

Modeling Context in Answer Sentence Selection Systems on a Latency Budget

Rujun Han*

University of Southern California
Los Angeles, CA, USA
rujunhan@usc.edu

Luca Soldaini

Amazon Alexa
Manhattan Beach, CA, USA
lssoldai@amazon.com

Alessandro Moschitti

Amazon Alexa
Manhattan Beach, CA, USA
amosch@amazon.com

Abstract

Answer Sentence Selection (AS2) is an efficient approach for the design of open-domain Question Answering (QA) systems. In order to achieve low latency, traditional AS2 models score question-answer pairs individually, ignoring any information from the document each potential answer was extracted from. In contrast, more computationally expensive models designed for machine reading comprehension tasks typically receive one or more passages as input, which often results in better accuracy. In this work, we present an approach to efficiently incorporate contextual information in AS2 models. For each answer candidate, we first use unsupervised similarity techniques to extract relevant sentences from its source document, which we then feed into an efficient transformer architecture fine-tuned for AS2. Our best approach, which leverages a multi-way attention architecture to efficiently encode context, improves 6% to 11% over non-contextual state of the art in AS2 with minimal impact on system latency. All experiments in this work were conducted in English.

1 Introduction

AS2 models for open-domain QA typically consider sentences from webpages as independent candidate answers for a given question. For any webpage containing potential answer candidates for a question, AS2 models first extract individual sentences, then independently estimate their likelihood of being correct answers; this approach enable highly efficient processing of entire documents. However, under this framework, context information from the entire webpage (global context), which could be crucial for selecting correct answers, is ignored. Conversely, current systems in Machine Reading (MR) (Huang et al., 2019;

* Work was conducted while the author was an intern at Amazon Alexa.

The math of pi explained, as simply as possible How many digits of pi we really need?

Pi, you may also remember from grade school, is not an ordinary number. It's irrational, meaning it has an endless number of decimals that never repeat. **Though even cutting off pi at 15 digits allows for extremely precise measurements.**

If you were to draw a circle with a diameter of 25 billion miles, using 15 digits of pi, you'd only arrive at a measurement of the circumference that's off by 1.5 inches, NASA's Marc Rayman explained in a post on NASA's JPL website. And that's good enough.

Of course, that hasn't stopped people from looking for more and more digits of pi. **Currently, there are more than 22.4 trillion known digits, which show no hint of ending or repeating.** Further reading: pi and pie.

Figure 1: Candidate sentence (in orange) for question “how many digits are in pi?”. Local context is shown in green, while global context is shown in red.

Kwiatkowski et al., 2019a; Lee et al., 2019; Joshi et al., 2020a) uses a much larger context from the retrieved documents. MR models receive a question and one or more passages retrieved through a search engine as input; they then select one or more spans from the input passages to return as answer.

While being potentially more accurate, MR models typically have higher computational requirements (and thus higher latency) than AS2 models. That is because MR models need to process passages in their entirety before an answer can be extracted; conversely, AS2 systems break down paragraphs in candidate sentences, and evaluate them all at once in parallel. Therefore, in many practical applications, MR models are only used to examine 10 to 50 candidate passages; in contrast, AS2 approaches can potentially process hundreds of documents, e.g., (Matsubara et al., 2020; Soldaini and Moschitti, 2020).

In this work, we study techniques that can combine the efficacy of MR models with the efficiency of AS2 approaches, while keeping a single sentence as target answer, as in related AS2 works¹. In particular, we focus our efforts on improving accuracy of AS2 systems without affecting their latency.

¹MR systems have a different aim than AS2.

Early neural models for retrieval-based QA focused on incorporating neighboring sentences (local context) to improve performance. For example, Tan et al. (2017) proposed a neural architecture based on gated recurrent units to encode question, answer, and local context; their approach, while effective at the time, shows a significant gap to the current state-of-the-art models (Garg et al., 2020). Min et al. (2018) studied neural efficient models for MR by optimizing answer candidate extraction. More recently, researchers have focused in including source document information in transformer models. For example, Joshi et al. (2020b) proposed a contextualized model for MR that augments named entities in candidate passages with snippets extracted from Wikipedia pages. Their approach, while interesting, is limited to entities-based context, and specific to Wikipedia and MR domain. For AS2, Lauriola and Moschitti (2021) proposed a model that uses local context as defined by the preceding and following sentences of the target answer. They also introduced a simple bag-of-words representation of documents as global context, which did not show significant improvement over non-contextual AS2 models.

