

Improving Sequence-to-Sequence Pre-training via Sequence Span Rewriting

Wangchunshu Zhou^{1*} Tao Ge² Canwen Xu³ Ke Xu⁴ Furu Wei²

¹Stanford University ²Microsoft Research Asia

³University of California, San Diego ⁴Beihang University

wcszhou@stanford.edu.cn, cxu@ucsd.edu kexu@nlsde.buaa.edu.cn

{tage, fuwei}@microsoft.com

Abstract

In this paper, we propose Sequence Span Rewriting (SSR), a self-supervised task for sequence-to-sequence (Seq2Seq) pre-training. SSR learns to refine the machine-generated imperfect text spans into ground truth text. SSR provides more fine-grained and informative supervision in addition to the original text-infilling objective. Compared to the prevalent text infilling objectives for Seq2Seq pre-training, SSR is naturally more consistent with many downstream generation tasks that require sentence rewriting (e.g., text summarization, question generation, grammatical error correction, and paraphrase generation). We conduct extensive experiments by using SSR to improve the typical Seq2Seq pre-trained model T5 in a continual pre-training setting and show substantial improvements over T5 on various natural language generation tasks.¹

1 Introduction

Text infilling (e.g., masked language modeling) has become a prevalent learning objective for pre-trained language models (PTLMs) (Peters et al., 2018; Radford et al., 2018; Devlin et al., 2019; Yang et al., 2019; Liu et al., 2019; Lan et al., 2020; Lewis et al., 2020b; Raffel et al., 2019). It provides self-supervision by masking out tokens or spans in text, and trains a model to infill the masked content based on the contexts, accordingly guiding the model for representation learning, as Figure 1(a) shows.

In this paper, we propose to extend the conventional text infilling to a novel sequence-to-sequence (Seq2Seq) pre-training objective, namely Sequence Span Rewriting (SSR). We train a model to rewrite machine-generated imperfect text spans into the

ground truth text, as illustrated in Figure 1(b). SSR has two advantages over text infilling: **(1) SSR provides better supervision signals**, as SSR trains the model with diverse and fine-grained rewriting patterns beyond filling the blanks; **(2) SSR bridges the gap between pre-training and fine-tuning**, because many downstream Seq2Seq tasks like summarization and paraphrase generation are naturally sequence span rewriting tasks where a source sentence is mapped to the target sentence following specific rewriting patterns.

The key element in implementing SSR is how to generate imperfect text spans that are both diverse and informative. Inspired by ELECTRA (Clark et al., 2020), we use a powerful pre-trained text infilling model – T5-large (Raffel et al., 2019) – as the imperfect span generator. Compared with random or rule-based noising approaches, the T5-based imperfect span generator can derive various informative text spans that benefit the model to learn meaningful and diverse rewriting patterns including paraphrasing and enhancing the fluency and contextual consistency through correcting grammatical, commonsense and factual errors, to improve a text sequence. These rewriting patterns resemble the goal of various NLG tasks and thus strengthen the ability of pre-trained model for downstream applications.

In our experiments, we apply SSR to the typical Seq2Seq pre-trained model – T5 (Raffel et al., 2019) in a continual learning fashion. We show SSR outperforms both the original pre-trained T5 models and their continual training counterparts with the conventional text infilling objective on various Seq2Seq tasks, including text summarization, question generation, and grammatical error correction, with a small number of optimization steps with moderate amount of machine-generated data, which confirms the potential of SSR to serve as a plug-and-play method to improve various existing pre-trained Seq2Seq models. Notably, we find SSR

^{*}This work was done during the first author’s internship at Microsoft Research Asia.

¹Code for pre-training SSR is available at https://github.com/MichaelZhouwang/Sequence_Span_Rewriting.

