

Conciseness: An Overlooked Language Task

Felix Stahlberg and Aashish Kumar and Chris Alberti and Shankar Kumar

Google Research

{fstahlberg,kumaraashish,chrisalberti,shankarkumar}@google.com

Abstract

We report on novel investigations into training models that make sentences concise. We define the task and show that it is different from related tasks such as summarization and simplification. For evaluation, we release two test sets, consisting of 2000 sentences each, that were annotated by two and five human annotators, respectively. We demonstrate that conciseness is a difficult task for which zero-shot setups with large neural language models often do not perform well. Given the limitations of these approaches, we propose a synthetic data generation method based on round-trip translations. Using this data to either train Transformers from scratch or fine-tune T5 models yields our strongest baselines that can be further improved by fine-tuning on an artificial conciseness dataset that we derived from multi-annotator machine translation test sets.

1 Introduction

“Vigorous writing is concise. A sentence should contain no unnecessary words, a paragraph no unnecessary sentences, for the same reason that a drawing should have no unnecessary lines and a machine no unnecessary parts.”

Strunk and White (1918)
The Elements of Style

Conciseness is a writing principle of removing redundant information in text. Even though conciseness is highly valued in expository English writing and is often considered good writing style (Brock and Walters, 1992; Zinsser, 2016), it is still an understudied topic in the natural language processing (NLP) community, mainly due to the lack of annotated data sets. However, automatic methods for improving conciseness have the potential to improve the writing experience even for native speakers, or to provide useful tools for editorial

tasks. In this work we take initial steps towards conciseness from an NLP perspective. We release¹ two hand-annotated test sets for conciseness – *Concise-Lite* (2-way annotated) and *Concise-Full* (5-way annotated). *Concise-Lite* annotators were asked to make minimal changes to the original sentence, whereas *Concise-Full* annotators were given the option to make larger rewrites. Table 1 contains examples from both test sets. For evaluation, we compute $F_{0.5}$ -scores of edit spans, a metric that is also commonly used for grammatical error correction (GEC) (Dahlmeier and Ng, 2012; Felice et al., 2016; Bryant et al., 2017). Given that both the test sets and the evaluation tool we employ are publicly available, we hope our setup will encourage NLP researchers to investigate models for conciseness.

We evaluate a range of models on our newly collected conciseness test sets. Our initial approach follows the recent paradigm of using massively pre-trained neural models with either no or very little task-specific training data. Inspired by Brown et al. (2020) we report on zero-shot experiments with the large language model LaMDA (Thoppilan et al., 2022). We also fine-tune the large sequence model T5 (Raffel et al., 2020) on small conciseness data sets. We achieve our best results using an unsupervised synthetic data generation method based on round-trip translations, i.e. sentence pairs that were generated by translating an English sentence into another language (e.g. German) and back, a technique that was previously proposed for GEC pre-training (Lichtarge et al., 2019). We construct additional data sets by creating mappings from the longest to the shortest reference in multi-reference machine translation (MT) test sets. Our experiments suggest that conciseness is a hard task for current NLP models. We conclude with a thorough investigation into the similarities and differences of our systems and map out the challenges ahead.

¹<https://github.com/google-research-datasets/wiki-conciseness-dataset>

Input sentence	Concise-Lite	Concise-Full
Gemco had a version called Memco, also owned by Lucky Stores, that operated stores in the Chicago and Washington, D.C., areas.	Gemco had a version called Memco, owned by Lucky Stores, operating stores in the Chicago and Washington, D.C.	Memco was a version of Gemco operated by Lucky Stores in Chicago and Washington, D.C.
The film was adapted from a best-selling biography of the brothers, and was well presented and well received.	The film was adapted from a best-selling biography of the brothers, and was well presented and received .	The film, adapted from the brothers' best-selling biography, was well presented and received .

Table 1: Example sentences from our *Concise-Lite* and *Concise-Full* test sets.

Input sentence	Abstractive sentence summarization	Conciseness model output
Exxon corp. and Mobil corp. have held discussions about combining their business operations, a person involved in the talks said Wednesday.	Exxon and Mobil discuss combining business operations; possible merger.	Exxon Corp. and Mobil Corp. have discussed combining their business operations, a person involved in the talks said Wednesday.
Chuck Knoblauch and Tino Martinez were as popular as squeegee men a week ago, the speculation rampant that one or the other or both might be exiled if the Yankees' historic year crumbled in the post-season.	Knoblauch and Martinez home run hits cinch Yankee's First World Series game	Chuck Knoblauch and Tino Martinez were as popular as squeegee men a week ago, the speculation rampant that either or both could be exiled if the Yankees' historic year crumbled in the postseason .

Table 2: Example outputs of one of our conciseness models on sentences from an abstractive sentence summarization data set (Over et al., 2007, DUC2004).

Input sentence	Sentence simplification	Conciseness model output
A mutant is a type of fictional character that appears in comic books published by Marvel comics.	A mutant is a form of imaginary character that is seen in comic books published by Marvel comics.	A mutant is a fictional character that appears in comics published by Marvel comics.
It will then dislodge itself and sink back to the river bed in order to digest its food and wait for its next meal.	It will then get away from its place and sink back into the river bed in order to digest its food and wait for its next meal.	It will then dislodge and return to the riverbed to digest its food and wait for the next meal.

Table 3: Example outputs of one of our conciseness models on sentences from a text simplification data set (Zhang and Lapata, 2017, WikiLarge).

2 The conciseness task

In this work we define the conciseness task as *applying the required edits to make a sentence less wordy without changing its meaning, intent or sentiment*. We will shed more light on the limitations of this definition in Sec. 6. We expect conciseness models to be useful mainly for native or advanced non-native writers who wish to improve their writing style. Conciseness is related to several other NLP tasks, but we argue below that each of these tasks has a different focus and deserves an independent treatment.

Summarization and sentence compression

Abstractive sentence summarization (Over et al., 2007) attempts to produce a condensed version of the input text. Summaries are similar to headlines with a maximum length that is independent of the input sentence length (Rush et al., 2015). Thus, generating a summary often requires a much more severe compression compared to conciseness.

Unlike summarization, conciseness is faithful to the input and aims to avoid the loss of any information – the goal is to generate a shorter sentence that can replace the original sentence within continuous text (see Table 2 for examples). Furthermore, most work on summarization focuses on the compression of entire documents or paragraphs (Zhang et al., 2020) and not on single sentences.

