 3). A strikingly different observation is high similarity and low diversity of Pre-trained variants, reinforcing the issues with existing automatic measures.

5 Extrinsic Evaluation

The intrinsic evaluations show that ERAP produces quality paraphrases while adhering to the specified entailment relation. Next, we examine the utility of entailment-aware paraphrasing models over unaware models for a downstream application, namely paraphrastic data augmentation for textual entailment task. Given two sentences, a premise p and a hypothesis h , the task of textual entailment is to determine if a human would infer h is true from p . Prior work has shown that paraphrastic augmentation of textual entailment datasets improve performance (Hu et al. 2019a); however, these approaches make the simplifying assumption that entailment relations are preserved under paraphrase, which is not always the case (see Figure 1 and 30% of ParaBank pairs were found to be semantically-divergent using Oracle). We use SICK NLI dataset for this task because we have a paraphrasing system trained on similar data distribution⁶.

We hypothesize that entailment-aware augmentations will result in fewer label violations, and thus overall improved performance on the textual entailment task. Moreover, explicit control over the entailment relation allows for greater variety of augmentations that can be generated (an exhaustive list of label preserving augmentations based on entail-

⁶Note that we retained the train, test, development sets of SICK NLI dataset in the recasted SICK dataset and therefore the paraphrasing models have only seen train set.

Type (Label)	Augmentation Pairs
\equiv (E/NE)	$\langle p', h \rangle \langle p, h' \rangle \langle p', h' \rangle$
\sqsupset (E/NE)	$\langle pr, h \rangle \langle pr, h' \rangle$
\sqsubset (E/U)	$\langle p, hf \rangle \langle p', hf \rangle \langle pr, hf \rangle$
Unknown (U/U)	$\langle pf, h' \rangle \langle pf, hf \rangle \langle p, hr \rangle \langle p', hr \rangle$
	$\langle pr, hr \rangle \langle pf, r \rangle \langle pf, h \rangle$

Table 6: Various augmentations for $\langle p, h \rangle$ with label as E/NE (entails/does not entail) grouped as per the type (and corresponding projected labels) according to the entailment composition rules defined in MacCartney (2009). $p'(h')$, $pr(hr)$, $pf(hf)$ denote \equiv , \sqsupset , and \sqsubset paraphrase, resp.

ment relation between a p (or h) and its paraphrase is presented in Table 6) with entailment-aware models.

Paraphrastic Data Augmentation We generate paraphrases for all premises $p \in P$ and hypotheses $h \in H$ present in the train set of SICK NLI using entailment-aware and unaware models. We obtain augmentation data by combining all the paraphrases (generated using entailment-aware models) with original data and label them as per Table 6. Augmentation paraphrases generated from entailment-unaware models are (naïvely) assumed to hold the \equiv relation. RoBERTa-based binary classifiers are trained on original dataset along with the paraphrastic augmentations to predict whether p entails h .

Susceptibility to Augmentation Artifacts. If paraphrastic augmentations introduce noisy training examples with incorrectly projected labels, this could lead to, what we call *augmentation artifacts* in downstream models. We posit that paraphrastically augmented textual entailment (henceforth, PATE) models trained on entailment-aware augmentations will be less susceptible to such artifacts than the models trained with entailment-unaware augmentations. To test this, we generate augmentations for the test set of SICK NLI and manually annotate 1253 augmented samples to obtain 218 incorrectly labeled examples. We evaluate PATE models on these examples (referred to as *adversarial test examples*).

Extrinsic Results. We report accuracy of PATE models on original SICK development and test sets as well as on *adversarial test examples* in Table 7. As per our hypothesis, models trained with augmentations generated using entailment-aware models result in improved accuracy on both original as well as adversarial test samples over those trained with entailment-unaware augmentations. Textual entailment model trained only on SICK NLI data performs the best on *adversarial test set* as expected and proves that although augmentation helps in boosting the performance of a model, it introduces *augmentation artifacts* during training.

6 Related Work

Paraphrase generation is a common NLP task with widespread applications. Earlier approaches are rule-based (Barzilay, McKeown, and Elhadad 1999; Ellsworth and Janin 2007) or data-driven (Madnani and Dorr 2010). Recent, supervised deep learning approaches use

Data	\mathcal{R} -T	O-Dev \uparrow	O-Test \uparrow	A-Test \uparrow
SICK NLI	-	95.56	93.78	83.02
+FT-U(\equiv)	\times	95.15	93.68	69.72
+FT-A(\equiv)		95.35	94.62	77.98
+FT-A(\equiv, \sqsupset)	\checkmark	95.76	93.95	75.69
+ERAP-A(\equiv)		95.15	94.58	78.44
+ERAP-A(\equiv, \sqsupset)		95.15	93.86	69.72

Table 7: Accuracy results of PATE models for Original (O-) and Adversarial (A-) datasets. FT/ERAP refers to the Fine-tuned/proposed model used for generating augmentations. Type of augmentation used as per Table 6 in parenthesis. U/A denote entailment-unaware (aware) variant.