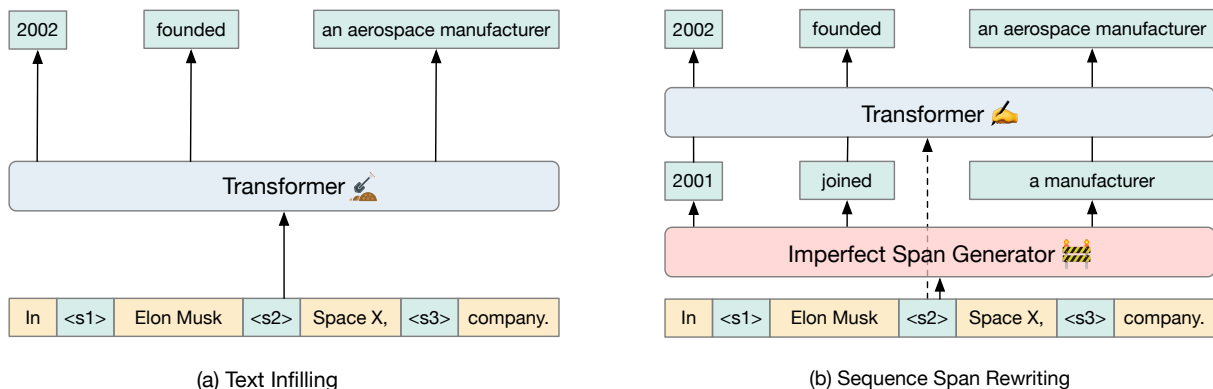


Figure 1: The comparison of (a) Text Infilling and (b) Sequence Span Rewriting. Instead of learning to directly fill the blanks, Sequence Span Rewriting first exploits an imperfect span generator to generate imperfect spans within the text and then feeds the filled text to the model to learn how to rewrite it into the ground truth.

especially useful for pre-training smaller Seq2Seq models, with the help of a powerful imperfect span generator. This observation sheds light on a new approach for knowledge transfer from large models to smaller ones.

2 Related Work

Pre-training in NLP BERT (Devlin et al., 2019) introduced the masked language modeling objective by masking out certain tokens in a text and predicting them based on their left and right side contexts. Recent work has shown that BERT’s performance can be further improved by training for longer (Liu et al., 2019), by tying parameters across layers (Lan et al., 2020), and by replacing a consecutive span of tokens with the mask token for MLM training (Joshi et al., 2020). Our approach is also related to ELECTRA (Clark et al., 2020), which uses a pre-trained masked language model to generate fake tokens and train a discriminator to detect them. The key difference is that our approach focuses on span-level texts and trains the model to correct the mistakes instead of simply detecting them, which includes more diverse and informative signals and enables the model to perform text generation tasks in a Seq2Seq fashion.

To enable mask language models for natural language generation tasks, Song et al. (2019) used a decoder to generate the masked tokens autoregressively. UniLM (Dong et al., 2019) multitasks MLM and language modeling objectives. More recently, BART (Lewis et al., 2020b) and T5 (Raffel et al., 2019) pre-train Seq2Seq models with the text span infilling objective, which removes text spans in the input texts and train the models to recover the original texts in an auto-regressive fashion.

More recently, CALM (Zhou et al., 2021a) introduces concept-to-sentence generation and concept order recovery as two self-supervised objectives that encourage Seq2Seq PTLMs to acquire generative commonsense reasoning ability. MARGE (Lewis et al., 2020a) pre-trains a Seq2Seq model with an unsupervised multi-lingual cross-document paraphrasing objective. Their approach is related to our text rewriting objective. However, MARGE requires multi-lingual paraphrase documents and needs to train a separate retrieval model while our method can simply used an off-the-shelf model pre-trained with text infilling to generate training data. Also, MARGE is pre-trained to generate a paraphrase-like document in another language, thus mainly helpful for translation tasks and multi-lingual tasks. In contrast, SSR focus on monolingual text rewriting and improve general text generation tasks.

SSR departs significantly from the aforementioned methods for Seq2Seq pre-training as it employ machine-generated noises instead of rule-based ones, thus introducing more diverse training signals. Also, SSR receives complete inputs without artificial masks during pre-training relying solely on monolingual corpus.

Model Acceleration for PTLMs Recently, many attempts have been made to speed up a large pre-trained language model (PTLM). To name a few, Shen et al. (2020) quantized BERT to 2-bit using Hessian information; Michel et al. (2019) pruned unnecessary attention heads in the transformer layers to reduce the parameters of a BERT model. DistilBERT (Sanh et al., 2019) and uses knowledge distillation (Hinton et al., 2015;

Romero et al., 2015) to compress BERT. More recently, (Zhou et al., 2021b) proposed Meta Distillation to improve the performance of knowledge distillation for compression BERT. In addition, Xu et al. (2020) introduced progressive module replacing to train more compact BERT models by encouraging the student model to behave similarly with the teacher model. In addition, Zhou et al. (2020c); Schwartz et al. (2020) proposed to accelerate the inference stage of pre-trained models via input-adaptive inference. However, to the best of our knowledge, few studies have been done for accelerating large sequence-to-sequence PTLMs. Our approach can also be used for model compression by using a large pre-trained model as the imperfect span generator. In this way, SSR also exploits the knowledge of a larger model to improve the training of a compact model.

3 Methodology

The key idea of SSR is to train a Seq2Seq model to rewrite machine-generated text spans that may contain a variety of noise such as paraphrase, grammatical and factual errors, into ground truth that are correct and appropriate in the context. As illustrated by Figure 1(b), SSR involves three steps: (1) masking out parts of the text; (2) generating imperfect text to fill in the masked spans; (3) training the Seq2Seq model to rewrite the imperfect spans to the ground truth. We will introduce the technical details of SSR in Section 3.1 and an advanced training strategy for SSR in Section 3.2.