Unlike previous efforts, our approaches consider both *local context* (that is, the sentences immediately preceding or succeeding a candidate answer), as well as *global context* (phrases in documents that represent the overall topics of a page), as they can both uniquely contribute to the process of selecting the right answer. As shown in the example in Figure 1, local context can help disambiguate cases where crucial entities are not present in the candidate answer (there’s no mention of “*pi*” in “[c]urrently, there are more than 22.4 trillion known digits”); conversely, global context can help reaffirm the relevance of a candidate answer in cases where noisy information is extracted as local context (in the example, “[f]urther reading: *pi* and *pie*” does not contain any relevant information).

The contributions of this work are: (i) first, we introduce two effective techniques to extract relevant local and global contexts for a given question and candidate answer; (ii) then, we propose three different methods for combining contextual information for AS2 tasks; (iii) finally, we evaluate our approaches on two AS2 datasets: ASNQ (Garg et al., 2020) and a benchmark dataset we built to evaluate real-world QA systems. Results show that our most efficient system, which lever-

Question: “where did the potter’s wheel first develop”

Correct Answer: “Tournettes, in use around 4500 BC in the Near East, were turned slowly by hand or by foot while coiling a pot”

Sentence selected by N-grams: “In the Iron Age, the potter’s wheel in common use had a turning platform about one metre (3 feet) above the floor, connected by a long axle to a heavy flywheel at ground level. Use of the potter’s wheel became widespread throughout the Old World but was unknown in the Pre-Columbian New World, where pottery was handmade by methods that included coiling and beating.”

Question: “where do pineapples come from in the world”

Correct answer: “In 2016, Costa Rica, Brazil, and the Philippines accounted for nearly one-third of the world’s production of pineapple.”

Sentence selected by Cosine Similarity: “The plant is indigenous to South America and is said to originate from the area between southern Brazil and Paraguay”

Table 1: Examples of **global context** selected via N-gram similarity (top) and cosine similarity (bottom). Overall, the N-gram approach tends to select longer context sentences than Cosine’s, which in turn leads to fewer context sentences being included in the global context (as we limit it to 128 tokens). Empirically, we also noticed that N-gram selected context sentences also contain more noise.

ages a multi-way attention architecture, can improve over the previous non-contextual state of the art model for AS2 by up to 11%; furthermore, these results are achieved while maintaining similar efficiency to the best-performing, non-contextual AS2 systems, making our approach a viable strategy for latency-sensitive applications. Code and models are made available at <https://github.com/alexa/wqa-contextual-qa>.

2 Methodology

Our approach to ranking candidate answer consists of two components: the first (Section 2.1) is responsible for extracting context for each candidate answer, while the second (Section 2.2) encodes information from local and global contexts to score each question / candidate answer pair.

2.1 Context Construction

As previously mentioned, our proposed method for contextualizing answers relies on enriching them with information encoded in sentences adjacent to them, as well as from sentences throughout the document each potential answer comes from; we will define these extraction processes in this section.

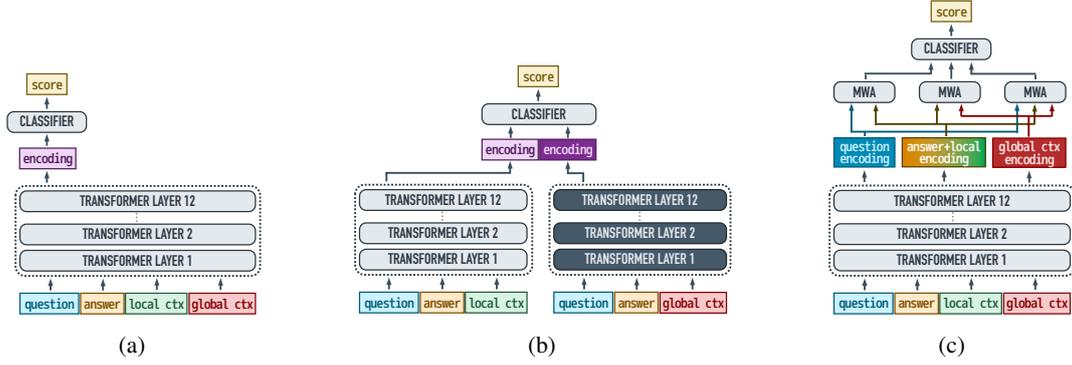


Figure 2: From left to right, the three approaches we evaluated in this work: context concatenation (Figure 2a), context ensemble (Figure 2b), and multi-way attention (Figure 2c).