Similarly to sentence summarization, *sentence compression* also aims to generate a shorter version of the input text. Many sentence compression models only allow the deletion of words without the ability to rephrase parts of the sentence (Knight and Marcu, 2000; Jing, 2000; Filippova et al., 2015). Perhaps closest to our work, Mallinson et al. (2018) trained sentence compression models on round-trip translations and thereby avoided this restriction. The main difference to us is that we evaluate a broader range of methods on human-annotated test sets which we release for future research.

Sentence simplification The task of reducing the linguistic complexity of text to improve readability is known as *sentence simplification* (Sagion, 2017). It can be subdivided into lexical (e.g. replacing uncommon words with synonyms) and syntactic (e.g. changing passive to active) simplification (Devlin, 1999; Carroll et al., 1999). Most forms of syntactic simplification result in concise outputs,² but lexical simplification may yield even more verbose outputs. For example, replacing ‘to portray’ with a simpler but verbose phrase such as ‘to describe very vividly’ would be an instance of lexical simplification but not of conciseness. Conversely, a conciseness system may substitute a phrase with another that is concise but less common and thereby deteriorate readability. Another difference is that simplification often targets people with cognitive disabilities (Devlin, 1999; Carroll et al., 1999; Rello et al., 2013) or low literacy (Watanabe et al., 2009) or second language learners (Petersen and Ostendorf, 2007; Siddharthan, 2002; Xia et al., 2016) whereas conciseness can be thought as writing assistance for proficient writers. Table 3 contrasts simplification and conciseness with the help of example sentences.

Style transfer Text style is an important consideration for several NLP tasks (Fu et al., 2018). For example, it is desirable for MT output to match the stylistic properties of the source sentence (Senrich et al., 2016; Lohar et al., 2017). Natural language generation systems not only need to take into account the content of generated utterances but also other attributes such as style and sentiment (Li et al., 2018). Text-to-text style transfer systems have been used to change Shakespearean English to modern English (Jhamtani et al., 2017). We consider conciseness as a special case of style transfer with a single source style (wordy) and one target style (concise). However, while most style transfer systems attempt to change attributes like sentiment or political slant (Li et al., 2018; Fu et al., 2018; Prabhumoye et al., 2018; Shen et al., 2017), our conciseness models aim to keep them unchanged.

Paraphrasing Paraphrasing databases such as PPDB (Ganitkevitch et al., 2013; Pavlick et al., 2015) that store pairs of phrases with the same meaning have proven useful for various NLP tasks such as textual entailment (Bjerva et al., 2014) and

²An exception would be *sentence splitting* since it is a syntactic simplification strategy that often makes the text longer.

semantic similarity (Han et al., 2013). In this work we include a paraphrasing system for comparison.

3 Modeling conciseness

The approaches in this section cover a wide range of NLP models to convey a better sense for the task. They are intended to serve as baselines to compare against, and as a starting point for future research.

3.1 Giant language models (LaMDA)

Large language models (LMs) such as OpenAI’s GPT-3 (Radford et al., 2019), Google’s Meena (Adiwardana et al., 2020) and PaLM (Chowdhery et al., 2022) and Microsoft’s Turing NLG³ have recently captured the interest of the general public through their ability to generate text that is sometimes astonishingly difficult to distinguish from text written by humans. While these models are useful for building open-domain dialog agents, they also have the potential to solve specific NLP problems when provided with an appropriate preamble (LM history) (Brown et al., 2020). We expect general dialog agents to understand the nuances of language such as grammar, conciseness, etc. Thus, we explored using the large LM LaMDA (Thoppilan et al., 2022) with a zero-shot preamble that steers the model towards making a sentence more concise. We use the following template to provide the LM context:

Here is some text:
“[INPUT_SENTENCE]”. Rewrite it to be more concise.

where [INPUT_SENTENCE] is replaced by the source sentence.⁴ We post-process the output to a) discard any additional comment that the model generated besides the rewrite, and b) retain only the first suggestion if multiple rewrites are generated.

3.2 Transformers pre-trained on round-trip translations

This method employs synthetic training data generated using MT. Fig. 1 illustrates the approach. First, we translate an English sentence into a pivot language such as German, and then translate it back

³<https://msturing.org/>

⁴This prompt was best among a small number of zero-shot and few-shot prompts we explored. Systematic prompt engineering could potentially improve LaMDA results at a significantly higher computational cost, but we have not explored this option in this work since we focus on conciseness as an NLP task.

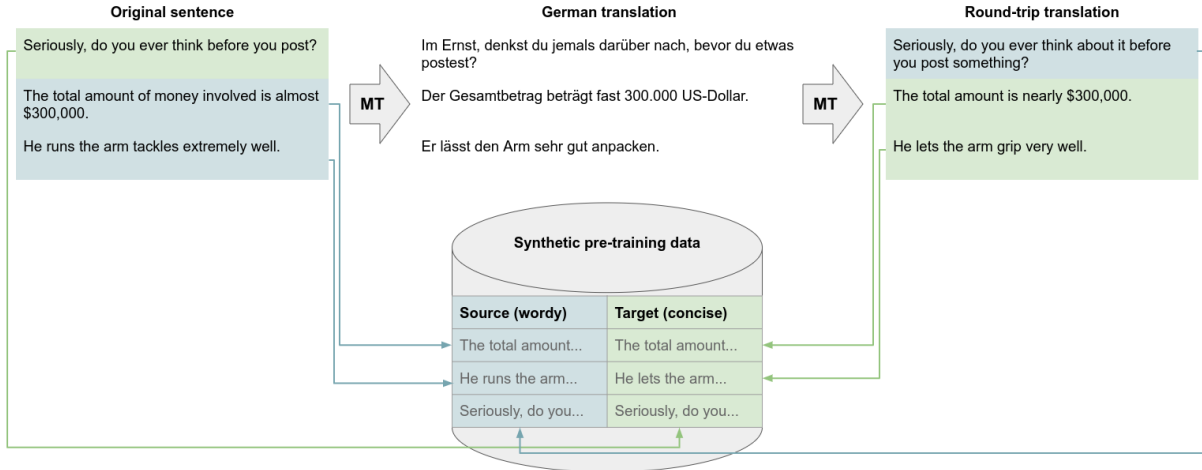


Figure 1: Synthetic pre-training data generation using round-trip translations.

