

Uncertainty Determines the Adequacy of the Mode and the Tractability of Decoding in Sequence-to-Sequence Models

Felix Stahlberg¹ and Ilia Kulikov^{2*} and Shankar Kumar¹

¹Google Research, ²New York University
{fstahlberg, shankarkumar}@google.com,
ik1147@nyu.edu

Abstract

In many natural language processing (NLP) tasks the same input (e.g. source sentence) can have multiple possible outputs (e.g. translations). To analyze how this ambiguity (also known as *intrinsic uncertainty*) shapes the distribution learned by neural sequence models we measure sentence-level uncertainty by computing the degree of overlap between references in multi-reference test sets from two different NLP tasks: machine translation (MT) and grammatical error correction (GEC). At both the sentence- and the task-level, intrinsic uncertainty has major implications for various aspects of search such as the inductive biases in beam search and the complexity of exact search. In particular, we show that well-known pathologies such as a high number of beam search errors, the inadequacy of the mode, and the drop in system performance with large beam sizes apply to tasks with high level of ambiguity such as MT but not to less uncertain tasks such as GEC. Furthermore, we propose a novel exact n -best search algorithm for neural sequence models, and show that intrinsic uncertainty affects model uncertainty as the model tends to overly spread out the probability mass for uncertain tasks and sentences.

1 Introduction

With the advent of deep learning, many applications of machine learning have converged on a similar set of methods and models. For example, the Transformer (Vaswani et al., 2017) sequence-to-sequence architecture is ubiquitous in various fields of natural language processing (NLP) such as machine translation (MT), grammatical error correction (GEC), speech recognition (Karita et al., 2019), etc., and has also been applied successfully to other tasks such as computer vision (Dosovitskiy et al., 2021). Recent large pre-trained NLP models

such as BERT (Devlin et al., 2019), GPT-3 (Brown et al., 2020), T5 (Raffel et al., 2020), RoBERTa (Liu et al., 2019), and XLNet (Yang et al., 2019) are all based on the Transformer, with relatively minor changes to the architecture itself.

We show that despite this architectural uniformity the learned distribution over sequences has strikingly different characteristics for different NLP tasks. Inspired by Ott et al. (2018) we identify *intrinsic uncertainty* – the nature of some NLP tasks to allow multiple viable outputs for a given input¹ – to be a major factor that shapes the search space of Transformer models and determines its tractability. In machine translation (MT) – a task known to have high intrinsic uncertainty (Padó et al., 2009; Dreyer and Marcu, 2012; Ott et al., 2018) – Transformer models suffer from a high number of beam search errors (Stahlberg and Byrne, 2019), an inadequacy of the mode (Eikema and Aziz, 2020), and translation performance degradation with large beam sizes (Koehn and Knowles, 2017) (also known as the “beam search curse”). In contrast, for the correction of writing errors in text (grammatical error correction – GEC) (Brockett et al., 2006), a task with a lower level of uncertainty (Bryant and Ng, 2015), none of these pathologies are evident. This pattern holds even at the sequence-level: input sentences with high uncertainty tend to result in more search errors and a less tractable search space. To study the influence of uncertainty on sequences around the mode, we propose an exact n -best search algorithm for neural sequence models. We show that the probability mass covered by the n -best candidates differs markedly between certain and uncertain tasks and sentences, which shows that intrinsic uncertainty also affects the spread of probability mass and thus the model uncertainty. We confirm recent work showing that beam search has drawbacks as a decoding scheme for MT. Nevertheless,

*Research done during internship at Google Research, now at Meta AI.

¹This is sometimes referred to as *aleatoric* uncertainty in the literature (Der Kiureghian and Ditlevsen, 2009).




Figure 1: Average uncertainty factor u for GEC (blue) and English-German MT (purple) grouped by the sentence length. The error bars show the standard error of the mean (SEM).

it is effective for GEC, a problem where modes are adequate, search errors are rare, and the n -best lists cover a large fraction of the probability mass.