In the rest of this work, we will use Q to refer to a question and $\mathcal{D} = \{D_1, \dots, D_i, \dots, D_N\}$ to indicate a collection of documents containing potential answers for Q . Each document D_i is comprised of an ordered sequence of sentences $D_i = \langle C_{i,1}, \dots, C_{i,j}, \dots, C_{i,M} \rangle$; each sentence $C_{i,j}$ could be used either as candidate answer, or as context for another candidate.

2.1.1 Local Context

Similarly to previous work (Tan et al., 2017; Lauriola and Moschitti, 2021), we define local context $\text{Loc}_k(C_{i,j})$ for candidate $C_{i,j}$ as the sentences immediately preceding and succeeding each answer candidate within a window of $2k + 1$ sentences, i.e., $\text{Loc}_k(C_{i,j}) = \langle C_{i,j-k}, \dots, C_{i,j-1}, C_{i,j+1}, \dots, C_{i,j+k} \rangle$. In our experiments, we tried constructing a local context of up to six sentences; however, we observed diminishing return when using more than the previous and next sentences (i.e., $k = 1$) at the expense of more computational complexity. Therefore, the results presented in this work use two adjacent sentences as local context.

2.1.2 Global Context

Unlike local context, there are many potential approaches to extracting information that can be used to understand relevancy of a candidate answer to a question. We proposed and evaluated two different techniques for extracting global context $\text{Glo}_h(C_{i,j})$ (examples for both are shown in Table 1).

N-gram Overlap Similarly, to Joshi et al. (2020b), we experimented with selecting sentences as global context based on their n-gram overlap with question and candidates.

In detail, we first extract the set of all unigrams, bigrams, and trigrams from question Q and can-

didate $C_{i,j}$, which we denote as $\text{Ng}_{1,2,3}(Q, C_{i,j})$; then, we repeat the same procedure for all $\{C_{i,p} \in \mathcal{D} \text{ where } p \neq j\}$ to obtain $\text{Ng}_{1,2,3}(C_{i,p})$. Finally, we score each sentence as follows:

$$\text{Score}_{\text{ngrams}}(C_{i,p}|Q, C_{i,j}) = \frac{|\text{Ng}_{1,2,3}(C_{i,p}) \cap \text{Ng}_{1,2,3}(Q, C_{i,j})|}{|\text{Ng}_{1,2,3}(Q, C_{i,j})|} \quad (1)$$

and pick the top h sentences as global context.

Semantic Similarity N-grams overlap can only extract spans of text that are lexically similar to either the query or candidate. To better capture context that is topically relevant to an answer, we also propose to use cosine distance between sentence embeddings to approximate semantic similarity.

Given a sentence encoder model² \mathcal{M} , we first obtain a representations for the question-answer pair $\mathcal{M}(Q \oplus C_{i,j})$ and context sentences $\{\mathcal{M}(C_{i,p}) \text{ for all } p = \{1, \dots, M\}, p \neq j\}$; then we pick the top h sentences maximizing the following cosine similarity score as global context:

$$\text{Score}_{\text{sim}}(C_{i,p}|Q, C_{i,j}) = \frac{\mathcal{M}(Q \oplus C_{i,j}) \cdot \mathcal{M}(C_{i,p})}{\|\mathcal{M}(Q \oplus C_{i,j})\| \|\mathcal{M}(C_{i,p})\|}, \quad (2)$$

where \oplus indicates string concatenation.

2.2 Contextual AS2 Models

Once local context $\text{Loc}_k(C_{i,j})$ and global context $\text{Glo}_h(C_{i,j})$ are extracted for candidate $C_{i,j}$, we encode them in conjunction with candidate answer and question to estimate the likelihood of $C_{i,j}$ being a correct answer for Q . Our approaches (summarized in Figure 2) consume up to $h = 5$ sen-

²In our experiments, we use non-finetuned RoBERTa_{BASE} model (Liu et al., 2019).

Model	P@1	MAP	MRR	Latency
No Context (Garg et al., 2020)	0.596	0.685	0.706	5.49
Local Context (Lauriola and Moschitti, 2021)	0.653	0.732	0.752	5.62
Global Context	0.622	0.698	0.722	5.67
Context Concat	0.631	0.726	0.743	5.76
Context Ensemble	0.668	0.743	0.765	7.41
MWA	0.661	0.742	0.758	5.82

Table 2: Results on the ASNQ Test Set. The latency is reported in micro-seconds (μs) per sample ($\pm 0.1 \mu s$ with 95% CI).

tences as global context so not to exceed 128 tokens.³ Similarly, to other efforts in this area (e.g., (Garg et al., 2020)), we leverage state-of-the-art transformer models to estimate said probability. Specifically, we studied three approaches to encode question and answer context. Although the methods we proposed can be easily combined with any transformer architecture, all models described here are initialized from a RoBERTa_{BASE} checkpoint.

