

Paraphrase Acquisition from Image Captions

Marcel Gohsen

Bauhaus-Universität Weimar

Matthias Hagen

Friedrich-Schiller-Universität Jena

Martin Potthast

Leipzig University and ScaDS.AI

Benno Stein

Bauhaus-Universität Weimar


Abstract

We propose to use image captions from the Web as a previously underutilized resource for paraphrases (i.e., texts with the same “message”) and to create and analyze a corresponding dataset. When an image is reused on the Web, an original caption is often assigned. We hypothesize that different captions for the same image naturally form a set of mutual paraphrases. To demonstrate the suitability of this idea, we analyze captions in the English Wikipedia, where editors frequently relabel the same image for different articles. The paper introduces the underlying mining technology, the resulting Wikipedia-IPC dataset, and compares known paraphrase corpora with respect to their *syntactic and semantic paraphrase similarity* to our new resource. In this context, we introduce characteristic maps along the two similarity dimensions to identify the style of paraphrases coming from different sources. An annotation study demonstrates the high reliability of the algorithmically determined characteristic maps.

1 Introduction

Two texts that convey “semantically equivalent information” via a different wording are called *paraphrases*, or said to *be in a paraphrase relation*. A variety of natural language processing tasks have been approached by using paraphrase resources or paraphrase generation and detection algorithms: textual entailment (Dagan et al., 2013; Marelli et al., 2014; Izadinia et al., 2015), semantic similarity (Agirre et al., 2015; Li and Srikumar, 2016), machine translation (Mehdizadeh Seraj et al., 2015), text reformatting (Stein et al., 2014), or question answering (Fader et al., 2013). Paraphrasing can help to improve sentence compression and text summarization (Cordeiro et al., 2007) as well as natural language understanding (Ganitkevitch, 2013).

To tackle the task of paraphrasing, large amounts of paraphrases usually are helpful. Hence, para-



Flourine-containing durable water repellent makes a fabric water-resistant.

Water bead on a fabric that has been made non-wetting by chemical treatment.

Figure 1: Paraphrased captions of an image in the Wikipedia pages ‘Durable water repellent’ (left) and ‘Wetting’ (right).

phrase acquisition has gained popularity over the recent years and paraphrase datasets have been created manually, (semi-)automatically, and via distant supervision resulting in different levels of “naturalness”, diversity (topic, genre, language, etc.), granularity (word, sentence, or passage level), and scale (see Section 2 for details). For model training, Zhang et al. (2019b) strongly emphasize the importance of negative examples that are lexically similar but semantically inappropriate as paraphrases.

With image captions, we propose an as yet underutilized source of natural paraphrases. The underlying hypothesis is that different captions for the same image form paraphrase candidates because the captions may describe the same image content. We study the phenomenon of ‘image reuse and captioning’ in the English Wikipedia (Figure 1 shows an example) by mining captions of images that have been reused. Carefully optimized filtering heuristics for captions and associated images yield a large set of captions, which we analyze qualitatively and quantitatively for their usefulness as paraphrases. The set of mined image captions forms our Wikipedia-IPC dataset,^{1,2} which contains caption pairs as paraphrases along with their image source three different quality levels; 30,237 caption pairs are gold-, 229,877 are silver-, and 656,560

¹Data: <https://doi.org/10.5281/zenodo.7621320>

²Code: <https://github.com/webis-de/EACL-23>

are bronze-quality paraphrase candidates.

The usefulness of an image caption as a paraphrase depends strongly on its corresponding image. Table 1 shows example image–caption pairs from Wikipedia, organized by content and purpose classes. Icons, symbols, and pictograms rarely have captions that are useful as paraphrase candidates, while for pictures, graphics, paintings, maps, and schematics, the purpose of the image in a page is crucial. In this regard, “ideological purpose” means that the image is mainly used for aesthetic reasons, resulting in captions that do not describe the shown content but rather some context, atmosphere, or sentiment. By contrast, explanatory or pragmatic images try to illustrate or transport knowledge. Their captions describe and explain (parts of) the visual content and form a particularly promising source of natural paraphrases.

2 Related Work

Paraphrase Corpora To collect high-quality paraphrases, some studies applied manual acquisition. For instance, for the Webis-CPC-11 (Burrows et al., 2013), crowdworkers were asked to rewrite one of 4,096 passage-level literature excerpts “*so that the rewritten version has the same meaning, but a completely different wording and phrasing.*” The resulting 7,859 pairs form one of the few passage-level paraphrase corpora. Later, also Xu et al. (2014) crowdsourced sentences with the same meaning as a shown sentence from some tweet for their PIT-2015 dataset (18,862 sentence pairs, 5,641 considered paraphrases).

Since crowdsourcing is costly to scale, automatic paraphrase acquisition often simply relies on machine translation. For instance, Zhang et al. (2019b) scrambled words and forth-and-back translated sentences to obtain a gold dataset of 108,463 paraphrase pairs and a silver dataset of 656,000 pairs. Also, the ParaNMT-50M dataset (Wieting and Gimpel, 2018) with 30 million “strong” paraphrases (assessment by the authors), was created by translating sentences from English-Czech parallel pairs.

Trying to combine the advantages of manual and automatic methods (quality vs. scale), most paraphrase acquisition methods use distant supervision. For instance, Barzilay and McKeown (2001) suggested aligned sentences from monolingual parallel corpora as paraphrases—an idea later used by Dolan and Brockett (2005) to create the Microsoft Research Paraphrase Corpus (5,801 man-

ually annotated sentence pairs, 3,900 considered paraphrases). Instead of monolingual sources, [Barnard and Callison-Burch \(2005\)](#) exploited bilingual parallel corpora by aligning phrases that translate into the same pivot phrase in another language. Based on this idea, the English/Spanish PPDB 1.0 ([Ganitkevitch et al., 2013](#)) and its extension to 23 additional languages ([Ganitkevitch and Callison-Burch, 2014](#)) with an English portion of 7.6 million lexical (i.e., synonymous) and 68.4 million phrasal paraphrases form the biggest paraphrase collection.