Name	Number of sentence pairs	Average source sentence length in words	Average target sentence length in words	Compression ratio
Pre-training and fine-tuning data sets				
RoundTrip-French	169M	20.6	19.4	0.94
RoundTrip-German	169M	20.4	19.4	0.95
RoundTrip-Japanese	169M	20.4	17.9	0.88
RoundTrip-Russian	169M	20.9	19.5	0.93
MultiRefMT-FineTune	9K	31.9	26.1	0.82
Development sets				
MultiRefMT-Dev	820	33.3	25.8	0.77
Hand-annotated test sets				
Concise-Lite	2K	23.7	21.2	0.89
Concise-Full	2K	23.7	20.1	0.85

Table 4: Data set statistics. The compression ratio is the number of target words divided by the number of source words.

into English. This idea of generating sentence pairs via round-trip translation was initially proposed by Lichtarge et al. (2019) to pre-train GEC systems. In this work, we construct synthetic parallel data for conciseness by using the longer sentence as the source and the shorter sentence as the target sentence. We then train a standard neural sequence-to-sequence Transformer (Vaswani et al., 2017) on the synthetic data until convergence.⁵ This approach is simple and enables us to generate large quantities of data, but the resulting data set contains noise. For example, round-trip translation pairs often contain synonym substitutions (see the replacement of *almost* with *nearly* in the second sentence in Fig. 1) that do not help conciseness. Furthermore, MT may fail to translate the sentence properly, resulting in an undesirable change of meaning (see the third sentence in Fig. 1). Another problem is that it is hard to control the compression ratio in the data set. Despite these limitations we show in Sec. 5 that

⁵More details about the Transformer model implementation are provided in Appendix A.

round-trip translations are useful for pre-training.

3.3 Fine-tuning T5

The final method considered in this work employs T5 (Raffel et al., 2020). Very large sequence-to-sequence models have been found to be extremely powerful, even for challenging language tasks with a limited amount of training data. We fine-tuned the publicly available 11B parameter version (xxl) of T5⁶, with a batch size of 1,024 sentences and a learning rate of 10^{-4} .

4 Data sets

Table 4 lists the data sets used in this work. Table 5 contains information about their provenance.

Round-trip translations (RoundTrip-*) Our Transformer system is pre-trained on round-trip translations of sentences crawled from news websites following the recipe of Lichtarge et al. (2019)

⁶https://github.com/google-research/text-to-text-transfer-transformer/blob/main/released_checkpoints.md

Name	Reference	Type
RoundTrip-*	Lichtarge et al. (2019)	Round-trip translations (news)
MultiRefMT-FineTune	LDC2010T10, LDC2010T11, LDC2010T12, LDC2010T14	4-annotator MT test sets (Arabic-English, Chinese-English)
MultiRefMT-Dev	LDC2013T03	4-annotator MT test set (Chinese-English)
Concise-Lite	This work	2-way hand-annotated conciseness test set
Concise-Full	This work	5-way hand-annotated conciseness test set

Table 5: Synthetic and hand-annotated conciseness data sets used in this work.

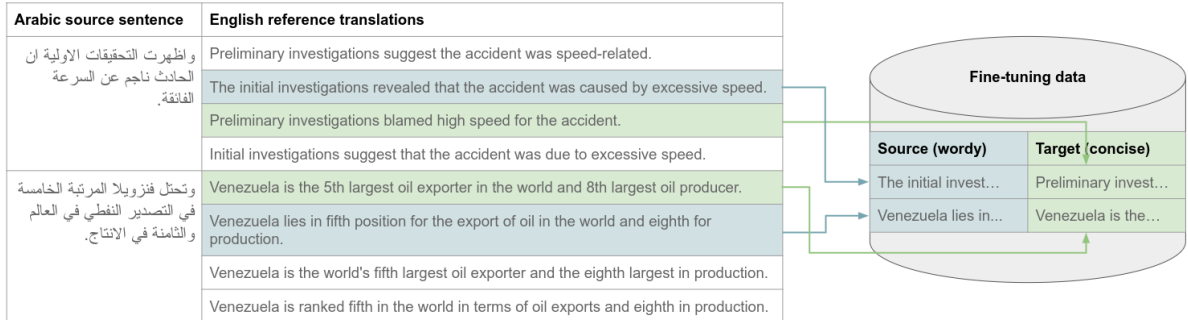


Figure 2: Fine-tuning data generation using multi-reference MT test sets.

that were prepared as described in Sec. 3.2. For fine-tuning T5 on round-trip translations we randomly sample 1M sentence pairs from the full data set to limit computation.

OpenMT-based fine-tuning and development sets (MultiRefMT-*) We derive fine-tuning and development sets from existing publicly available MT test sets. It is common practice in several NLP areas to collect reference sentences from multiple annotators to increase the trustworthiness of automatic evaluation measures, for example in grammatical error correction (Ng et al., 2014; Bryant and Ng, 2015; Napoles et al., 2017), MT (Fretag et al., 2020), and image caption generation (Zheng et al., 2018). Multi-reference MT test sets have been used in the past to evaluate paraphrasing or sentence compression systems (Ganitkevitch et al., 2011; Pang et al., 2003). We make use of these multi-annotator test sets by selecting the longest reference sentence as the (wordy) source sentence and the shortest reference sentence as the golden (concise) target sentence (Fig. 2). Our MultiRefMT-FineTune set uses all Arabic-English and Chinese-English NIST Open Machine Translation (OpenMT) evaluation sets from 2002-2005. The MultiRefMT-Dev set is based on the Chinese-English 2012 OpenMT evaluation set.

Hand-annotated test sets (Concise-*) Deriving conciseness test sets from multi-reference MT evaluation sets is viable as a first approximation given

that all references have similar meaning, intent, and sentiment by design (apart from annotation errors). However, it does not allow us to determine how wordy the sentence is in the first place. If all MT references agreed, it would suggest that the original source sentence has a single obvious translation, not that the references are already concise.

Therefore, we collected two new data sets, consisting of 2000 sentences each, that were explicitly annotated for conciseness – *Concise-Lite* and *Concise-Full*. Both data sets used the same set of source sentences drawn from Wikipedia. Sentences that a) were ungrammatical, b) contained fewer than 15 words or c) included mismatched quotation marks were not selected. While *Concise-Lite* annotators were asked to make minimal changes to the original sentence, *Concise-Full* annotators were given the flexibility to make larger changes to the original sentence. The exact annotator guidelines are listed in Appendix B.