3.1 Sequence Span Rewriting

Text Span Masking To generate training data of sequence span rewriting in a self-supervised fashion, we first randomly sample a number of text spans and mask them. Specifically, the spans are masked with special mask tokens by order (e.g., $\langle s_1 \rangle$, $\langle s_2 \rangle$ and $\langle s_3 \rangle$) in Figure 1(b) as in T5, with span lengths drawn from a Poisson distribution ($\lambda = 3$). The number of spans is controlled so that approximately 30% of all tokens are masked. Specially, 0-length spans correspond to an insertion of a mask token.

For example, as shown in Figure 1, given a sentence “In 2002, Elon Musk founded SpaceX, an aerospace manufacturer company.”, we randomly sample three text spans (two of them are of length 1). The masked sentence becomes “In $\langle s_1 \rangle$, Elon Musk $\langle s_2 \rangle$ SpaceX, $\langle s_3 \rangle$ company.”

Imperfect Span Generation With masked spans, we can generate imperfect text to fill in the spans. Specifically, we feed the masked input into the imperfect span generator to generate predictions in an auto-regressive fashion. To improve the diversity of generation, we use nucleus sampling (Holtzman et al., 2020) that truncates the unreliable tail of the probability distribution and samples from the dynamic nucleus of tokens containing the vast majority of the probability mass. For instance, given the previous masked input sentence, a T5-large model generates “2001”, “joined”, and “a manufacturer” as imperfect spans.

Span Rewriting After we obtain imperfect spans within the text, we pre-train the Seq2Seq model to rewrite imperfect text spans into the ground truth. Specifically, we use special tokens $\langle s_i \rangle$ and $\langle /s_i \rangle$ to denote the starting and ending of i -th text span to be rewritten in the source sequence, which gives “In $\langle s_1 \rangle$ 2001 $\langle /s_1 \rangle$, Elon Musk $\langle s_2 \rangle$ joined $\langle /s_2 \rangle$ SpaceX, $\langle s_3 \rangle$ a manufacturer $\langle /s_3 \rangle$ company.” as the input for SSR pre-training. Similarly, we use $\langle s_i \rangle$ to separate different text spans in the target sequence, which gives “ $\langle s_1 \rangle$ 2002 $\langle s_2 \rangle$ founded $\langle s_3 \rangle$ an aerospace manufacturer” as the target sequence. We train the model to generate target text spans from left to right auto-regressively by maximum likelihood estimation.

We can see that the SSR objective involves using a pre-trained model to generate imperfect spans, which will lead to increased computational cost. In practice, we suggest starting SSR pre-training based on checkpoints of existing Seq2Seq pre-trained models. In this way, we only need to generate a few amount of imperfect spans and continually pre-train the models for a few steps. In this perspective, SSR can be viewed as a general approach that be used to improve various Seq2Seq pre-trained models before fine-tuning them on downstream text generation tasks.

For fine-tuning SSR, we simply denote the entire input sequence with the same span identifier (e.g., $\langle s_1 \rangle$) used during SSR pre-training. Therefore, the model would learn to rewrite the entire input sequence, alleviating the gap caused by the $\langle \text{mask} \rangle$ token during text infilling pre-training. For example, for grammatical error correction, the input is formatted as “ $\langle s_1 \rangle$ I go to school yesterday. $\langle /s_1 \rangle$ ” and the output is “ $\langle s_1 \rangle$ I went to school yesterday.”, which exactly corresponds to the pre-training format of SSR.

In addition, for some constrained text generation tasks (Lin et al., 2020) and controlled text generation (Hu et al., 2017) tasks, we can specify which part of input text to be rewritten with span identifiers. This enables more flexible text generation with Seq2Seq pre-trained models. Taking text attribute transfer as an example, an input example would look like “Great food $\langle s_1 \rangle$ but very rude $\langle /s_1 \rangle$ waiters.” and the corresponding target sequence is “ $\langle s_1 \rangle$ and very friendly”. The inductive bias of span rewriting learned by SSR pre-training naturally benefit these kind of NLG applications.

3.2 Curriculum SSR

As mentioned above, we apply SSR as a continual training objective for pre-trained Seq2Seq models that were originally trained with the text infilling objective. However, continually training a pre-trained Seq2Seq model with a different objective may result in drastic adaption of its parameters. To make this transition smoother and reduce the difficulty of optimization, we propose to schedule the SSR training examples with curriculum learning (Bengio et al., 2009) according to their difficulties. Specifically, we measure the difficulty of rewriting a certain imperfect text span with both the length of the imperfect span and the uncertainty (i.e., perplexity) of the imperfect span generator when generating this span.