2 Measuring Intrinsic Uncertainty

Intrinsic uncertainty refers to the inherent nature of some NLP tasks to allow for more than one feasible output for a given input. For example, intrinsic uncertainty in MT stems from the fact that there are often several semantically equivalent translations for the same source sentence, or that the translation into a highly inflected language is sometimes underspecified (Ott et al., 2018). Studies have shown that even for tasks like GEC, annotators do not always agree (Tetreault and Chodorow, 2008; Rozovskaya and Roth, 2010; Bryant and Ng, 2015), but the level of intrinsic uncertainty is arguably lower than for MT because there is a limited number of ways to correct an ungrammatical sentence.

We propose a simple way to measure sentence-level output uncertainty by making use of multi-reference test sets. For an n -way annotated sentence with references $\mathbf{y}_1, \dots, \mathbf{y}_n$ we define the uncertainty u as the average relative edit distance between two references:

$$\begin{aligned}
 u &:= \frac{1}{\underbrace{\frac{1}{n} \sum_{i=1}^n |\mathbf{y}_i|}_{\text{Avg. ref. length}}} \underbrace{\frac{\sum_{i=1}^{n-1} \sum_{j=i+1}^n d_{\text{edit}}(\mathbf{y}_i, \mathbf{y}_j)}{\frac{n(n-1)}{2}}}_{\text{Avg. edit distance between refs.}}} \\
 &= \frac{2}{(n-1) \sum_{i=1}^n |\mathbf{y}_i|} \sum_{i=1}^{n-1} \sum_{j=i+1}^n d_{\text{edit}}(\mathbf{y}_i, \mathbf{y}_j)
 \end{aligned} \tag{1}$$

where $d_{\text{edit}}(\cdot, \cdot)$ denotes the Levenshtein distance. Fig. 1 presents this uncertainty score for one MT test set and two GEC test sets. *MT-ende* is the official WMT19 English-German test set (Barrault et al., 2019) paired with the additional human-

annotated “newstest2019 AR” references provided by Freitag et al. (2020).² *GEC-conll14* uses the 10 references published by Bryant and Ng (2015) for the CoNLL-2014 shared task on GEC (Ng et al., 2014), and *GEC-jfleg* is a 4-reference GEC test set that represents “a broad range of language proficiency levels” (Napoles et al., 2017). Our uncertainty measure reflects our intuition that MT is a significantly more uncertain task than GEC.³ For both tasks the uncertainty increases with the sentence length as longer sentences typically have more feasible mappings than shorter ones. We use the edit distance rather than task-specific metrics like BLEU (Papineni et al., 2002) or BLEURT (Sellam et al., 2020) since they are designed to be robust against uncertainty effects such as reordering or semantically equivalent references, precisely the kinds of effects we aim to capture with u . We follow Bryant and Ng (2015) by not using inter-annotator agreement statistics like Cohen’s κ (Cohen, 1960) since they are more appropriate for the classification into single, well-defined categories.

3 Mode-seeking Search

Neural sequence-to-sequence models define a probability distribution $P(\mathbf{y}|\mathbf{x})$ over target sequences \mathbf{y} given a source sequence \mathbf{x} :

$$\log P(\mathbf{y}|\mathbf{x}) = \sum_{j=1}^{|\mathbf{y}|} \log P(y_j | y_1^{j-1}, \mathbf{x}). \tag{2}$$

Sequences are typically computed over a subword (Sennrich et al., 2016; Kudo and Richardson, 2018) vocabulary \mathcal{V} and end with a special end-of-sentence symbol $\langle /s \rangle$:

$$\mathbf{x}, \mathbf{y} \in \{\mathbf{z} \cdot \langle /s \rangle | \mathbf{z} \in \mathcal{V}^*\} \tag{3}$$

where \mathcal{V}^* is the Kleene closure over \mathcal{V} which includes the empty sequence ϵ . Since sequence models are usually trained to maximize the probability of the sequences in the training set, a common strategy to use such a model for inference is to search for the most likely output sequence \mathbf{y}^* , also known as the *mode* of the model distribution:⁴

$$\mathbf{y}^* = \arg \max_{\mathbf{y}} P(\mathbf{y}|\mathbf{x}). \tag{4}$$

²The AR references are created from scratch, unlike the other paraphrasing references by Freitag et al. (2020).

³The mean u value differs significantly between GEC and MT in each length bucket (t -test p -value of less than 0.0001).

⁴In a Bayesian framework this is often referred to as maximum a posteriori (MAP) inference.