Translation pivoting was also used for the Opusparcus ([Creutz, 2018](#)) and TaPaCo ([Scherrer, 2020](#)) datasets. Opusparcus was created from movie and TV subtitles in six European languages. The automatically filtered training set contains 7 million English sentence and phrase pairs (“good” or “mostly good” paraphrases according to the authors) while the manually labeled development and test sets contain 3,088 pairs (1,997 paraphrases). Analogously, TaPaCo is created from the multilingual Tatoeba data, a crowdsourced collection of sentence translations.³ The English part of TaPaCo contains 158,000 sentences that are clustered in 62,000 paraphrase sets. Instead of translation, [Lan et al. \(2017\)](#) used a different pivoting method for their Twitter News URL Corpus. They mined tweets that contain the same hyperlink and labeled 51,524 respective sentence pairs via crowdsourcing. Our paraphrase acquisition methodology is also based on a pivoting technique, but the pivot media are images.

Captions and Paraphrases Our idea to use images from the web as pivots for distant supervision paraphrase acquisition is inspired by studies on crowdsourced caption datasets. For instance, [Marelli et al. \(2014\)](#) extracted a collection of 9,840 paraphrase candidates from the Flickr 8k dataset ([Rashtchian et al., 2010](#)) of 8,000 images each with 5 crowdsourced captions and the MSR Video Paraphrase Corpus ([Chen and Dolan, 2011](#)) of 2,000 short video clips with 85,000 English crowdsourced descriptions of what can be seen. The larger MSCOCO dataset ([Lin et al., 2014](#)) with 120,000 images each having 5 crowdsourced captions was even already used to train a neural paraphrase generation model ([Prakash et al., 2016](#)). Interestingly, the existing caption-paraphrasing datasets were all created in dedicated crowdsourcing processes that are costly to scale. We are the first to explore mining “real” image cap-

³<https://tatoeba.org/eng/about>


















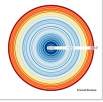







| Image content | | Image purpose | |
|---|---|---|--|
| Icons, symbols, pictograms | Pictures, graphics, paintings, maps, schematics | | |
| | Ideological (used to convey atmosphere and sentiment) | Specific, explanatory, pragmatic (used in discourse) | |
|  Brazilian flag  Up arrow  Wikimedia community logo  Yellow card  Icon for plain stage in cycling stage race  Motorway exit icon  Pictograms of Olympic sports Athletics  |  Paraguayan goalkeeper Justo Villar was awarded as the best goalkeeper of the tournament.  AC-130 Spectres were highly effective during the battle.  Toronto's MaRS Discovery District is a centre for research in biomedicine.  J. C. A. Corea was born in the historic town of Chilaw, on the west coast of Sri Lanka.  Dostana was filmed almost entirely in Miami.  February 10: New Delhi becomes India's capital.  The Armenian genocide (pictured) was the first event officially condemned as "crimes against humanity".  "I Blame Myself" was co-written by Ariel Rechtshaid (pictured), who additionally assisted with the songwriting and production of each track on Night Time, My Time. |  Semi hollow-body electric guitar Gibson ES-335 has a "solid center block" inside a body.  This "circular warming stripes" graphic depicts average global temperature using chronologically ordered, concentric coloured rings.  Water bead on a fabric that has been made non-wetting by chemical treatment.  Dikes are used to protect the rice paddy fields from the channels of saltwater which overflow during high tide.  The camouflage grouper Epinephelus polyphkadion is the type-host of <i>P. viscosus</i> ; the photograph shows a specimen from New Caledonia.  The stone is against the half-timbered wall of Exmewed Hall, on St. Peter's Square, Ruthin.  A sub-surface Metropolitan line A Stock train (left) passes a deep-tube Piccadilly line 1973 Stock train (right) in the siding at Rayners Lane.  The X-38 CRV prototype makes a gentle lakebed landing at the end of a July 1999 test flight at the Dryden Flight Research Center with a fully deployed parafoil.  Tylosaurus proriger mounted skeleton in the Rocky Mountain Dinosaur Resource Center in Woodland Park, Colorado. | |

Table 1: Example Wikipedia images and captions. Following Bär (2008), we categorize images by content as icons, symbols, pictograms (left), or pictures, graphics, etc. (middle and right). Pictures, graphics, etc. can be further subdivided by the purpose they serve in a text as either ideological (i.e., to convey sentiment; middle) or pragmatic (i.e., to explain and specify; right).

tions from the web (i.e., not crowdsourced) as a paraphrase acquisition method.

3 Caption-Based Paraphrase Acquisition

We collect captions of images that are used on the English Wikipedia, grouping those captions that describe the same image. This section outlines the mining and filtering steps.

Data Source: Wikipedia As a source for paraphrases, we use captions from Wikipedia articles. Wikipedia is a collection of well-formulated, human-written text that is constantly being improved by a community of volunteer editors. The guidelines for the use of images on Wikipedia encourage editors to write a caption to explain the relevance of an image in a particular context, resulting in 60% of images having a caption.

We use the English Wikipedia for the paraphrase

| Dump | Pages | Revisions | Images | Image refs. |
|-------|------------|-------------|------------|---------------|
| Small | 22,311,116 | 22,311,116 | 4,508,808 | 7,384,368 |
| Full | 54,898,564 | 987,527,838 | 11,183,853 | 3,629,447,851 |

Table 2: Image mining statistics of Wikipedia dumps. "Small" refers to all articles, templates, media descriptions, and primary meta-pages without edit histories, whereas "Full" includes edit histories.

acquisition (Wikimedia, September 1st, 2022).⁴ Two versions of Wikipedia dumps are available for the paraphrase acquisition—a reduced dataset with all articles, templates, media descriptions, and primary meta pages and a full-size dump with all pages, including their full revision history. Table 2 compares the number of pages, revisions, images, and image references between these two datasets.