We will make the test sets publicly available to establish a benchmark for researchers to evaluate conciseness models.

5 Results

We use the GEC evaluation toolkit ERRANT (Bryant et al., 2017; Felice et al., 2016) to compute $F_{0.5}$ -scores on spaCy⁷-tokenized text. Like in GEC, precision is weighted twice as high as recall using the $F_{0.5}$ -score, which matches our intuition

⁷<https://spacy.io/>

System	Concise-Lite			Concise-Full		
	P	R	$F_{0.5}$	P	R	$F_{0.5}$
Other NLP tasks						
a Summarization: Pegasus	0.8	1.4	0.9	2.0	3.9	2.2
b Summarization: Long-T5	1.7	6.3	2.0	3.5	11.7	4.1
c Simplification: T5	7.4	5.4	6.9	13.8	9.9	12.8
d Paraphrasing: ParaNMT	9.3	21.4	10.4	15.4	25.1	16.7
Conciseness models						
e Giant-LM (zero-shot LaMDA)	4.4	13.5	5.1	8.5	20.0	9.6
f Transformer (RT)	13.6	21.3	14.6	21.1	25.5	21.9
g Transformer (RT→MT)	15.0	25.8	16.4	24.4	29.6	25.2
h T5 (RT)	18.4	19.5	18.6	29.1	24.2	28.0
i T5 (RT→MT)	16.0	26.8	17.4	26.6	30.6	27.3

Table 6: System comparison on our two conciseness test sets. “RT” denotes models trained on round-trip translations. “RT→MT” configurations are subsequently fine-tuned on MultiRefMT-FineTune.

System	Number of parameters
Giant-LM (LaMDA)	137B
T5	11B
Transformer	313M

Table 7: Number of model parameters.




Figure 3: Transformer models trained from scratch on round-trip translations via different pivot languages.

that a conciseness system should act as a minimally intrusive writing assistant for which false positives are far worse than false negatives.

5.1 System comparison

Table 6 compares all approaches from Sec. 3 and the following baselines from other NLP tasks:

- Summarization: Long-T5 (Guo et al., 2022) and Pegasus (Zhang et al., 2020).
- Simplification: T5 fine-tuned on the Wiki-Large simplification dataset (Zhang and Lapata, 2017) using a procedure similar to our T5-conciseness system from Sec. 3.3.⁸
- Paraphrasing: A Transformer model trained on the full ParaNMT-50M (Wieting and

⁸Our simplification baseline achieves 33.1 SARI on the WikiLarge test set.

Gimpel, 2018) training set using the hyperparameters in Appendix A.

The summarization baselines (rows a and b) perform poorly since they are mostly trained on full documents. The simplification system achieves a slightly higher performance but is weaker than the paraphrasing or the Transformer/T5 based conciseness systems. The paraphrasing system (row d) achieved a recall of over 20% on both test sets, but the precision is relatively low because the ParaNMT training set contains various types of edits such as synonym replacements or word reorderings that do not necessarily help conciseness.

The zero shot Giant-LM (LaMDA) setup (row e) was not able to match either the precision or recall of the other conciseness systems. Round-trip translations are useful for both training a Transformer model from scratch (row f) and fine-tuning T5 (row h). Subsequent fine-tuning on MultiRefMT-FineTune yields large precision and recall gains for the Transformer model (row g). MultiRefMT-FineTune also improves the recall for T5, but the precision suffers (row i).⁹ T5 outperforms the Transformers in terms of $F_{0.5}$ -score by achieving higher precision on both sets but has many more parameters (Table 7).

5.2 Ablation studies and analyses

The following analyses were carried out on the *Concise-Lite* and *Concise-Full* test sets.

Round-trip translation languages Our final models in Table 6 use round-trip translations from four different pivot languages: French, German,

⁹T5 is fine-tuned for 4K steps on the 1M round-trip translations and for 1K steps on the smaller MultiRefMT-FineTune set.

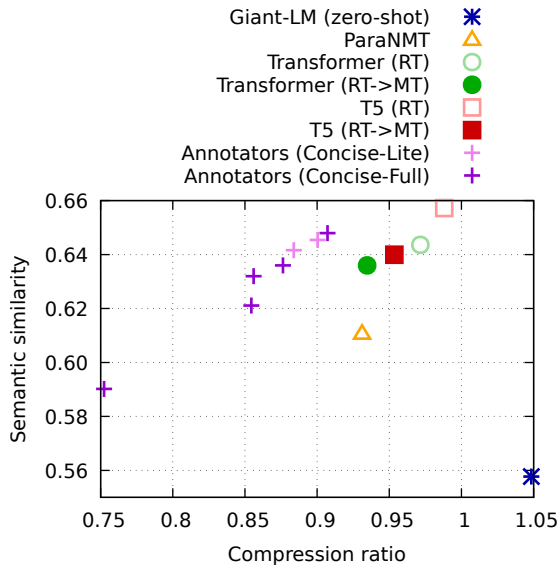


Figure 4: Trade-off between semantic similarity and the sentence compression ratio.

Japanese, and Russian. Fig. 3 shows that combining all languages yields consistent gains on both test sets over using any single language.

Preserving semantics To measure how well our systems retain the meaning of the original sentence we computed semantic similarity scores between the input and the output sentences using the models provided by the Semantic Reactor toolkit (Yang et al., 2018; Cer et al., 2018). Systems and annotators trade off compression against semantic similarity differently (Figure 4). There is a large variability in compression ratio (i.e. the number of target words divided by the number of source words) and semantic similarity between the *Concise-Full* annotators (dark purple). The Giant-LM (blue) is more prone to meaning change than other systems, and is not effective in reducing the sentence length. Fine-tuning on MultiRefMT-FineTune (empty vs. filled circle/square) improves the compression ratio but hurts semantic similarity. T5 (red) preserves semantics better than the Transformer but outputs slightly longer sentences.

Readability Fig. 5 shows that our systems often improve the readability of the sentence, in particular the Giant-LM system. The Giant-LM prefers simpler language as it was originally designed for dialog applications (Thoppilan et al., 2022). In contrast, the *Concise-Full* annotators tend to achieve concision using longer and more complex words, resulting in a decline in readability (dark purple).