Intuitively, a short imperfect span generally includes some simple word substitution (e.g., big \rightarrow large) or grammatical error (e.g., is \rightarrow was) while a longer imperfect span may require more complicated paraphrasing (e.g., what is happening \rightarrow what’s up). Also, an imperfect span with larger perplexity suggests the span may be of lower quality or more uncommon, thus more difficult to be rewritten into ground truth. Therefore, we consider longer imperfect spans and spans with a higher perplexity under the imperfect span generator to be more difficult. We split the SSR training examples into k ($k = 5$ in our experiments) groups according to the sum of per-token loss of the imperfect span generator when it generates an SSR training example. We then start pre-training the model with the easiest group of SSR training examples and then gradually switch to more difficult groups during pre-training. Intuitively, this will make the transition from the original text infilling objective to the sequence span rewriting objective more smooth.

4 Experiments

4.1 Experimental Settings

SSR is implemented as a text-to-text transformer model with a bidirectional encoder and a left-to-right auto-regressive decoder. For pre-training, we minimize the negative log-likelihood of the original ground truth text spans. We describe details of the architecture, pre-training, and fine-tuning of SSR in this section.

Architecture We use the same architecture as T5 (Raffel et al., 2019) which is roughly equivalent to the original Transformer proposed by Vaswani et al. (2017), with the exception of removing the Layer Norm bias, placing the layer normalization outside the residual path, and using a different relative position embedding scheme.

Following the design choice of T5 (Raffel et al., 2019), we train three sizes of SSR:

- SSR-small: 60M parameters, 6 Transformer layers, 8 attention heads, 512 hidden size
- SSR-base: 220M parameters, 12 Transformer layers, 12 attention heads, 768 hidden size
- SSR-large: 770M parameters, 24 Transformer layers, 16 attention heads, 1024 hidden size

Pre-training Details As we propose SSR to serve as a general plug-and-play approach to improve existing Seq2Seq pre-trained models without intensive computation like pre-training from scratch, we initialize each size of SSR model with the corresponding pre-trained T5 model of the same size, and continually pre-train the models with the SSR objective.

For imperfect span generation, we use the off-the-shelf T5-large model with nucleus sampling ($p = 0.9$) to sample generated text spans. For SSR learning, we sample 4GB of text from Wikipedia corpus, BookCorpus (Zhu et al., 2015), and RealNews (Zellers et al., 2019), which are commonly used for pre-training language models. Our implementation is based on Hugging Face Transformers (Wolf et al., 2020). We use text sequences with a maximum length of 256 tokens to sample masked text spans and generate imperfect text spans. We then continually pre-train different variants of SSR for 100k updates², with a maximum sequence length of 256, a batch size of 512, and a

²We empirically find 100k updates to be enough since the models’ performance on downstream tasks begin to saturate.

Model	Architecture	CNN/DM			XSum		
		RG-1	RG-2	RG-L	RG-1	RG-2	RG-L
<i>Performance of models without pre-training</i>							
Lead-3	-	40.42	17.62	36.67	16.30	1.60	11.95
PTGEN (See et al., 2017)	-	36.44	15.66	33.42	29.70	9.21	23.24
<i>Performance of state-of-the-art models based on pre-trained models of comparable size</i>							
MASS (Song et al., 2019)	L=6, H=1024	42.12	19.50	39.01	39.75	17.24	31.95
BERTSumAbs (Liu and Lapata, 2019)	L=12, H=768	41.72	19.39	38.76	38.76	16.33	31.15
UniLMv2 (Bao et al., 2020)	L=12, H=768	43.16	20.42	40.14	44.00	21.11	36.08
<i>Performance of comparable models based on T5-base</i>							
T5-base (Raffel et al., 2019)	L=12, H=768	42.25	20.22	39.45	43.12	20.84	34.98
T5-base-cont	L=12, H=768	42.49	20.33	39.65	43.32	20.94	35.21
DistilT5-base	L=12, H=768	42.37	20.25	39.53	43.25	20.89	35.14
DenoiseT5-base	L=12, H=768	42.22	20.18	39.41	43.14	20.82	35.03
SSR-base	L=12, H=768	43.53*	20.79*	40.44*	44.05	21.19	35.88

Table 1: Abstractive summarization results. We also present the transformer architecture for the methods using pre-trained models. For example, L=12, H=768 means both the encoder and decoder are built with 12 transformer layers with a hidden size of 768. *The asterisk denotes statistically significant improvement with p-value < 0.05 upon all compared models.

learning rate of 5e-5 with a linear warm-up for the first 8,000 updates.