⁴<https://dumps.wikimedia.org/enwiki/>

| Filter | Images | | | | Image references | | | Image captions | | Paraphrase cand. | |
|-----------------------------|-----------|----------|----------------|----------------|------------------|----------|--------------------|----------------|----------|------------------|----------|
| | Remaining | Δ | Refs. ≥ 2 | Refs. ≥ 5 | Remaining | Δ | μ_{img} | Remaining | Δ | Remaining | Δ |
| 0. No filter | 4,508,808 | – | 941,339 | 102,046 | 7,384,368 | – | 1.64 | 4,571,671 | – | 355,619,088 | – |
| 1. References ≥ 2 | 941,339 | -79% | 941,339 | 102,046 | 3,816,899 | -48% | 4.05 | 2,053,338 | -55% | 355,619,088 | 0% |
| 2. References ≤ 10 | 918,758 | -2% | 918,758 | 79,465 | 2,494,271 | -35% | 2.71 | 1,569,327 | -23% | 1,350,050 | -100% |
| 3. Has caption | 713,815 | -22% | 494,321 | 32,949 | 1,499,479 | -40% | 2.10 | 1,569,327 | 0% | 1,350,050 | 0% |
| 4. Caption words ≥ 6 | 518,851 | -27% | 300,244 | 18,436 | 987,541 | -34% | 1.90 | 1,014,351 | -35% | 760,448 | -44% |
| 5. Caption is sentence | 97,496 | -81% | 23,062 | 665 | 128,893 | -87% | 1.32 | 129,776 | -87% | 44,005 | -94% |
| 6. References ≥ 2 | 23,062 | -76% | 23,062 | 665 | 54,459 | -58% | 2.36 | 54,940 | -58% | 44,005 | 0% |
| 7. Unique candidates | 22,643 | -2% | 22,582 | 507 | 52,394 | -4% | 2.31 | 52,394 | -5% | 39,995 | -9% |
| 8. Divergent captions | 18,979 | -16% | 18,979 | 391 | 43,616 | -17% | 2.30 | 43,616 | -17% | 32,830 | -18% |
| 9. Significant caption diff | 17,293 | -9% | 17,293 | 386 | 40,131 | -8% | 2.32 | 40,131 | -8% | 30,237 | -6% |

Table 3: Effects of the filtering steps in the paraphrase acquisition pipeline on the number of images, references, captions, and paraphrase candidates mined from the reduced Wikipedia dump. Images that have been referenced at least 2 or 5 times are called “Refs. ≥ 2 ” and “Refs. ≥ 5 ”, respectively. μ_{img} describes the average number of references per img.

Image and Caption Mining Wikimedia Commons is a service which hosts free-to-use images and other media files which are used across all Wikimedia projects (like Wikipedia). The markup language for Wikipedia articles, Wikitext, features the *extended image syntax*, which allows editors to insert images into Wikipedia pages from Wikimedia Commons. The specified image is referenced by its unique identifier in Wikimedia Commons, which is a concatenation of a generic media type prefix (e.g., Image) followed by a colon and its (also unique) filename. The extended image syntax allows the specification of style properties, meta-information, and a caption. If a caption is present, we store it along the image reference.

However, not all images in Wikipedia pages are included by the use of the extended image syntax. Wikitext allows the use of (user-created) style templates, often with inconsistent image reference syntax. *Infoboxes* on Wikipedia are template-based tables placed at the beginning of an article that contain a collection of the most important information about its subject. They are commonly used in Wikipedia—20% of all the pages contain one. Important to us is that one third of the infoboxes contain one or even multiple images, and about 56% of these images have a caption. Wikitext represents infoboxes as lists of key-value pairs whose sets of allowed keys and values are defined in templates. There are close to 2,000 different infobox templates based on semantic categories of an article. The responsible keyword for image referencing varies among these templates, and to cover all cases is not feasible. Thus, we identified the most commonly

used keyword which is “image”, followed by an optional counter for multiple images. These captions are mined as individual image–caption pairs. No Wikitext parser other than Wikipedia’s own implementation (which is monolithically intertwined with Wikipedia’s software stack) is capable of reliably parsing all the different template syntaxes. Our heuristic for image mining from infoboxes yields an increase of 24% of image references.

In addition to a caption, an image reference may also have an *alternative text* that accurately describes the visual content of an image. Alternative texts are intended to serve as a substitute in case the user is visually impaired, or the image cannot be displayed correctly. According to Wikipedia’s policy, all images are obliged to have an alternative text, except those with a purely decorative purpose. We extract them as paraphrase candidates, too. Altogether, 355 million caption pairs are mined that we subsequently filter as follows (see Table 3).

Image Filter A necessary condition for an image to be useful is that it has been referenced at least twice, so we discard all image captions whose images have only one reference. As illustrated in Table 1, the likelihood of image–caption pairs to form a pair of paraphrase candidates depends on the image’s content and purpose. A useful feature to discard icons, symbols, and pictograms is the number of references. These images occur more frequently than any other kind of image. To discard them, we set an upper bound of 10 references per image for the corresponding image–caption pairs to be retained. After these steps, 1.35 million image–caption pairs remain (see Table 3).

Caption Filter The selection of caption pairs has a substantial effect on the quality of the acquired paraphrases. As a consequence, all mined image captions pass through various carefully optimized pre-processing and filtering steps. The pre-processing pipeline comprises cleansing of Wiki-text markup, removal of line-breaks, and white-space normalization (e.g., removal of sequences).

To discard trivial image captions, we determined a lower bound for the caption length through manual observation and found that most captions that represent a phrasal expression contain at least 6 words. Captions with less than that usually comprise the name of an entity that can be seen in the image and are consequently discarded.

Through manual review, we identified three common practices for writing image captions with respect to style and syntax. Captions on Wikipedia are either: (1) simple noun phrases (e.g., “*Last Supper by Dieric Bouts*”), (2) sentence fragments containing a verb phrase (e.g., “*Last Supper drawn by Dieric Bouts*”), or (3) grammatically correct sentences (e.g., “*The Last Supper was drawn by Dieric Bouts*”). The latter two caption types are more likely to be a source of high-quality paraphrases, whereas grammatical sentences are most desirable. To discard noun phrases, it is sufficient to exclude captions that do not contain verbs. The detection of proper sentences is far more difficult.