LSTMs (Prakash et al. 2016), VAEs (Gupta et al. 2018), pointer-generator networks (See, Liu, and Manning 2017), and transformer-based (Li et al. 2019) sequence-to-sequence models. Li et al. (2018) use RL for supervised paraphrasing. **Unsupervised paraphrasing** is a challenging and emerging NLP task with limited efforts. Bowman et al. (2015b) train VAE to sample less controllable paraphrases. Others use metropolis-hastings (Miao et al. 2019), simulated annealing (Liu et al. 2020) or dynamic-blocking (Niu et al. 2020) to add constraints to the decoder at test time. Siddique, Oymak, and Hristidis (2020) use RL to maximize expected reward based on adequacy, fluency and diversity. Our RL-based approach draws inspiration from this work by introducing oracle and hypothesis-only adversary.

Controllable text generation is a closely related field with efforts been made to add lexical (Hu et al. 2019a; Garg et al. 2021) or syntactic control (Iyyer et al. 2018; Chen et al. 2019; Goyal and Durrett 2020) to improve diversity of paraphrases. However, ours is the first work which introduces a semantic control for paraphrase generation.

Style transfer is a related field that aims at transforming an input to adhere to a specified target attribute (*e.g.*, sentiment, formality). RL has been used to explicitly reward the output to adhere to a target attribute (Gong et al. 2019; Sancheti et al. 2020; Luo et al. 2019; Liu, Neubig, and Wieting 2020; Goyal et al. 2021). The target attributes are only a function of the output and defined at a lexical level. However, we consider a relation control which is a function of both the input and the output, and is defined at a semantic level.

7 Conclusion

We introduce a new task of entailment-relation-aware paraphrase generation and propose a RL-based weakly-supervised model (ERAP) that can be trained without a task-specific corpus. Additionally, an existing NLI corpora is recasted to curate a small annotated dataset for this task, and provide performance bounds for it. A novel Oracle is proposed to obtain weak-supervision for relation control for existing paraphrase corpora. ERAP is shown to generate paraphrases conforming to the specified relation while maintaining quality of the paraphrase. Intrinsic and Extrinsic experiments demonstrate the utility of entailment-relation control, indicating a fruitful direction for future research.

References

- Androutsopoulos, I.; and Malakasiotis, P. 2010. A survey of paraphrasing and textual entailment methods. *Journal of Artificial Intelligence Research*, 38: 135–187.
- Barzilay, R.; McKeown, K.; and Elhadad, M. 1999. Information fusion in the context of multi-document summarization. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics*, 550–557.
- Bhagat, R.; and Hovy, E. 2013. What is a paraphrase? *Computational Linguistics*, 463–472.
- Bhagat, R.; Pantel, P.; and Hovy, E. 2007. LEDIR: An unsupervised algorithm for learning directionality of inference rules. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015a. A large annotated corpus for learning natural language inference. In *Conference on Empirical Methods in Natural Language Processing, EMNLP 2015*, 632–642. Association for Computational Linguistics (ACL).
- Bowman, S. R.; Vilnis, L.; Vinyals, O.; Dai, A. M.; Jozefowicz, R.; and Bengio, S. 2015b. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*.
- Carpuat, M.; Vyas, Y.; and Niu, X. 2017. Detecting cross-lingual semantic divergence for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, 69–79.
- Chen, M.; Tang, Q.; Wiseman, S.; and Gimpel, K. 2019. Controllable Paraphrase Generation with a Syntactic Exemplar. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5972–5984.
- Ellsworth, M.; and Janin, A. 2007. Mutaphrase: Paraphrasing with framenet. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, 143–150.
- Ganitkevitch, J.; Van Durme, B.; and Callison-Burch, C. 2013. PPDB: The paraphrase database. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 758–764.
- Garg, S.; Prabhu, S.; Misra, H.; and Srinivasaraghavan, G. 2021. Unsupervised contextual paraphrase generation using lexical control and reinforcement learning. *arXiv preprint arXiv:2103.12777*.
- Glockner, M.; Shwartz, V.; and Goldberg, Y. 2018. Breaking NLI Systems with Sentences that Require Simple Lexical Inferences. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 650–655.
- Gong, H.; Bhat, S.; Wu, L.; Xiong, J.; and Hwu, W.-M. 2019. Reinforcement Learning Based Text Style Transfer without Parallel Training Corpus. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Goyal, N.; Srinivasan, B. V.; Anandhavelu, N.; and Sancheti, A. 2021. Multi-Style Transfer with Discriminative Feedback on Disjoint Corpus. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3500–3510.
- Goyal, T.; and Durrett, G. 2020. Neural Syntactic Preordering for Controlled Paraphrase Generation. In *Annual Meeting of the Association for Computational Linguistics*.
- Gupta, A.; Agarwal, A.; Singh, P.; and Rai, P. 2018. A deep generative framework for paraphrase generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Holtzman, A.; Buys, J.; Du, L.; Forbes, M.; and Choi, Y. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Hu, J. E.; Khayrallah, H.; Culkin, R.; Xia, P.; Chen, T.; Post, M.; and Van Durme, B. 2019a. Improved lexically constrained decoding for translation and monolingual rewriting. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 839–850.
- Hu, J. E.; Rudinger, R.; Post, M.; and Van Durme, B. 2019b. ParaBank: Monolingual bitext generation and sentential paraphrasing via lexically-constrained neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6521–6528.
- Iyyer, M.; Wieting, J.; Gimpel, K.; and Zettlemoyer, L. 2018. Adversarial Example Generation with Syntactically Controlled Paraphrase Networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1875–1885.
- Kobus, C.; Crego, J. M.; and Senellart, J. 2017. Domain Control for Neural Machine Translation. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, 372–378.
- Kusner, M.; Sun, Y.; Kolkin, N.; and Weinberger, K. 2015. From word embeddings to document distances. In *International conference on machine learning*, 957–966. PMLR.
- Li, Z.; Jiang, X.; Shang, L.; and Li, H. 2018. Paraphrase Generation with Deep Reinforcement Learning. In *EMNLP*.
- Li, Z.; Jiang, X.; Shang, L.; and Liu, Q. 2019. Decomposable neural paraphrase generation. *arXiv preprint arXiv:1906.09741*.
- Liu, X.; Mou, L.; Meng, F.; Zhou, H.; Zhou, J.; and Song, S. 2020. Unsupervised Paraphrasing by Simulated Annealing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 302–312.
- Liu, Y.; Neubig, G.; and Wieting, J. 2020. On Learning Text Style Transfer with Direct Rewards. *arXiv preprint arXiv:2010.12771*.

- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Luo, F.; Li, P.; Zhou, J.; Yang, P.; Chang, B.; Sui, Z.; and Sun, X. 2019. A dual reinforcement learning framework for unsupervised text style transfer. *arXiv preprint arXiv:1905.10060*.
- MacCartney, B. 2009. *Natural language inference*. Stanford University.
- Madnani, N.; and Dorr, B. J. 2010. Generating phrasal and sentential paraphrases: A survey of data-driven methods. *Computational Linguistics*, 36(3): 341–387.
- Marelli, M.; Menini, S.; Baroni, M.; Bentivogli, L.; Bernardi, R.; Zamparelli, R.; et al. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Lrec*, 216–223. Reykjavik.
- McCoy, T.; Pavlick, E.; and Linzen, T. 2019. Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3428–3448.
- Miao, N.; Zhou, H.; Mou, L.; Yan, R.; and Li, L. 2019. Cgmh: Constrained sentence generation by metropolis-hastings sampling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6834–6842.
- Min, J.; McCoy, R. T.; Das, D.; Pitler, E.; and Linzen, T. 2020. Syntactic Data Augmentation Increases Robustness to Inference Heuristics. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2339–2352.
- Niu, T.; Yavuz, S.; Zhou, Y.; Wang, H.; Keskar, N. S.; and Xiong, C. 2020. Unsupervised paraphrase generation via dynamic blocking. *arXiv preprint arXiv:2010.12885*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Pavlick, E.; Bos, J.; Nissim, M.; Beller, C.; Van Durme, B.; and Callison-Burch, C. 2015. Adding semantics to data-driven paraphrasing. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1512–1522.
- Pham, M. Q.; Crego, J. M.; Senellart, J.; and Yvon, F. 2018. Fixing translation divergences in parallel corpora for neural mt. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2967–2973.
- Poliak, A.; Naradowsky, J.; Haldar, A.; Rudinger, R.; and Van Durme, B. 2018. Hypothesis Only Baselines in Natural Language Inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*.
- Prakash, A.; Hasan, S. A.; Lee, K.; Datla, V.; Qadir, A.; Liu, J.; and Farri, O. 2016. Neural paraphrase generation with stacked residual LSTM networks. *arXiv preprint arXiv:1610.03098*.
- Sancheti, A.; Krishna, K.; Srinivasan, B. V.; and Natarajan, A. 2020. Reinforced rewards framework for text style transfer. *Advances in Information Retrieval*, 12035: 545.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Siddique, A.; Oymak, S.; and Hristidis, V. 2020. Unsupervised paraphrasing via deep reinforcement learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1800–1809.
- Sun, H.; and Zhou, M. 2012. Joint learning of a dual SMT system for paraphrase generation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 38–42.
- Valencia, V. M. S. 1991. *Studies on natural logic and categorical grammar*. Universiteit van Amsterdam.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wieting, J.; and Gimpel, K. 2018. ParaNMT-50M: Pushing the Limits of Paraphrastic Sentence Embeddings with Millions of Machine Translations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 451–462.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1112–1122.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3): 229–256.
- Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Zhao, W.; Peyrard, M.; Liu, F.; Gao, Y.; Meyer, C. M.; and Eger, S. 2019. MoverScore: Text Generation Evaluating with Contextualized Embeddings and Earth Mover Distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 563–578.