It is noteworthy that although SSR requires using a pre-trained Seq2Seq model for imperfect span generation, the computation cost of using SSR to improve a Seq2Seq pre-trained model is still considerably smaller than the pre-training cost. This is because SSR requires much smaller training corpus and optimization steps when employed in a continual pre-training setting. This also reduces recent concerns (Strubell et al., 2019; Bender et al., 2021) about the carbon footprint and energy consumption in LM pre-training.

4.2 Tasks and Datasets

Abstractive Summarization aims to rewrite a long document into a short summary. To provide a comparison with the recent work in pre-trained models for this task, we present results on two widely used summarization datasets: CNN/DailyMail (Hermann et al., 2015) and XSum (Narayan et al., 2018), and report evaluation results in terms of ROUGE-1, ROUGE-2 and ROUGE-L (Lin, 2004).

Question Generation is to generate valid and fluent questions according to a given passage and target answers. It can be considered as rewriting a target answer and its surrounding context into a question form. Following previous work (Dong

et al., 2019), we concatenate the passage and an answer as the input of the model to learn to generate the corresponding question in the fine-tuning stage. We use SQUAD (Rajpurkar et al., 2016) dataset to train and test question generation following the data split in (Du and Cardie, 2018). We report evaluation results in terms of BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and CIDEr (Vedantam et al., 2015).

Grammatical Error Correction is a task that rewrites a potentially erroneous input sentence into a fluent sentence that is grammatical error free without changing the original meaning of the input sentence. Following the recent work (Grundkiewicz et al., 2019; Kiyono et al., 2019; Zhou et al., 2020a) in GEC, we use the public Lang-8 (Mizumoto et al., 2011), NUCLE (Dahlmeier et al., 2013), FCE (Yanakoudakis et al., 2011) and W&I+LOCNESS datasets (Bryant et al., 2019; Granger, 1998) for fine-tuning without using any synthetic GEC data, and then evaluate Max-Match (M^2) precision, recall, and $F_{0.5}$ score on the CoNLL-2014 (Ng et al., 2014) test set.

4.3 Compared Models

We compare SSR with the following models:

- **T5**: the original pre-trained text-to-text transformer based on the text infilling objective.

Model	Architecture	Question Generation			GEC		
		BLEU-4	METEOR	CIDEr	P	R	F _{0.5}
<i>Performance of baseline models without pre-training</i>							
Zhang and Bansal (2019)	-	18.37	22.65	46.68	-	-	-
Xfmr-big (Chen et al., 2020)	L=12, H=1024	-	-	-	64.9	26.6	50.4
Xfmr-big + Synthetic Data (Zhou et al., 2020b)	L=12, H=1024	-	-	-	69.1	33.7	57.1
<i>Performance of state-of-the-art models based on pre-trained models of comparable size</i>							
UniLMv2 (Bao et al., 2020)	L=12, H=768	24.43	26.34	51.97	-	-	-
<i>Performance of comparable models based on T5-base</i>							
T5-base (Raffel et al., 2019)	L=12, H=768	23.74	25.95	51.61	68.6	33.5	56.7
T5-base-cont	L=12, H=768	23.93	26.11	51.78	69.6	33.6	57.3
DistilT5-base	L=12, H=768	23.86	25.93	51.64	69.3	33.1	56.9
DenoiseT5-base	L=12, H=768	23.70	25.91	51.58	69.5	33.4	57.1
SSR-base	L=12, H=768	24.35	26.51*	52.11*	70.5*	34.9*	58.7*

Table 2: Question generation and GEC results. We also present the transformer architecture for the methods using transformer models. For example, L=12, H=768 means both the encoder and decoder are built with 12 transformer layers with a hidden size of 768. *The asterisk denotes statistically significant improvement with p-value < 0.05 upon all compared models.

- **T5-cont**: text-to-text transformer initialized by T5 and continually pre-trained with the original text infilling objective with additional training steps. The total number of additional training steps is equal to that of SSR.
- **DistilT5**: the variant that continually pre-trains T5 by text infilling with sequence-level knowledge distillation (Kim and Rush, 2016). This is implemented by using the imperfect text spans generated by T5-large as target outputs for text infilling. DistilT5-small and DistilT5-base are similar to conventional sequence-level knowledge distillation while DistilT5-large can be viewed as continually pre-trained with self-distillation.
- **DenoiseT5**: the variant that injects rule-based noises into plain text and continually pre-train a T5 model to output the original text. The rule-based noises include token shuffling, deletion, and replacement. We adopt the same noise strategy as described in Wang et al. (2019).

For reference, we also compare against two state-of-the-art base-sized pre-trained models for NLG including MASS (Song et al., 2019) and UniLMv2 (Bao et al., 2020).

4.4 Experimental Results

We first present experimental results of SSR-base and comparable baselines on different datasets.