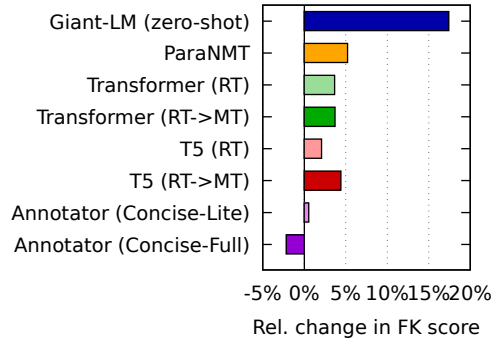


Figure 5: Relative change in Flesch-Kincaid readability scores (Kincaid et al., 1975).




Figure 6: Relative change in information density.

Information density We expect the outputs of a high-performing conciseness system to have a high information content per word. This information density can be measured using per-token inverse document frequency (Jones, 1973):

$$\text{idf}(t) = \log \frac{N}{|\{d \in D : t \in d\}|},$$

where t is the token, N is the total number of documents, and D is the document collection. In our case, the document frequencies are derived from the C4 corpus (Raffel et al., 2020). Fig. 6 shows that the reference sentences from the *Concise-Lite* and *Concise-Full* annotators indeed have a higher per-token IDF than the input sentences (pink and dark purple bars). The results on the system outputs are mixed, but fine-tuning on MultiRefMT-FineTune improves the per-token IDF for the Transformer and T5 (“RT” vs. “RT → MT”).

Synonym substitutions One problem with using round-trip translations for training and multi-reference test sets for evaluation is that both may contain synonym substitutions that do not help conciseness. We counted synonym substitutions by extracting all 1:1 substitutions and checking

	Without A1			Without A2			Without A3			Without A4			Without A5		
	P	R	$F_{0.5}$	P	R	$F_{0.5}$	P	R	$F_{0.5}$	P	R	$F_{0.5}$	P	R	$F_{0.5}$
Annotator A1	45.8	52.0	46.9												
Annotator A2				16.3	32.0	18.1									
Annotator A3							51.5	48.4	50.9						
Annotator A4										23.1	32.6	24.5			
Annotator A5													33.5	27.1	32.0
Transformer	22.7	27.6	23.5	19.7	28.6	21.0	23.6	27.7	24.3	20.8	26.8	21.8	23.0	25.9	23.6
T5	25.3	28.9	26.0	20.7	29.2	22.0	25.7	28.9	26.3	23.1	28.0	23.9	25.4	27.0	25.7

Table 8: Measuring annotator agreement on *Concise-Full* by evaluating each single annotator using the other four annotations as references. We list the Transformer and T5 system outputs (“RT→MT”) for comparison.

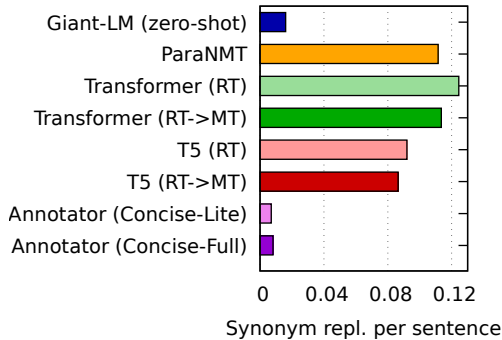


Figure 7: Number of 1:1 synonym substitutions.

whether these were marked as synonyms in WordNet (Miller, 1995). Fig. 7 shows that most of our systems replace synonyms on an average in every 10th sentence. Fine-tuning the Transformer or T5 on MultiRefMT-FineTune reduces the number of synonym substitutions. Synonyms are much less of a problem with the Giant-LM (blue bar) which was not trained on round-trip translations.

6 Limitations

In terms of both information density (Fig. 6) and number of unnecessary synonym replacements (Fig. 7), the annotators are clearly separated from most of our automatic systems, illustrating the gap to human performance on this task.

Our experiments showed that the Giant-LM (zero-shot) underperformed the other approaches. Preliminary experiments using few-shot learning did not yield improvements over the zero-shot setting. We expect the performance of Giant-LM to improve via systematic prompt engineering.

Another challenge lies in the intrinsic uncertainty (Ott et al., 2018; Stahlberg et al., 2022) of the conciseness task, i.e. the existence of multiple viable ways to make a sentence more concise. Table 8 demonstrates that the five *Concise-Full* annotators usually did not agree on a single concise

version of a sentence, leading to great variability in $F_{0.5}$ -scores when evaluated against each other.¹⁰ Therefore, adequate system outputs may get penalized if they do not agree with one of the human references. We mitigate this concern by using multiple annotators, but – like in other intrinsically uncertain NLP tasks such as MT – a certain level of noise remains in our evaluation.

Limitations of our task definition We acknowledge that there are various aspects of conciseness that are not covered by our definition in Sec. 2 (“applying the required edits to make a sentence less wordy without changing its meaning, intent or sentiment”). First, we intentionally did not include the use of context in our definition. In practice, however, appropriate levels of conciseness can be highly context dependent. Treating the problem on the sentence-level is limiting because using inter-sentential cross-references for conciseness requires access to the document-level context such as the previous sentence. Furthermore, the sentence-level restriction prevents the systems from improving conciseness through sentence splitting (Botha et al., 2018) or merging (Geva et al., 2019). In real-life situations, the context may also be provided through other channels such as physical medium (e.g. pointing to things) or social factors (e.g. does person B know person A?). We also noticed that our *Concise-Full* annotators occasionally relied on common knowledge to shorten sentences (see Appendix C for examples), a strategy that is *not* covered by our definition and thus makes our evaluation slightly more noisy. Exploring the various forms of context for conciseness is a promising potential direction for future research.

Another limitation of our definition is that it does

¹⁰On some of the setups in Table 8 (e.g. “Without A2” or “Without A4”), T5 achieves scores comparable to the human annotators. We emphasize that this is a sign of low inter-annotator agreement and does not allow us to claim human parity since this pattern is not consistent across annotators.

not allow for a change of semantics, intent, or sentiment. In practice, however, conciseness or the lack of it may reflect the intent of the speaker, for example in indicating emergency situations (signalling urgency through brevity) or in detecting lying (Vrij, 2005). Another manner in which conciseness can carry meaning is when used as a rhetorical device to persuade or inspire the audience, a well-known strategy in legal writing (Osbeck, 2011) that was perhaps most famously demonstrated by Abraham Lincoln in the Gettysburg Address (Oseid, 2009). Furthermore, our ablation studies in Sec. 5.2 revealed that systems and human annotators alike sometimes accepted a minor loss of (irrelevant) information to achieve better compression, which, despite being contrary to our definition, may be acceptable in practice.