Grammaticality and sentence fragment detection is a broad research field and most state-of-the-art methods use neural models. However, to ensure high precision in the resulting paraphrase dataset, we resorted to a more efficient and effective heuristic that is well suited for captions: rule-based sentence classification. Since the authors of existing rule-based grammar checkers rarely publish their full rule sets, we create our own rules to identify grammatical sentences. To find suitable rules, we take 500 and 100 random captions from Wikipedia to create a training and test set, respectively. An expert with 20 years of experience in the field of linguistics manually annotated the 600 captions to determine whether they are sentences or sentence fragments. Based on the 500 annotations from the training set and the associated automatically assigned POS tags, the expert manually created a set of rules to distinguish fragments from sentences.

Table 4 shows our sentence classification rules based on POS tags from Penn Treebank (Taylor et al., 2003). To be considered a sentence, a cap-

| Rule | Premise | Pattern |
|------|-------------------------|--|
| 1 | . *MD.* | . * MD RB? VB.* |
| 2 | . * (WDT WP WRB) .* | [¬(WDT WP WRB)]* (VBP VBZ VBD) .* |
| 3 | . * IN.* | [¬IN]* (VBP VBZ VBD) .* |
| 4 | ⊥ | . * (VBP VBZ VBD) .* |

Table 4: POS-based rules for sentence classification.

tion must satisfy one of these rules. Which rule should be applied is determined by its premises, which are tested in the order given. The first rule targets sentences that contain modals. A modal must be directly followed by an optional adverb and an obligatory verb in the base form. A sentence that would be correctly classified as such is “*Last Supper might be drawn by Dieric Bouts*”. The second and third rules deal with subordinate clauses at the end of a sentence. These rules state that an inflected verb must precede a subordinating conjunction. For example, “*Last Supper was drawn by Dieric Bouts which is an exceptional artwork*” would be correctly classified as a sentence under Rule 2. Rule 4 requires that sentences contain an inflected verb and is applied when no other premise is satisfied. “*Dieric Bouts drew the Last Supper*” is a correctly classified sentence according to Rule 4.

Our heuristic sentence classifier performs satisfactorily with a precision of 94% and a recall of 79%, evaluated on the 100 annotated captions in the test set. Problematic cases are verbs that have identical base and inflected forms, which are often mislabeled by the POS tagger of the Stanford CoreNLP toolkit (Manning et al., 2014).

After application of these caption filters, in total 44,005 image–caption pairs remain (see Table 3).

Image Equivalence Wikipedia hosts all of its media files on Wikimedia Commons, and therefore we can identify duplicate images without requiring physical copies of an image. All images have a unique name that, along with a file type prefix, points to the URI of their description page on Wikimedia Commons. When an editor of a Wikipedia article wants to include an image from Wikimedia Commons, the image is referenced by this URI. Therefore, this URI is the criterion for deciding whether images are equivalent. All file type prefixes are derived from the common “File” prefix. Therefore, all image prefixes are normalized to “File” prefixes to detect image duplicates with inconsistent prefix usage.

Paraphrase Construction All references that refer to the same image are grouped into clusters. Paraphrase candidates are constructed by forming all possible unique pairs of captions of the same caption type (i.e., alternative text or regular caption) within each cluster. Combining regular captions and alternative texts rarely results in meaningful paraphrases due to their different purposes. Alternative texts describe the visual content of an image, while regular captions describe the image in context and from the perspective of an article.

Of these paraphrase candidates, we discard caption pairs that are exact or near-duplicates. We consider a caption pair to be a near-duplicate if it differs only by punctuation, capitalization, or whitespaces. (Near-)duplicates of captions are artifacts of reusing centrally hosted captions for the same image in different articles on Wikipedia. Removing duplicates results in 30,237 caption pairs as paraphrase candidates (see Table 3).

Corpus Construction We create three individual paraphrase datasets with different quality levels which we assign based on the used Wikipedia dump and the set of applied filtering heuristics.

Gold-quality paraphrase candidates are pairs of acquired image captions from the reduced Wikipedia dump which are classified as (sequences of) sentences by our rule-based classifier. Table 3 shows the loss of images, references, captions, and paraphrase candidates due to the application of our filter heuristics. In total, 30,237 gold-quality paraphrase candidates were acquired with our pipeline.

The most “aggressive” filter heuristic in the pipeline is the sentence classification of image captions. It discards 87% of the mined image captions. Since phrases can also be in a paraphrase relation, we construct a silver-quality paraphrase dataset by exchanging the sentence classifier for a heuristic that requires image captions to contain a verb. This heuristic excludes 50% of image captions for the paraphrase acquisition. This altered pipeline mined 229,877 silver-quality paraphrase candidates.

For the bronze-quality paraphrase dataset, we collect as many reasonable caption-pairs as possible, using the full Wikipedia dump with edit history as source for the paraphrase acquisition. We set the upper bound of allowed number of references per image to 18, the average number of revisions per page. Thus, an image can be referenced at most 180 times. The verb heuristic is applied. This yields 656,560 bronze-quality paraphrase candidates.

4 Corpus Analysis and Evaluation

We analyze and evaluate our paraphrase corpus based on the quantification of paraphrase similarity, for which we select well-known syntactic and semantic measures, and compare our corpus with five state-of-the-art paraphrase and three image–caption corpora from the literature. Therefore, we propose a new measure $\Delta_{\text{sem},\text{syn}}$ that measures paraphrase sophistication based on the difference of semantic and syntactic similarities. To validate the proposed measure we analyze the correlation between human judgments and automatically computed semantic and syntactic similarities.