Then we show additional results of SSR-small and SSR-large for further analysis.

Summarization Results According to Table 1, it is observed that SSR-base substantially improves the original T5-base model and its continual training variants on both CNN/DM and XSum datasets, and achieves state-of-the-art results for the models of the same size in the abstractive summarization benchmarks. It is notable that our models are only continually pre-trained on a relatively small dataset for only a few number of updates. This confirms the potential of our approach as a general “plug-and-play” approach for improving various kinds of sequence-to-sequence pre-trained models. In contrast, using T5-large as a teacher model fails to improve the training of a T5-base student with sequence-level knowledge distillation. This shows SSR can better exploit the capability of a large Seq2Seq pre-trained model to improve a smaller one, indicating its potential to serve as a model compression technique for Seq2Seq pre-trained models.

Question Generation and GEC Results Similar results are observed for question generation and GEC tasks, as shown in Table 2: SSR-base substantially outperforms all the other T5 variants and achieves comparable or even better results than the other base-size pre-trained models. Surprisingly, continually pre-training T5-base with SSR can achieve significant improvement over a transformer-big model pre-trained on rule-based

Model	CNN/DM		
	RG-1	RG-2	RG-L
T5-large	43.09	20.68	40.15
T5-large-cont	43.14	20.71	40.21
DistilT5-large	43.05	20.63	40.07
SSR-large	43.65*	20.98*	40.69*

Table 3: Abstractive summarization results on CNN/DailyMail for SSR-large and corresponding T5 models of the same size. *The asterisk denotes statistically significant improvement with p-value < 0.05 upon all compared models.

synthetic data. We attribute this to the closer relationship between the task of GEC and our proposed SSR objective and more diverse grammatical errors introduced by the machine-generated spans. Interestingly, we observe the improvement of SSR on the GEC task is even more significant than that on question generation and summarization datasets, because SSR is intuitively more similar to the challenge of GEC which can be well addressed by span correction (Chen et al., 2020).

4.5 Analysis

Impact of Model Size To analyze the effectiveness of the proposed SSR objective for Seq2Seq pre-trained models with different sizes, we report the performance comparison of small-size and large-size SSR and different T5-based baselines. Note that we focus on analysis of SSR on the same T5 backbone model and do not compare against other large-sized Seq2Seq PTLMs because they are pre-trained with different data and number of steps, thus are not comparable with our models.

We present the results of large-size models and small-size models in Table 3 and Table 4, respectively.³ We find that the sequence span rewriting objective improves both large-size and small-size models. However, the improvement upon small-size models is significantly larger than that upon large-size models. This suggests that our method is more effective when the infilling model is significantly larger than the rewriting model. The performance of SSR-small is also significantly better than DistilT5-small sequence-level knowledge distillation. That indicates SSR’s potential on exploiting the knowledge from large pre-trained Seq2Seq

³We do not compare against the variant with the denoising-based objective since its performance is consistently lower than the baseline in the previous experiments.

Model	CNN/DM		
	RG-1	RG-2	RG-L
T5-small	40.22	19.36	37.85
T5-small-cont	40.43	19.55	38.08
DistilT5-small	40.38	19.49	38.01
SSR-small	41.95*	20.06*	39.01*

Table 4: Abstractive summarization results on CNN/DailyMail for SSR-small and corresponding T5 models of the same size. *The asterisk denotes statistically significant improvement with p-value < 0.05 upon all compared models.

Model	Imperfect Span Generator	CNN/DM		
		RG-1	RG-2	RG-L
T5-base	-	42.25	20.22	39.45
SSR-base	T5-base	42.78	20.51	39.97
SSR-base	T5-large	43.47	20.74	40.37
T5-small	-	40.22	19.36	37.85
SSR-small	T5-base	41.03	19.74	38.68
SSR-small	T5-large	41.95	20.06	39.01

Table 5: Abstractive summarization results on CNN/DailyMail for SSR with imperfect span generator of different sizes.

transformers to improve the training of smaller models in a task-agnostic fashion.

Impact of Imperfect Span Generator We also investigate the impact of the size of the imperfect span generator. This time, we generate imperfect text spans for pre-training using T5-base model and continually pre-train SSR-base and SSR-small. The results are shown in Table 5. We find that our approach performs better with a larger imperfect span generator, which seems in contradiction to the findings in the replaced token detection objective introduced in ELECTRA (Clark et al., 2020). We suspect the reason is that the task of span-level infilling is more challenging than its token-level counterpart. Therefore, a small imperfect span generator may not be powerful enough to generate imperfect text spans that are meaningful and of relatively high quality. Consequently, the rewriting model may simply learn to ignore the imperfect spans and the SSR objective will degrade into text infilling. Moreover, we can see that the improvement yielded by the SSR objective is more significant when the size of the imperfect span generator is larger than the rewriting model that we aim to train. This confirms that SSR can effectively exploit the knowledge of a large model to better train a smaller