Algorithm 1 NbestDFS($n, \mathbf{x}, \mathbf{y}, p, \gamma, \tilde{Y}$)

Input: n : Search for n global best sequences
 \mathbf{x} : Source sequence
 \mathbf{y} : Target prefix (default: ϵ)
 p : $\log P(\mathbf{y}|\mathbf{x})$ (default: 0.0)
 γ : Lower bound (default: $-\infty$)
 \tilde{Y} : Priority queue (default: \emptyset)

- 1: **if** $y_{|\mathbf{y}|} = \langle /s \rangle$ **then**
- 2: $\tilde{Y}.push((p, \mathbf{y}))$
- 3: **if** $|\tilde{Y}| > n$ **then**
- 4: $(\gamma, -) \leftarrow \tilde{Y}.pop()$
- 5: **end if**
- 6: **return** (γ, \tilde{Y})
- 7: **end if**
- 8: **for all** $w \in \mathcal{V}$ **do**
- 9: $p' \leftarrow p + \log P(w|\mathbf{x}, \mathbf{y})$
- 10: **if** $p' > \gamma$ **then**
- 11: $(\gamma, \tilde{Y}) \leftarrow \text{NbestDFS}(n, \mathbf{x}, \mathbf{y} \cdot w, p', \gamma, \tilde{Y})$
- 12: **end if**
- 13: **end for**
- 14: **return** (γ, \tilde{Y})

Eq. 4 is usually approximated using beam search. For analysis purposes, Stahlberg and Byrne (2019) proposed an exact depth-first search (DFS) algorithm that is guaranteed to find the mode.

4 N -best Search

In addition to our investigations into the mode we also examine the cumulative probability mass that is covered by the n best hypotheses. If a hypothesis set covers a large fraction of the entire probability mass it approximates the full model distribution well. Approximating the full model distribution is useful for various methods such as minimum risk training (Shen et al., 2016), reinforcement learning (Williams, 1992; Ranzato et al., 2015), minimum Bayes risk decoding (Kumar and Byrne, 2004; Stahlberg et al., 2017; Eikema and Aziz, 2020), etc. Ott et al. (2018) argued that the fraction of probability mass which is covered by a fixed number of candidates reflects the *model* uncertainty on the sequence level. We show that this model uncertainty is in line with our notion of intrinsic uncertainty that we measure with u (Sec. 2). To that end, we propose a generalization of the exact search algorithm of Stahlberg and Byrne (2019) that is able to find the n global best hypotheses rather than the single best one. Similarly to the single-best algorithm, we use the monotonicity of neural sequence model scores:

$$\forall j \in [2, |\mathbf{y}|] : \log P(y_1^{j-1}|\mathbf{x}) > \log P(y_1^j|\mathbf{x}). \quad (5)$$

Stahlberg and Byrne (2019) keep track of the best

Parameter	Value
Attention dropout rate	0.1
Attention layer size	512
Batch size	256
Dropout rate	0.1
Embedding size	512
MLP dimension	2,048
Number of attention heads	8
Number of layers	6
Total number of parameters	121M

Table 1: Transformer hyper-parameters.

Language pair	#Training sentence pairs	
	Unfiltered	Filtered
German-English	39M	33M
Finnish-English	6.6M	5.5M
Lithuanian-English	2.3M	2.0M

Table 2: MT training set sizes.

complete (i.e. ending with the end-of-sentence symbol $\langle /s \rangle$) hypothesis score during search, and use it to safely prune entire subspaces using Eq. 5. In contrast, we keep track of the n -th best complete hypothesis score by keeping the n best complete hypotheses in a priority queue. Our exact n -best search algorithm is listed in Algorithm 1. Note that we recover the DFS scheme of Stahlberg and Byrne (2019) with $n = 1$.