Analyzing Paraphrase Similarity

According to Wahle et al. (2022), high-quality paraphrases are texts pairs with high semantic similarity and high lexical and syntactic diversity, as paraphrasing models trained on linguistically diverse examples tend to be more robust (Qian et al., 2019). Similarly, Niu et al. (2021) measures paraphrase quality by rewarding high semantic similarity while penalizing high lexical overlap. The relation between semantic and syntactic similarity characterizes the “sophistication” of a paraphrase. Based on this assumption, we propose $\Delta_{\text{sem},\text{syn}}$ to quantify the sophistication of paraphrases. $\Delta_{\text{sem},\text{syn}}$ computes the average difference between the average semantic and syntactic similarity scores.

Similarity Measures To measure syntactic similarities between pairs of paraphrases, we choose ROUGE-1, ROUGE-L (Lin, 2004), and BLEU (Papineni et al., 2002). ROUGE-1 assesses unigram overlap, ROUGE-L computes lexical similarity based on the longest-common-subsequence paradigm, and BLEU is computed up to 4-grams to quantify structural similarity. All three measures have been proposed and evaluated for paraphrase generation and acquisition.

The computation of semantic similarity is done by transformer-based models. One of our three measures is BERTScore (Zhang et al., 2019a) since it was found to correlate well with human judgments for the task of automatic image captioning. The second one is the cosine similarity of dense vector representations computed by a BERT-based Sentence Transformer (Reimers and Gurevych, 2019). It proved to be decisive at Sent-Eval on the MSRPC. The third one is Word Mover Distance (Kusner et al., 2015) which computes the minimum amount of distance that embedded words of a text need to

| Image Caption Pairs | Syntactic similarity | | | | Semantic similarity | | | | |
|--|----------------------|---------|------|------|---------------------|------|------|------|----------------------------|
| | ROUGE-1 | ROUGE-L | BLEU | Avg. | WMS | BERT | ST | Avg. | $\Delta_{\text{sem, syn}}$ |
| A1: An Easter postcard from 1907 depicting a rabbit. A2: A 1907 postcard featuring the Easter Bunny. | 0.53 | 0.13 | 0.14 | 0.27 | 0.76 | 0.76 | 0.90 | 0.81 | 0.54 |
| B1: Twelfth century illustration of a man digging. B2: An English serf at work digging, c. 1170. | 0.13 | 0.13 | 0.10 | 0.12 | 0.51 | 0.47 | 0.83 | 0.60 | 0.48 |
| C1: Troops clearing rubble after the May air raid on Belfast. C2: Soldiers clearing rubble after the May air raid on Belfast. | 0.90 | 0.90 | 0.99 | 0.93 | 0.92 | 0.89 | 0.98 | 0.93 | 0.00 |
| D1: System of a Down is composed of four Armenian-Americans. D2: Dolmayan drumming with System of a Down in 2011. | 0.42 | 0.42 | 0.33 | 0.39 | 0.50 | 0.07 | 0.32 | 0.30 | -0.09 |

Table 5: Examples from the Wikipedia-IPC dataset with particularly high and low $\Delta_{\text{sem, syn}}$. Caption pairs with high $\Delta_{\text{sem, syn}}$ are structurally and lexically diverse while maintaining high levels of semantic similarity. Caption pairs with low $\Delta_{\text{sem, syn}}$ are either semantically dissimilar or have a high degree of lexical overlap.

“travel” to reach the embedded words of another text. We use the inversion of this measure as a Word Mover-based similarity score (WMS). All transformer-based models were applied with default hyperparameter configurations using a single NVIDIA A100 Tensor Core GPU. We normalize all the syntactic and semantic similarity measures to a scale from zero to one before we compute the average syntactic and semantic similarity.

Table 5 shows particularly sophisticated and unsophisticated caption pairs as paraphrases quantified by $\Delta_{\text{sem, syn}}$ from the Wikipedia-IPC (silver quality). Caption A1 is a paraphrase of A2 with a significantly different structure and sufficiently high lexical diversity. There is some unigram overlap, which causes the ROUGE-1 score to be comparatively high, since these captions are quite short. Due to the high semantic similarity of A1 and A2 these mutual paraphrases score high in the $\Delta_{\text{sem, syn}}$ measure. B1 and B2 have almost no words in common, but they are nevertheless semantically close. One semantic difference is that an “English serf” is a different concept than “a man” is, which could be the reason for the comparatively lower semantic similarity. However, the semantic similarity is high enough to yield a high $\Delta_{\text{sem, syn}}$. In (C1, C2) and (D1, D2) we see caption pairs that do not qualify as “good” paraphrases for two different reasons. In C2, a single word is replaced by a synonym in C1 and therefore has almost identical wording to C1. Although D1 and D2 are captions that both refer to the band “System of a Down”, they describe entirely different aspects, which is reflected in the semantic similarity scores. Therefore, both pairs (C1, C2) and (D1, D2) receive a low $\Delta_{\text{sem, syn}}$ value.

| Syntax: | ROUGE-1 | ROUGE-L | BLEU | Average |
|------------|---------|---------|------|---------|
| r | 0.78 | 0.77 | 0.70 | 0.79 |
| Semantics: | WMS | BERT | ST | Average |
| r | 0.59 | 0.70 | 0.78 | 0.76 |

Table 6: Top: Pearson correlation between manual judgements of syntactic similarity and the measures ROUGE-1, ROUGE-L, BLEU, and their average. Bottom: Pearson correlation between manual judgements of semantic similarity and the measures Word Mover Score (WMS), BERT-Score (BERT), Sentence Transformer (ST), and their average.

Manual Similarity Judgments In order to investigate whether the chosen similarity measures accurately quantify syntactic and semantic similarity, we conduct a manual annotation study to measure the correlation between the measures and human similarity ratings. For that purpose, we sampled 100 paraphrases from each of the three datasets—Wikipedia-IPC, MSRRC, and TaPaCo. We draw paraphrase examples with a length of 10 to 30 words stratified by their $\Delta_{\text{sem, syn}}$ distribution in their corresponding datasets. Four expert annotators with at least six years of experience rate semantic and syntactic similarity of all the 300 examples on a 5-point Likert scale.