Model	CNN/DM		
	RG-1	RG-2	RG-L
SSR-base	43.53*	20.79*	40.47*
No curriculum	43.26	20.53	40.14
Anti-curriculum	43.09	20.48	40.01
Loss-only curriculum	43.40	20.67	40.25
Length-only curriculum	43.43	20.71	40.35

Table 6: Ablation study results on CNN/DailyMail for SSR-base with different curriculum learning strategies. *The asterisk denotes statistically significant improvement with p-value < 0.05 upon all compared ablation.

one. Interestingly, we find that using imperfect spans generated by T5-base to continually pre-train T5-base can still improve the performance, which is similar to the case of self-distillation (Furlanello et al., 2018).

Impact of Curriculum Pre-training We then analyze the effectiveness of the proposed curriculum pre-training technique. We continually pre-train SSR-base with three variants of the proposed curriculum pre-training method: *No curriculum* denotes the variant without curriculum pre-training; *Anti-curriculum* denotes the variant where pre-training starts with difficult examples; *Loss-only* and *Length-only curriculum* denote a curriculum based solely on per-token loss and the length of imperfect span, respectively. The results are shown in Table 6. We find that pre-training SSR from relatively easy examples to hard examples statistically significantly improve its performance on downstream tasks. More specifically, we find that scheduling the training examples by their length is slightly more effective than by per-token loss, while the combination of them can yield further improvements.

5 Discussion

Pre-training via Rewriting We discuss several key advantages of SSR over the conventional text infilling objectives here. **(1) SSR is closer to the downstream sequence transduction tasks.** This is because the model’s prediction is not only based on its bidirectional context but also conditioned on the imperfect spans. In this way, the gap between pre-training and fine-tuning stages, which is introduced by the masked tokens or spans in conventional pre-training objectives, is alleviated. Indeed, many NLG tasks can be viewed as sequence span rewriting problems that rewrite the

input text into another language, more compact format, grammatically correct sentences, or another style. **(2) SSR introduces more diverse noise patterns.** These patterns include paraphrasing and simplification of the text span, missing or redundant information, grammatical errors, and errors in terms of world knowledge or commonsense knowledge. In fact, many of the rewriting patterns introduced by SSR resemble training examples in the downstream tasks. In contrast, conventional self-supervised Seq2Seq pre-training techniques rely on rule-based noise functions like text span masking, token masking, token deletion, token rotation, sentence shuffling, etc. **(3) SSR enables the model to learn from informative examples.** SSR enables the model to learn from informative examples, where the span generator makes an error. This provides more meaningful supervision and is also similar to the idea of active learning (Settles, 2009).

Distillation via Rewriting SSR sheds light on a new perspective of exploiting the knowledge of a large pre-trained model to improve smaller models. Similar to knowledge distillation (KD), this can be achieved by using a large-size teacher model pre-trained with the text infilling objective as the imperfect span generator, and pre-train or refine a small-size student model with the generated data using SSR. Different from conventional KD (Hinton et al., 2015) or sequence-level KD (Kim and Rush, 2016), SSR enables the student model to exploit both teacher outputs and the ground truth at the same time. It is also related to boost learning (Schapire, 2003) and residual learning (He et al., 2016) in a sense that the model only needs to learn the prediction error of the teacher model, instead of the original task, text infilling, which may be too difficult for smaller-size models.

6 Conclusion

We present sequence span rewriting (SSR), a novel self-supervised objective for improving sequence-to-sequence transformers pre-trained with conventional text infilling objectives. SSR introduces more diverse and fine-grained learning signals and also bridges the gap between self-supervised pre-training and task-specific fine-tuning on common NLG datasets. Our experiments on continual T5 pre-training confirm the effectiveness of SSR on improving pre-trained T5 models of different sizes across different tasks and datasets. Also, the large

improvements achieved on small models with a larger imperfect span generator indicates a new perspective of exploiting the knowledge of a large pre-trained model to help train smaller ones.

Ethical Considerations

Our approach is proposed to improve existing sequence-to-sequence pre-training techniques. It does not involve the collection and release of data except that generated by a pre-trained model, nor inference of information or judgments about individuals. That being said, since an improved sequence-to-sequence pre-trained model may be used in various downstream applications, it is still an important future direction to investigate the bias, fairness, and privacy issue in various kinds of pre-trained models.

Acknowledgments

We appreciate all anonymous reviewers and the meta-reviewer, for their insightful comments. Tao Ge is the corresponding author.