Context Concatenation A simple baseline (van Aken et al., 2019; Joshi et al., 2020b) for encoding multiple contexts is to concatenate each question, candidate answer, and local/global context text and feed them through transformer model (Fig. 2a); the resulting encoding is then projected to a probability distribution using a dense feed-forward layer. This baseline relies on the transformer self-attention mechanism to implicitly model relations between local and global context.

Context Ensemble As mentioned in Section 1, local and global contexts might capture different aspects of the source document of a candidate answer. To empirically verify this hypothesis, we evaluated an ensemble model that encodes local and global contexts separately using two independent transformer models (Figure 2b). The two models are independently trained for AS2; then, their encodings are concatenated and passed to a feed-forward layer to estimate relevance of candidate $C_{i,j}$ for question Q . The top 3 layers⁴ model is once again fine tuned on the training set.

³Using up 5 sentences resulted in 3.02 (n-gram context) and 2.85 (cosine context) sentences being selected on average for the ASNQ dataset.

⁴We tested with approaches for gradual unfreezing of the top k layers; $k = 3$ yielded the best validation scores.

Model	P@1	MAP	MRR
No Context		<i>baseline</i>	
Global Context	+3.95%	+2.52%	+2.29%
Local Context	+3.59%	+2.89%	+2.21%
Context Concat	+1.52%	+2.33%	+1.34%
Context Ensemble	+5.92%	+4.10%	+3.39%
MWA	+5.56%	+3.92%	+3.08%

Table 3: Results on the WQA Test Set.

Multi-way Attention (MWA) While leveraging independent encoders for local and global contexts can lead to an improvement in performance compared to using a single encoder, it also doubles computational requirements. Therefore, we also explored techniques that incorporate inductive biases into transformer models and achieve efficiency comparable to the context concatenation approach. One such approach is, as shown in Figure 2c, to combine a transformer model with a multi-way attention mechanism (Tan et al., 2018), which has been shown to be effective for commonsense reasoning tasks (Huang et al., 2019). This approach still uses a single transformer model to produce an encoding for a sequence of question, candidate answer, local context, and global context; however, similarly to the ensemble model, the additional attention mechanism forces the encoder to selectively look at local and global contexts separately.

3 Experiments

3.1 Setup

In order to validate the effectiveness of the proposed context modeling techniques, we evaluated our results on two datasets: ASNQ and WQA.

ASNQ The Answer Sentence Natural Questions dataset (Garg et al., 2020) is a large collection of 59,914 questions and 24,732,396 candidate answers. It was obtained by extracting sentence candidates from the Google Natural Question (NQ) benchmark (Kwiatkowski et al., 2019b). We use the train, development, and test splits proposed by Soldaini and Moschitti (2020).

WQA The Web-based Question Answering is an in-house dataset built by Alexa AI as part of the effort of understanding and benchmarking QA systems. The creation process includes the following steps: (i) Given a set of questions, a search engine is used to retrieve up to 100 web pages from an index containing hundreds of million pages. (ii) From the set of retrieved documents, all candidate

Technique	P@1	MAP	MRR
N-gram Overlap	0.638	0.725	0.746
Cosine Similarity	0.661	0.742	0.758

Table 4: Comparison of global context extraction techniques on the ASNQ test set when used with MWA.

sentences are extracted and ranked using AS2 models from Garg et al. (2020); and (iii) at least 25 candidates for each question are annotated by humans. Overall, the version of WQA we used contains 6,962 questions and 283,855 candidate answers. We reserved 3,000 questions for evaluation, 808 for development, and used the rest for training⁵.

Models were trained on a single machine with 8 NVIDIA Tesla V100 GPUs with 32 GB of memory each. We used model implementations from the Transformers library when available (Wolf et al., 2020). All our experiments were computed using mixed precision through NVIDIA apex⁶. Latency was measured on single GPU with a fixed batch size of 128. Tokenization and time to transfer tensors to the GPU was not included in the latency values.

3.2 Results and Discussion

Results on the ASNQ and WQA are summarized in tables 2 and 3, respectively. Overall, we observe that the context ensemble model achieves the best performance; however, as observed in Section 2.2, this model is twice as large as a RoBERTa_{BASE} model, thus it is a rather expensive solution.

Among our baselines, we note that local context outperforms the model leveraging the global context. This observation suggests that local information carries more importance in understanding the semantic relationship between question and candidate answers. Surprisingly, we observe that simply concatenating local and global contexts achieves worse performance of local context alone, and it even underperforms the global context method on WQA. This suggests that, without any additional structure, the self-attention mechanism of the transformer cannot effectively distinguish and leverage information from the local and global contexts.