Figure 2 shows the distribution of semantic and syntactic similarities of the median rating from the annotation study. These figures show that the semantic and syntactic similarities of the examples from the three different datasets are uniquely distributed. Examples from the TaPaCo corpus are pairs with subtle differences at the syntactic level that ultimately result in minor semantic differ-

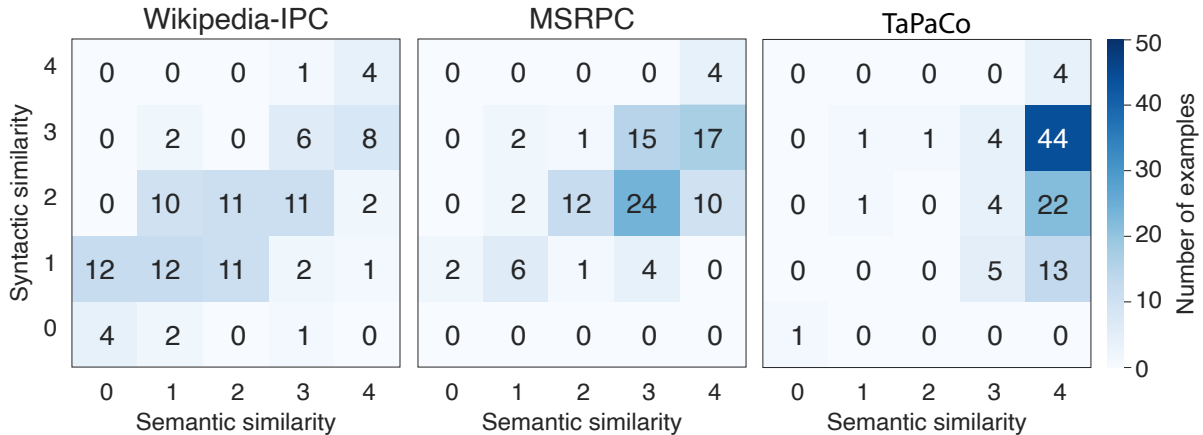


Figure 2: Distribution of semantic and syntactic similarity (characteristic map) for three datasets at sample size 100 for paraphrases with a length between 10 and 30 words. Among others, it can be seen that the TaPaCo corpus comprises paraphrases that result from subtle and minor changes within a paraphrase’s sentence pairs (44 x syn-4/sem-3), while Wikipedia-IPC shows a more uniform distribution along the diagonal syn-0/sem-0... syn-4/sem-4.

| Corpus | Corpus statistics | | | | Syntactic similarity | | | | | | | | Semantic similarity | | | | | | | | | | |
|---------------------------------|-------------------|-----------|--------|--------|----------------------|-------------|----------|-------------|----------|-------------|----------|-------------|---------------------|-------------|----------|-------------|----------|-------------|----------|-------------|--------------------|-------------|----------|
| | Subset | Type | Number | Length | ROUGE-1 | | ROUGE-L | | BLEU | | Avg. | | WMS | | BERT | | ST | | Avg. | | $\Delta_{sem,syn}$ | | |
| ST > 0.8 | | <i>n</i> | % | μ | σ | μ | σ | μ | σ | μ | σ | μ | σ | μ | σ | μ | σ | μ | σ | μ | σ | μ | σ |
| Wikipedia-IPC _{gold} | C | 10,327 | 34% | 17.97 | 9.97 | 0.74 | 0.18 | 0.71 | 0.20 | 0.56 | 0.28 | 0.67 | 0.21 | 0.83 | 0.12 | 0.69 | 0.19 | 0.91 | 0.06 | 0.81 | 0.11 | 0.14 | 0.13 |
| Wikipedia-IPC _{silver} | C | 105,475 | 29% | 20.38 | 13.56 | 0.71 | 0.19 | 0.67 | 0.22 | 0.52 | 0.28 | 0.63 | 0.22 | 0.63 | 0.20 | 0.81 | 0.12 | 0.90 | 0.06 | 0.78 | 0.12 | 0.15 | 0.12 |
| Flickr8k | C | 29,011 | 36% | 11.65 | 3.70 | 0.53 | 0.16 | 0.48 | 0.17 | 0.22 | 0.14 | 0.41 | 0.14 | 0.59 | 0.13 | 0.73 | 0.10 | 0.86 | 0.05 | 0.73 | 0.08 | 0.32 | 0.09 |
| PASCAL | C | 3,329 | 33% | 9.83 | 3.30 | 0.51 | 0.16 | 0.47 | 0.17 | 0.22 | 0.14 | 0.40 | 0.15 | 0.59 | 0.13 | 0.72 | 0.10 | 0.86 | 0.05 | 0.73 | 0.08 | 0.32 | 0.09 |
| MS-COCO | C | 392,248 | 32% | 10.54 | 2.25 | 0.51 | 0.16 | 0.45 | 0.16 | 0.22 | 0.13 | 0.39 | 0.14 | 0.57 | 0.14 | 0.71 | 0.10 | 0.86 | 0.04 | 0.71 | 0.08 | 0.32 | 0.09 |
| PAWS | G | 64,940 | 99% | 21.36 | 5.47 | 0.94 | 0.03 | 0.79 | 0.12 | 0.69 | 0.18 | 0.81 | 0.10 | 0.82 | 0.10 | 0.96 | 0.04 | 0.97 | 0.03 | 0.92 | 0.04 | 0.11 | 0.08 |
| ParaNMT-5m | G | 2,607,580 | 49% | 11.97 | 6.18 | 0.63 | 0.15 | 0.60 | 0.16 | 0.33 | 0.19 | 0.52 | 0.16 | 0.60 | 0.14 | 0.75 | 0.17 | 0.87 | 0.05 | 0.74 | 0.09 | 0.22 | 0.13 |
| MSRPC | H | 3,720 | 64% | 20.02 | 4.94 | 0.73 | 0.12 | 0.69 | 0.13 | 0.54 | 0.20 | 0.65 | 0.14 | 0.72 | 0.11 | 0.82 | 0.08 | 0.90 | 0.05 | 0.82 | 0.07 | 0.16 | 0.10 |
| PPDB 2.0 | H | 654,531 | 24% | 4.53 | 0.69 | 0.64 | 0.18 | 0.63 | 0.18 | 0.32 | 0.18 | 0.53 | 0.17 | 0.64 | 0.22 | 0.63 | 0.30 | 0.89 | 0.05 | 0.72 | 0.14 | 0.19 | 0.15 |
| TaPaCo | H | 117,447 | 52% | 5.47 | 2.28 | 0.65 | 0.18 | 0.63 | 0.18 | 0.30 | 0.16 | 0.53 | 0.16 | 0.78 | 0.14 | 0.79 | 0.21 | 0.91 | 0.06 | 0.83 | 0.10 | 0.30 | 0.13 |