References

- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, and Hsiao-Wuen Hon. 2020. [Unilmv2: Pseudo-masked language models for unified language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 642–652. PMLR.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *FAccT*, pages 610–623. ACM.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, volume 382 of *ACM International Conference Proceeding Series*, pages 41–48. ACM.
- Christopher Bryant, Mariano Felice, Øistein E. Andersen, and Ted Briscoe. 2019. [The BEA-2019 shared task on grammatical error correction](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75, Florence, Italy. Association for Computational Linguistics.
- Mengyun Chen, Tao Ge, Xingxing Zhang, Furu Wei, and Ming Zhou. 2020. [Improving the efficiency of grammatical error correction with erroneous span detection and correction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7162–7169, Online. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. [Building a large annotated corpus of learner English: The NUS corpus of learner English](#). In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 22–31, Atlanta, Georgia. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. [Unified language model pre-training for natural language understanding and generation](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13042–13054.
- Xinya Du and Claire Cardie. 2018. [Harvesting paragraph-level question-answer pairs from Wikipedia](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1907–1917, Melbourne, Australia. Association for Computational Linguistics.
- Tommaso Furlanello, Zachary Chase Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. 2018. [Born-again neural networks](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of

- Proceedings of Machine Learning Research*, pages 1602–1611. PMLR.
- Sylviane Granger. 1998. *The computer learner corpus: a versatile new source of data for SLA research*. na.
- Roman Grundkiewicz, Marcin Junczys-Dowmunt, and Kenneth Heafield. 2019. [Neural grammatical error correction systems with unsupervised pre-training on synthetic data](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 252–263, Florence, Italy. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward controlled generation of text](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596. PMLR.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Shun Kiyono, Jun Suzuki, Masato Mita, Tomoya Mizumoto, and Kentaro Inui. 2019. [An empirical study of incorporating pseudo data into grammatical error correction](#). 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 1236–1242, Hong Kong, China. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020a. [Pre-training via paraphrasing](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020b. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. [Commongen: A constrained text generation challenge for generative commonsense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 1823–1840. 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.
- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). 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 3730–3740, Hong Kong, China. 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*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems*

- 32: *Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 14014–14024.
- Tomoya Mizumoto, Mamoru Komachi, Masaaki Nagata, and Yuji Matsumoto. 2011. [Mining revision log of language learning SNS for automated Japanese error correction of second language learners](#). In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 147–155, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. 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.
- 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.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2015. [Fitnets: Hints for thin deep nets](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Robert E Schapire. 2003. The boosting approach to machine learning: An overview. In *Nonlinear estimation and classification*, pages 149–171. Springer.
- Roy Schwartz, Gabriel Stanovsky, Swabha Swayamdipta, Jesse Dodge, and Noah A. Smith. 2020. [The right tool for the job: Matching model and instance complexities](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6640–6651, Online. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Burr Settles. 2009. Active learning literature survey.
- Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. 2020. Q-bert: Hessian based ultra low precision quantization of bert. In *AAAI*, pages 8815–8821.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tiejun Liu. 2019. [MASS: masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5926–5936. PMLR.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [Cider: Consensus-based image description evaluation](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4566–4575. IEEE Computer Society.
- Liang Wang, Wei Zhao, Ruoyu Jia, Sujian Li, and Jingming Liu. 2019. [Denoising based sequence-to-sequence pre-training for text 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 4003–4015, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. 2020. [BERT-of-theseus: Compressing BERT by progressive module replacing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7859–7869, Online. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. [A new dataset and method for automatically grading ESOL texts](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 180–189, Portland, Oregon, USA. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. [Defending against neural fake news](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 9051–9062.
- Shiyue Zhang and Mohit Bansal. 2019. [Addressing semantic drift in question generation for semi-supervised question answering](#). 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 2495–2509, Hong Kong, China. Association for Computational Linguistics.
- Wangchunshu Zhou, Tao Ge, Chang Mu, Ke Xu, Furu Wei, and Ming Zhou. 2020a. [Improving grammatical error correction with machine translation pairs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 318–328, Online. Association for Computational Linguistics.
- Wangchunshu Zhou, Tao Ge, and Ke Xu. 2020b. [Pseudo-bidirectional decoding for local sequence transduction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1506–1511, Online. Association for Computational Linguistics.
- Wangchunshu Zhou, Dong-Ho Lee, Ravi Kiran Selvam, Seyeon Lee, and Xiang Ren. 2021a. [Pre-training text-to-text transformers for concept-centric common sense](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. 2020c. [Bert loses patience: Fast and robust inference with early exit](#). In *NeurIPS 2020*.
- Wangchunshu Zhou, Canwen Xu, and Julian J. McAuley. 2021b. [Meta learning for knowledge distillation](#). *CoRR*, abs/2106.04570.
- Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 19–27. IEEE Computer Society.