We note that MWA achieves near identical performance to the ensemble model on both datasets,

⁵The public version of WQA will be released in the short-term future. Please search for a publication with title *WQA: A Dataset for Web-based Question Answering Tasks* on arXiv.org.

⁶<https://github.com/NVIDIA/apex>

suggesting that a controlled attention mechanism can overcome limitations in the representation for vanilla transformers, while reducing latency by 21.5% and memory usage by 89%. MWA also matches the latency of the context concatenation model, while improving it by 4.8% and 3.9% in P@1 on ASNQ and WQA, respectively.

Finally, we study the effect of our proposed global extraction techniques in Table 4. We observe that, among the two proposed algorithms, the cosine similarity approach significantly outperforms the N-gram based method. This confirms that pre-trained language models can better select context semantically related to question and candidates.

We note n-gram overlap is less computationally taxing, as it can be efficiently implemented as a set of sparse operations over bag of word representations of the question and answer candidates. On the other hand, for cosine similarity, it is necessary to compute $\text{Score}_{sim}(C_{i,p}|Q, C_{i,j})$ for all context sentences using a transformer model. Recently introduced transformer architecture variants could be used to either speed up this similarity computation (Cao et al., 2020), or compute query and text representation independently (Khattab and Zaharia, 2020). We leave the evaluation of these techniques to future work.

4 Conclusions

For efficiency reasons, traditional AS2 models are designed to estimate answer relevancy by only comparing question and candidates. In this work, we described and evaluate several techniques to incorporate local and global contexts in the answer selection process. The results of our experiments show that our proposed methods significantly outperform non-contextual approaches; further, we empirically demonstrate that local and global contexts can be effectively combined to further improve ranking performance.

References

- Betty van Aken, Benjamin Winter, Alexander Löser, and Felix A. Gers. 2019. *How does bert answer questions? a layer-wise analysis of transformer representations*. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 1823–1832, New York, NY, USA. Association for Computing Machinery.
- Qingqing Cao, Harsh Trivedi, Aruna Balasubramanian, and Niranjan Balasubramanian. 2020. *De-*

- Former: Decomposing pre-trained transformers for faster question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4497, Online. Association for Computational Linguistics.
- Siddhant Garg, Thuy Vu, and Alessandro Moschitti. 2020. **Tanda: Transfer and adapt pre-trained transformer models for answer sentence selection**. In *AAAI*, pages 7780–7788.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. **Cosmos QA: Machine reading comprehension with contextual commonsense reasoning**. 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 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020a. **SpanBERT: Improving pre-training by representing and predicting spans**. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Mandar Joshi, Kenton Lee, Yi Luan, and Kristina Toutanova. 2020b. Contextualized representations using textual encyclopedic knowledge. *arXiv preprint arXiv:2004.12006*.
- Omar Khattab and Matei Zaharia. 2020. **Colbert: Efficient and effective passage search via contextualized late interaction over bert**. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 39–48.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019a. **Natural questions: a benchmark for question answering research**. *Transactions of the Association of Computational Linguistics*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019b. **Natural questions: a benchmark for question answering research**. *Transactions of the Association of Computational Linguistics*.
- Ivano Lauriola and Alessandro Moschitti. 2021. Answer sentence selection using local and global context in transformer models. *ECIR*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. **Latent retrieval for weakly supervised open domain question answering**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**. *arXiv preprint arXiv:1907.11692*.
- Yoshitomo Matsubara, Thuy Vu, and Alessandro Moschitti. 2020. **Reranking for Efficient Transformer-Based Answer Selection**. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '20, page 1577–1580, New York, NY, USA. Association for Computing Machinery.
- Sewon Min, Victor Zhong, Richard Socher, and Caiming Xiong. 2018. **Efficient and robust question answering from minimal context over documents**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1725–1735, Melbourne, Australia. Association for Computational Linguistics.
- Luca Soldaini and Alessandro Moschitti. 2020. **The cascade transformer: an application for efficient answer sentence selection**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5697–5708, Online. Association for Computational Linguistics.
- Chuanqi Tan, Furu Wei, Wenhui Wang, Weifeng Lv, and Ming Zhou. 2018. **Multiway attention networks for modeling sentence pairs**. In *IJCAI*, pages 4411–4417.
- Chuanqi Tan, Furu Wei, Qingyu Zhou, Nan Yang, Bowen Du, Weifeng Lv, and Ming Zhou. 2017. **Context-aware answer sentence selection with hierarchical gated recurrent neural networks**. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(3):540–549.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.