Table 7: Comparison of syntactic similarity (measured as the average of ROUGE-1, ROUGE-L, and BLEU) and paraphrase semantic similarity (measured as the average of BERTScore, Sentence Transformer (ST) and Word Mover Score (WMS)) of semantically similar examples (Sentence Transformer > 0.8) from state-of-the-art paraphrase corpora, image caption corpora, and the new silver-quality Wikipedia-IPC. “Type” classifies the acquisition method in caption (C), automatically generated (G), and human-written (H) (e.g., acquired via crowd-sourcing).

ences. In contrast, examples from the Wikipedia-IPC (and MSRPC to some extent) are more evenly distributed along the diagonal syn-0/sem-0, ..., syn-4/sem-4. While Wikipedia-IPC has examples from all semantic similarity levels, semantic similarity is higher among examples from the MSRPC.

To justify the choice of the semantic and syntactic similarity measures, we compute Pearson correlation coefficients with manual similarity judgments. Table 6 shows high correlation values for syntactic and semantic similarity ratings, which renders our measures well suited for the given task.

Quantitative Similarity Analysis In the paraphrase literature, authors usually require pairs of texts to have varying degrees of semantic similarity

to consider them paraphrases of each other. The choice of a “semantic threshold” drastically affects the distribution of syntactic and semantic similarities of paraphrases within a dataset. To standardize the decision criteria for paraphrases, we manually set a semantic threshold that allows us to comparably evaluate paraphrases from each dataset.

We choose the semantic similarity score from the Sentence Transformer as the decision criterion for the subset of paraphrases under consideration because it has the highest correlation with human ratings of semantic similarity. To find a reliable threshold, we analyze the distributions of Sentence Transformer scores of paraphrases and non-paraphrases within the MSRPC and found that a threshold of 0.8

correctly classifies paraphrases with a precision and recall of 82% and 77%, respectively.

Table 7 shows the results of the comparison of syntactic and semantic similarities of paraphrase pairs with a high semantic similarity ($ST > 0.8$) grouped by type which is either image caption (C), automatically generated (G), or human-written (H). f calculates the frequency of text pairs that exceeds the semantic threshold of 0.8 relative to all examples in the dataset while $\Delta_{\text{sem,syn}}$ computes the average difference between the average semantic and syntactic similarity scores.

The high syntactic and semantic similarities of the examples from the PAWS dataset are striking. PAWS was designed as a benchmark corpus for paraphrase detection and contains lexically similar text pairs with subtle semantic inconsistencies that are difficult to distinguish from actual paraphrases. The high ROUGE-1 value of 94% with a low standard derivation of 0.03 translates to an almost identical wording. Hence, it is no surprise that the $\Delta_{\text{sem,syn}}$ is the lowest among all datasets.

From the gold and silver quality proportions of the Wikipedia-IPC, 34% and 29% of the examples exceed the semantic threshold, respectively. These rates are in line with those of the other caption corpora studied. Furthermore, we find that all caption datasets have similar overall syntactic and semantic similarities. However, the caption pairs from the Wikipedia-IPC tend to have higher syntactic similarity. We suspect that this observation is due to Wikipedia’s caption reuse policy, which allows editors to use slightly modified existing captions rather than inventing a new text. However, the sufficiently large $\Delta_{\text{sem,syn}}$ of the Wikipedia-IPC, which is on par with the commonly used MSRPC, and larger than that of the other caption corpora strongly suggests the suitability of image captions as a source of interesting paraphrases.

5 Conclusion and Future Work

We propose a new approach to use image captions as a resource for paraphrasing. From Wikipedia, we extracted 30,237 caption pairs in gold quality, 229,877 in silver quality, and 656,560 in bronze quality for our new Wikipedia-IPC dataset. As part of our analysis, we found that many caption pairs are “sophisticated” paraphrases in the sense of being semantically similar but dissimilar at the lexical and syntactic levels. We have introduced a respective measure for assessing paraphrase sophis-

tics based on semantic and syntactic similarity. The new measure correlates well with manual judgments, and we have shown that paraphrases from different sources have individual characteristics along the two similarity dimensions. Studying large sets of paraphrases with characteristics required by a specific task (e.g., sophisticated vs. only word order changed) will be useful for task-specific paraphrase recognition and generation approaches, and for natural language understanding as a whole.

In future work, we plan to extend the mining process to larger web resources than the Wikipedia and to apply image analysis as a new equivalence criterion. Identifying highly similar but not identical images may help to identify even more caption-based paraphrase candidates.

Limitations

Challenging for our proposed paraphrase acquisition approach are captions that describe the same image in very different contexts (e.g., maps). Such captions tend to focus on rather different aspects of the image (e.g., rivers vs. cities on a map) but since the captions often still share some commonalities, the semantic component of our proposed measure $\Delta_{\text{sem,syn}}$ might still classify the captions as similar. A “deeper” estimation of semantic similarity could help. With more research on image–text relationships, applying models such as CLIP (Radford et al., 2021) or Stable Diffusion (Rombach et al., 2022) to take image semantics into account might lead to improvements for our proposed paraphrase refinement measure.