7 Conclusion

Our work is an initial exploration of conciseness from an NLP point of view. We compared a variety of approaches to the problem using popular techniques based on synthetic data generation or giant pre-trained sequence models. Round-trip translations provide a useful data source for training conciseness models but can introduce undesirable synonym substitutions.¹¹ Our analyses show that our systems trade off the objectives in conciseness differently (e.g. reducing the sentence length vs. preserving semantics vs. improving readability vs. increasing information density). Further experiments are necessary to understand how these trade-offs would impact the user experience or potential downstream NLP tasks. We expect our study and our annotated test sets to provide impetus for researchers to explore this field further.

References

- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- Johannes Bjerva, Johan Bos, Rob van der Goot, and Malvina Nissim. 2014. [The meaning factory: Formal semantics for recognizing textual entailment and determining semantic similarity](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 642–646, Dublin, Ireland. Association for Computational Linguistics.
- Jan A. Botha, Manaal Faruqui, John Alex, Jason Baldridge, and Dipanjan Das. 2018. [Learning to split and rephrase from Wikipedia edit history](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 732–737, Brussels, Belgium. Association for Computational Linguistics.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2021. [JAX: composable transformations of Python+NumPy programs](#).
- Mark Newell Brock and Larry Walters. 1992. *Teaching composition around the pacific rim: Politics and pedagogy*, volume 88. Multilingual matters.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. [Automatic annotation and evaluation of error types for grammatical error correction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 793–805, Vancouver, Canada. Association for Computational Linguistics.
- Christopher Bryant and Hwee Tou Ng. 2015. [How far are we from fully automatic high quality grammatical error correction?](#) 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)*, pages 697–707, Beijing, China. Association for Computational Linguistics.
- John Carroll, Guido Minnen, Darren Pearce, Yvonne Canning, Siobhan Devlin, and John Tait. 1999. [Simplifying text for language-impaired readers](#). In *Ninth Conference of the European Chapter of the Association for Computational Linguistics*, pages 269–270, Bergen, Norway. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Céspedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng

¹¹Appendix C illustrates the strengths and weaknesses of our current systems with the help of some example outputs.

- Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#). *CoRR*.
- Daniel Dahlmeier and Hwee Tou Ng. 2012. [Better evaluation for grammatical error correction](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–572, Montréal, Canada. Association for Computational Linguistics.
- Siobhan Devlin. 1999. *Simplifying natural language text for aphasic readers*. Ph.D. thesis, Ph. D. thesis, University of Sunderland, UK.
- Mariano Felice, Christopher Bryant, and Ted Briscoe. 2016. [Automatic extraction of learner errors in ESL sentences using linguistically enhanced alignments](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 825–835, Osaka, Japan. The COLING 2016 Organizing Committee.
- Katja Filippova, Enrique Alfonseca, Carlos A. Colmenares, Lukasz Kaiser, and Oriol Vinyals. 2015. [Sentence compression by deletion with LSTMs](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 360–368, Lisbon, Portugal. Association for Computational Linguistics.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. [BLEU might be guilty but references are not innocent](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 61–71, Online. Association for Computational Linguistics.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. [Style transfer in text: Exploration and evaluation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Juri Ganitkevitch, Chris Callison-Burch, Courtney Napoles, and Benjamin Van Durme. 2011. [Learning sentential paraphrases from bilingual parallel corpora for text-to-text generation](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1168–1179, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- 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.
- Mor Geva, Eric Malmi, Idan Szpektor, and Jonathan Berant. 2019. [DiscoFuse: A large-scale dataset for discourse-based sentence fusion](#). 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 3443–3455, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontañón, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2022. [LongT5: Efficient text-to-text transformer for long sequences](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 724–736.
- Lushan Han, Abhay L. Kashyap, Tim Finin, James Mayfield, and Jonathan Weese. 2013. [UMBC_EBIQUITY-CORE: Semantic textual similarity systems](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 44–52, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Harsh Jhamtani, Varun Gangal, Eduard Hovy, and Eric Nyberg. 2017. [Shakespeareizing modern language using copy-enriched sequence to sequence models](#). In *Proceedings of the Workshop on Stylistic Variation*, pages 10–19, Copenhagen, Denmark. Association for Computational Linguistics.
- Hongyan Jing. 2000. [Sentence reduction for automatic text summarization](#). In *Sixth Applied Natural Language Processing Conference*, pages 310–315, Seattle, Washington, USA. Association for Computational Linguistics.
- K. Sparck Jones. 1973. [Index term weighting](#). *Information Storage and Retrieval*, 9(11):619–633.
- J Peter Kincaid, Robert P Fishburne Jr, Richard L Rogers, and Brad S Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and Flesch reading ease formula) for Navy enlisted personnel. Technical report, Naval Technical Training Command Millington TN Research Branch.
- Kevin Knight and Daniel Marcu. 2000. Statistics-based summarization - step one: Sentence compression. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, page 703–710. AAAI Press.

- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. [Delete, retrieve, generate: a simple approach to sentiment and style transfer](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1865–1874, New Orleans, Louisiana. Association for Computational Linguistics.
- Jared Lichtarge, Chris Alberti, Shankar Kumar, Noam Shazeer, Niki Parmar, and Simon Tong. 2019. [Corpora generation for grammatical error correction](#). 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 3291–3301, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pintu Lohar, Haithem Affi, and Andy Way. 2017. Maintaining sentiment polarity in translation of user-generated content. *Prague Bulletin of Mathematical Linguistics*, pages 73–84.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2018. [Sentence compression for arbitrary languages via multilingual pivoting](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2453–2464, Brussels, Belgium. Association for Computational Linguistics.
- George A. Miller. 1995. [WordNet: A lexical database for English](#). *Commun. ACM*, 38(11):39–41.
- Courtney Napoles, Keisuke Sakaguchi, and Joel Tetreault. 2017. [JFLEG: A fluency corpus and benchmark for grammatical error correction](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 229–234, Valencia, Spain. Association for Computational Linguistics.
- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. [The CoNLL-2014 shared task on grammatical error correction](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–14, Baltimore, Maryland. Association for Computational Linguistics.
- Mark K Osbeck. 2011. What is good legal writing and why does it matter. *Drexel L. Rev.*, 4:417.
- Julie A Oseid. 2009. The power of brevity: Adopt Abraham Lincoln’s habits. *J. Ass’n Legal Writing Directors*, 6:28.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018. Analyzing uncertainty in neural machine translation. In *International Conference on Machine Learning*, pages 3956–3965. PMLR.
- Paul Over, Hoa Dang, and Donna Harman. 2007. [DUC in context](#). *Information Processing & Management*, 43(6):1506–1520. Text Summarization.
- Bo Pang, Kevin Knight, and Daniel Marcu. 2003. [Syntax-based alignment of multiple translations: Extracting paraphrases and generating new sentences](#). In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 181–188.
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. [PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 425–430, Beijing, China. Association for Computational Linguistics.
- Sarah E Petersen and Mari Ostendorf. 2007. Text simplification for language learners: a corpus analysis. In *Workshop on Speech and Language Technology in Education*.
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. [Style transfer through back-translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876, Melbourne, Australia. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text Transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Luz Rello, Ricardo Baeza-Yates, Stefan Bott, and Horacio Saggion. 2013. [Simplify or help? text simplification strategies for people with dyslexia](#). In *Proceedings of the 10th International Cross-Disciplinary Conference on Web Accessibility, W4A ’13*, New York, NY, USA. Association for Computing Machinery.

- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Horacio Saggion. 2017. Automatic text simplification. *Synthesis Lectures on Human Language Technologies*, 10(1):1–137.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Controlling politeness in neural machine translation via side constraints](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40, San Diego, California. Association for Computational Linguistics.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- A. Siddharthan. 2002. [An architecture for a text simplification system](#). In *Language Engineering Conference, 2002. Proceedings*, pages 64–71.
- Felix Stahlberg, Iliia Kulikov, and Shankar Kumar. 2022. [Uncertainty determines the adequacy of the mode and the tractability of decoding in sequence-to-sequence models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8634–8645, Dublin, Ireland. Association for Computational Linguistics.
- William Strunk and E. B. White. 1918. *The Elements of style*. W.F. Humphrey, Ithaca, N.Y.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. [Lamda: Language models for dialog applications](#). *arXiv preprint arXiv:2201.08239*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Aldert Vrij. 2005. Criteria-based content analysis: A qualitative review of the first 37 studies. *Psychology, Public Policy, and Law*, 11(1):3.
- Willian Massami Watanabe, Arnaldo Candido Junior, Vinícius Rodriguez Uzêda, Renata Pontin de Matos Fortes, Thiago Alexandre Salgueiro Pardo, and Sandra Maria Aluísio. 2009. [Facilita: Reading assistance for low-literacy readers](#). In *Proceedings of the 27th ACM International Conference on Design of Communication, SIGDOC '09*, page 29–36, New York, NY, USA. Association for Computing Machinery.
- 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.
- Menglin Xia, Ekaterina Kochmar, and Ted Briscoe. 2016. [Text readability assessment for second language learners](#). In *Proceedings of the 11th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 12–22, San Diego, CA. Association for Computational Linguistics.
- Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. [Learning semantic textual similarity from conversations](#). In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 164–174, Melbourne, Australia. Association for Computational Linguistics.
- Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. 2020. [Large batch optimization for deep learning: Training BERT in 76 minutes](#). In *International Conference on Learning Representations*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. [PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.
- Xingxing Zhang and Mirella Lapata. 2017. [Sentence simplification with deep reinforcement learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 584–594, Copenhagen, Denmark. Association for Computational Linguistics.
- Renjie Zheng, Mingbo Ma, and Liang Huang. 2018. [Multi-reference training with pseudo-references for neural translation and text generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3188–3197, Brussels, Belgium. Association for Computational Linguistics.
- William Zinsser. 2016. *On writing well: The classic guide to writing nonfiction*. Harper Perennial, New York, N.Y.

Parameter	Value
Attention dropout rate	0.1
Attention layer size	1,024
Batch size	256
Beam size	10
Dropout rate	0.1
Embedding size	1,536
Learning rate	0.4
MLP dimension	4,096
Number of attention heads	4
Number of layers	6
Number of fine-tuning iterations	100-2,000 (early stopping)
Number of pre-training iterations	100,000
TPU topology	4x4

Table 9: Transformer hyper-parameters.

A Transformer hyper-parameters

Our round-trip translation based models (Sec. 3.2) are trained on TPUs with the LAMB optimizer (You et al., 2020) in JAX (Bradbury et al., 2021). We used the Transformer (Vaswani et al., 2017) implementation from the MT example in Flax¹² with the 32K SentencePiece vocabulary (Kudo and Richardson, 2018) from T5 (Raffel et al., 2020). Model hyper-parameters are listed in Table 9.

B Annotator instructions

The *Concise-Lite* annotators received the following instructions:

Rewrite the sentence to make it more concise, without changing the sentence structure. By sentence structure, we mean the general order of words in the sentence should not change, some sub-phrases could be rewritten/replaced/deleted (3-5 words). These should be relatively minor rewrites, such that you can replace a phrase with a shorter alternative without reorganizing the entire sentence. The sentences should be annotated in isolation without any assumptions on preceding or succeeding sentences.

The *Concise-Full* instructions are:

Rewrite the sentence to achieve maximum conciseness. These can be major rewrites that alter the sentence structure to make it as concise as possible. The annotator needs to make sure that the sentence stays the same semantically (meaning, intent & sentiment) and there is no loss of any critical information. The sentences should be annotated in isolation without any assumptions on preceding or succeeding sentences.

C Example outputs

Table 10 shows some example outputs of our systems and the baselines. The summarization (Long T5) system frequently changes the meaning of the source sentence. The simplification (Simplify T5) system performs slightly better but still changes the meaning in some instances (example c). The T5 system is mostly faithful to the meaning of the source sentence. We observe occasional slight meaning shifts with the Transformer and ParaNMT systems (see e.g. examples b) and g)). The Giant-LM often changes or expands the information in the source sentence (e.g. examples b) and d), f)) or adds certain artefacts (e.g. “*Here is a revision: ‘...’*” in example a)) that stem from its main use case as a user-facing dialog agent. Being a paraphrasing system, ParaNMT often falls short of actually improving the conciseness (examples c) and f)), and often uses unnecessary synonyms. Synonym replacements can also be found sometimes in Transformer and T5 outputs (examples a) and c)), but not in Giant-LM and human-annotated sentences. The pre-trained models Giant-LM and T5 are sometimes able to compress sentences by relying

¹²<https://github.com/google/flax/tree/master/examples/wmt/>

Example a)	
Source	A few other men fell and sneaked back, but I don't know their names.
Long T5	A few other men fell and sneaked back.
Simplify T5	A few other men fell and sneaked back, but I don't know their names.
Giant-LM	Here is a revision: Other men fell and snuck back, but their names I do not know.
ParaNMT	several other men fell and crept back, but I don't know their names.
Transformer	Some other men fell and sneaked, but I do not know their names.
T5	A few other men fell and escaped, but I do not know their names.
Human (C.-Lite)	Some other men fell and sneaked back, but I don't know their names.
Human (C.-Full)	Others fell and snuck back, but I don't know their names.
Example b)	
Source	"Helplessly Hoping" is a 1969 song by the American folk rock group Crosby, Stills, and Nash.
Long T5	Helplessly Hoping is a 1969 song by the American folk rock group Crosby, Stills and Nash.
Simplify T5	Helplessly Hoping is a song by the American folk rock group Crosby, Stills & Nash .
Giant-LM	"Helplessly Hoping" is a 1969 song by the American folk rock group Crosby, Stills, and Nash. The song was written for Stephen Stills' wife.
ParaNMT	"helpless Hoping" is a 1969 song from the American rock band Crosby, Stills and Nash.
Transformer	"Helpless Hoping" is a 1969 song by American folk rock group Crosby, Stills and Nash.
T5	"Helplessly Hoping" is a 1969 song by the American folk rock group Crosby, Stills and Nash.
Human (C.-Lite)	"Helplessly Hoping" is a 1969 song by American folk rock group Crosby, Stills, and Nash.
Human (C.-Full)	Crosby, Stills, and Nash's "Helplessly Hoping" was released in 1969.
Example c)	
Source	The NLA Tower, Britain's 88th tallest tower, is an example of original 1970s architecture in the town.
Long T5	The NLA Tower, Britain's 88th tallest tower, is an example of original 1970s architecture in the town.
Simplify T5	The NLA Tower is an example of original 1970s architecture in the town .
Giant-LM	NLA Tower, Britain's 88th tallest tower, is a 1970s example of architecture in the town.
ParaNMT	the NLA Tower, the 88th highest British tower, is an example of the architecture of the 1970s in the city.
Transformer	The NLA Tower, Britain's 88th tallest tower, is an example of the city's original 1970s architecture.
T5	The NLA Tower, Britain's 88th tallest building, is an example of the town's original 1970s architecture.
Human (C.-Lite)	The NLA Tower, Britain's 88th tallest tower, is an original 1970s architecture in town.
Human (C.-Full)	Britain's 88th tallest tower, the NLA Tower, is an original 1970s architecture.
Example d)	
Source	There are some claims that even the kaiser had been amused by the incident.
Long T5	Police say the kaiser used a special case kaiser to cover the snafu on a McDonald's breakfast.
Simplify T5	There are some claims that even the kaiser had been amused by the incident.
Giant-LM	the kaiser was reportedly amused by the incident – This is a pretty vague phrase, in fact. Was the kaiser laughing? Grinning?
ParaNMT	there are some claims that even the kaiser was amused by the incident.
Transformer	There are some claims that even Kaiser was amused by the incident.
T5	Some claim that even the Kaiser was amused by the incident.
Human (C.-Lite)	There are claims that even the Kaiser was amused by the incident.
Human (C.-Full)	Some claim that the incident amused even the kaiser.
Example e)	
Source	Approximately 50% of viewers spend between one and three hours watching TV per day.
Long T5	TV Viewership Up to 50%.
Simplify T5	Approximately 50% of viewers spend between one and three hours watching TV per day.
Giant-LM	About half of viewers watch TV for 1 to 3 hours each day.
ParaNMT	approximately 50% of viewers spend one to three hours watching TV every day.
Transformer	About 50% of viewers spend between one and three hours watching TV a day.
T5	About 50% of viewers watch TV for one to three hours a day.
Human (C.-Lite)	Approximately 50% of viewers spend between one and three hours watching TV daily.
Human (C.-Full)	50% of viewers spend 1-3 hours watching TV daily.
Example f)	
Source	He studied at Tübingen, Freiburg, Berlin, and Heidelberg, where he earned his medical degree in 1910.
Long T5	German Doctor Emil Henckel von Donnersmarck Dies at 87
Simplify T5	He studied at Tübingen, Freiburg, Berlin, and Heidelberg, where he earned his medical degree in 1910.
Giant-LM	He studied medicine in Tübingen, Freiburg, Berlin and Heidelberg.
ParaNMT	he studied at Tübingen, Freiburg, Berlin, and Heidelberg, where he earned a medical degree in 1910.
Transformer	He studied in Tübingen, Freiburg, Berlin and Heidelberg, where he graduated in medicine in 1910.
T5	He studied in Tübingen, Freiburg, Berlin and Heidelberg, where he received his medical degree in 1910.
Human (C.-Lite)	He studied at Tübingen, Freiburg, Berlin, and Heidelberg, where he earned his medical degree in 1910.
Human (C.-Full)	He studied at Tübingen, Freiburg, Berlin, and Heidelberg, earning his medical degree in 1910.
Example g)	
Source	Almost without exception, the Keetoowahs sided with the Northern States during the Civil War.
Long T5	Keetoowahs Speak Out on Civil War.
Simplify T5	Almost without exception, the Keetoowahs sided with the Northern States during the Civil War.
Giant-LM	Almost without exception, the Keetoowahs sided with the Union.
ParaNMT	almost without exception, Keetoowah sailed with the Northern States during the Civil War.
Transformer	Almost without exception, the Keetoowahs joined the northern states during the civil war.
T5	Almost without exception, the Keetoowahs sided with the North during the Civil War.
Human (C.-Lite)	The Keetoowahs sided with the Northern States during the Civil War.
Human (C.-Full)	During the Civil War, the Keetoowahs sided with the North.

Table 10: Example sentences from our conciseness systems and other baselines (summarization: Long T5, simplification: Simplify T5, ParaNMT). We use the “RT→MT” setups for the Transformer and T5 systems. We show one *Concise-Lite* and one *Concise-Full* human reference.

on background knowledge, e.g. by replacing “the Northern States” with “the Union” or “the North” in example g).