References

- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Iñigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uria, and Janyce Wiebe. 2015. SemEval-2015 task 2: Semantic textual similarity, English, Spanish and pilot on interpretability. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 252–263, Denver, Colorado. Association for Computational Linguistics.
- Colin J. Bannard and Chris Callison-Burch. 2005. Paraphrasing with Bilingual Parallel Corpora. In *ACL 2005, 43rd Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 25-30 June 2005, University of Michigan, USA*, pages 597–604.
- Janina Bär. 2008. *Analyse zur automatischen*

- Klassifizierung von Bildern*. Ph.D. thesis, Technische Universität Chemnitz.
- Regina Barzilay and Kathleen McKeown. 2001. Extracting Paraphrases from a Parallel Corpus. In *Association for Computational Linguistic, 39th Annual Meeting and 10th Conference of the European Chapter, Proceedings of the Conference, July 9-11, 2001, Toulouse, France*, pages 50–57.
- Steven Burrows, Martin Potthast, and Benno Stein. 2013. Paraphrase Acquisition via Crowdsourcing and Machine Learning. *Transactions on Intelligent Systems and Technology (ACM TIST)*, 4(3):43:1–43:21.
- David Chen and William Dolan. 2011. Collecting Highly Parallel Data for Paraphrase Evaluation. In *Proceedings of the Forty-Ninth Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 190–200, Portland, Oregon. Association for Computational Linguistics.
- Joao Cordeiro, Gael Dias, and Pavel Brazdil. 2007. New Functions for Unsupervised Asymmetrical Paraphrase Detection. *Journal of Software*, 2(4):12–23.
- Mathias Creutz. 2018. Open Subtitles Paraphrase Corpus for Six Languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto. 2013. *Recognizing Textual Entailment: Models and Applications*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- William B. Dolan and Chris Brockett. 2005. Automatically Constructing a Corpus of Sentential Paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing*, pages 1–8, Jeju, South Korea. Kazuhide Yamamoto.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. 2013. Paraphrase-Driven Learning for Open Question Answering. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1608–1618, Sofia, Bulgaria. Association for Computational Linguistics.
- Juri Ganitkevitch. 2013. Large-Scale Paraphrasing for Natural Language Understanding. In *Proceedings of the 2013 NAACL HLT Student Research Workshop*, pages 62–68, Atlanta, Georgia. Association for Computational Linguistics.
- Juri Ganitkevitch and Chris Callison-Burch. 2014. The Multilingual Paraphrase Database. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4276–4283, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 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*, pages 758–764, Atlanta, Georgia. Association for Computational Linguistics.
- Hamid Izadinia, Fereshteh Sadeghi, Santosh K. Divvala, Hannaneh Hajishirzi, Yejin Choi, and Ali Farhadi. 2015. Segment-Phrase Table for Semantic Segmentation, Visual Entailment and Paraphrasing. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From Word Embeddings To Document Distances. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 957–966, Lille, France. PMLR.
- Wuwei Lan, Siyu Qiu, Hua He, and Wei Xu. 2017. A Continuously Growing Dataset of Sentential Paraphrases. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1224–1234, Copenhagen, Denmark. Association for Computational Linguistics.
- Tao Li and Vivek Srikumar. 2016. Exploiting Sentence Similarities for Better Alignments. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2193–2203, Austin, Texas. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. A SICK cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Ramtin Mehdi-zadeh Seraj, Maryam Siahbani, and Anoop Sarkar. 2015. Improving Statistical Machine

- Translation with a Multilingual Paraphrase Database. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1379–1390, Lisbon, Portugal. Association for Computational Linguistics.
- Tong Niu, Semih Yavuz, Yingbo Zhou, Nitish Shirish Keskar, Huan Wang, and Caiming Xiong. 2021. Unsupervised Paraphrasing with Pretrained Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5136–5150, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Aaditya Prakash, Sadid A. Hasan, Kathy Lee, Vivek Datla, Ashequl Qadir, Joey Liu, and Oladimeji Farri. 2016. Neural Paraphrase Generation with Stacked Residual LSTM Networks. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2923–2934, Osaka, Japan. The COLING 2016 Organizing Committee.
- Lihua Qian, Lin Qiu, Weinan Zhang, Xin Jiang, and Yong Yu. 2019. Exploring Diverse Expressions for Paraphrase Generation. 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)*, pages 3173–3182, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Cyrus Rashtchian, Peter Young, Micah Hodosh, and Julia Hockenmaier. 2010. Collecting Image Annotations Using Amazon’s Mechanical Turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 139–147, Los Angeles. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. 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)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695.
- Yves Scherrer. 2020. TaPaCo: A Corpus of Sentential Paraphrases for 73 Languages. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6868–6873, Marseille, France. European Language Resources Association.
- Benno Stein, Matthias Hagen, and Christof Bräutigam. 2014. Generating Acrostics via Paraphrasing and Heuristic Search. In *25th International Conference on Computational Linguistics (COLING 2014)*, pages 2018–2029. Association for Computational Linguistics.
- Ann Taylor, Mitchell Marcus, and Beatrice Santorini. 2003. *The Penn Treebank: An Overview*, pages 5–22. Springer Netherlands, Dordrecht.
- Jan Philip Wahle, Terry Ruas, Frederic Kirstein, and Bela Gipp. 2022. How Large Language Models are Transforming Machine-Paraphrase Plagiarism. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 952–963, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 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)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- Wei Xu, Alan Ritter, Chris Callison-Burch, William B. Dolan, and Yangfeng Ji. 2014. Extracting Lexically Divergent Paraphrases from Twitter. *Transactions of the Association for Computational Linguistics*, 2:435–448.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019a. BERTScore: Evaluating Text Generation with BERT. *CoRR*, abs/1904.09675.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019b. PAWS: Paraphrase Adversaries from Word Scrambling. 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)*, pages 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.