5 Experimental Setup

We trained four Transformer neural machine translation (NMT) models (Table 1) – English-German (*MT-ende*), German-English (*MT-deen*), Finnish-English (*MT-fien*), and Lithuanian-English (*MT-lten*) – on the WMT19 (Barrault et al., 2019) training sets as provided by TensorFlow Datasets.⁵ We selected these language pairs to experiment with different training set sizes (Table 2). The MT training sets were filtered using language ID and simple length-based heuristics, and split into subwords using joint 32K SentencePiece (Kudo and Richardson, 2018) models. For training our GEC model we used the hyper-parameters from Table 1 and followed the three-stage training recipe of Stahlberg and Kumar (2021) using the 32K SentencePiece model from Raffel et al. (2020). All our models were trained until convergence on the development set using the LAMB (You et al., 2020) optimizer in JAX (Bradbury et al., 2018) by minimizing cross-entropy without label smoothing. Our NMT models are evaluated on the WMT19 test sets (Barrault

⁵https://www.tensorflow.org/datasets/catalog/wmt19_translate

System	ende	deen	fien	lten
Xia et al. (2019)	44.9	42.8	31.9	35.6
Our baselines	39.6	39.7	27.7	26.9

Table 3: BLEU scores of our NMT baselines and one of the best systems in the WMT19 evaluation campaign – MSRA . MADL (Xia et al., 2019).

System	conll14 (F _{0.5})	jfleg (GLEU)
Lichtarge et al. (2020)	66.8	64.9
Rothe et al. (2021)	68.9	-
Our baseline	60.0	62.1

Table 4: Comparison of our GEC baseline with the best results reported in the literature.

et al., 2019) using SacreBLEU (Post, 2018). Our GEC model is evaluated on the CoNLL14 (Ng et al., 2014, *GEC-conll14*) test set using F_{0.5}-scores computed with the M2 scorer (Dahlmeier and Ng, 2012) and on the JFLEG test set (Napoles et al., 2017, *GEC-jfleg*) using GLEU (Napoles et al., 2015).

6 Results

In this work our focus is to analyze the impact of intrinsic uncertainty on search. Thus we keep our setup simple, reproducible, and computationally economical rather than obtain new state-of-the-art results. Nevertheless, Tables 3 and 4 show that our baselines are not unreasonably far off from the best results in the literature given that the systems we compare with are often highly engineered and use many more parameters. Xia et al. (2019) used various techniques like back-translation, ensembling, dual learning, MASS pre-training, architecture search, larger models, etc. to improve their systems, and Rothe et al. (2021) used a 11B parameters T5 (Raffel et al., 2020) model.

6.1 Finding the Most Likely Hypothesis

Even though alternative decision rules like MBR have recently received some attention in the NMT literature (Eikema and Aziz, 2020; Müller and Senrich, 2021), mode-seeking decoding schemes such as beam search or Nucleus sampling (Holtzman et al., 2020) are by far the most common choices. In this section we explore how uncertainty changes the mode and the ability of beam search to find it.

A well-known pathology of NMT models is the “beam search curse” (Koehn and Knowles, 2017): Increasing the beam size improves the predictive log-probabilities of the hypotheses, but it leads to worse translation quality due to the NMT model error of preferring short translations. We replicate




Figure 2: Relative beam search improvements over greedy search. MT quality degrades with large beam sizes, but GEC saturates after a beam size of 10.

this result in Fig. 2: BLEU scores for MT initially improve over greedy search at smaller beam sizes but after reaching a peak at beam size of 4, we observe a dramatic drop in BLEU. The trajectory of the blue curves (GEC) is markedly different: the performance does *not* drop for large beams but saturates instead. The beam search curse affects tasks with high intrinsic uncertainty like MT but spares more certain tasks like GEC although both tasks use the same neural Transformer architecture.

To determine why the beam size affects NMT and GEC so differently we ran the exact decoding algorithm of Stahlberg and Byrne (2019) to find the global best hypotheses and counted search errors, i.e. the number of sentences in the test set for which beam search does not find the global best sequence. Our results confirm the findings of Stahlberg and Byrne (2019) that increasing the beam sizes leads to fewer NMT search errors (Fig. 3). Among our MT language pairs, English-German (MT-ende) suffers the most from the beam search curse and the proportion of search errors in the test set. This is possibly because translation from English to German typically results in a longer sequence and thus more uncertainty. GEC differs significantly from NMT in the total number of search errors. For MT, even with a very large beam size of 500, beam search does not find the mode for more than 20% of the sentences in any language pair. In contrast for GEC, we do not observe any search errors for beam sizes larger than 10. This suggests that task uncertainty determines the tractability of the search space and particularly the search for the mode.

Uncertainty also determines the computational costs of exact search. To abstract away from hardware and implementation details, we measure the time complexity of exact search by counting the




Figure 3: Number of beam search errors.

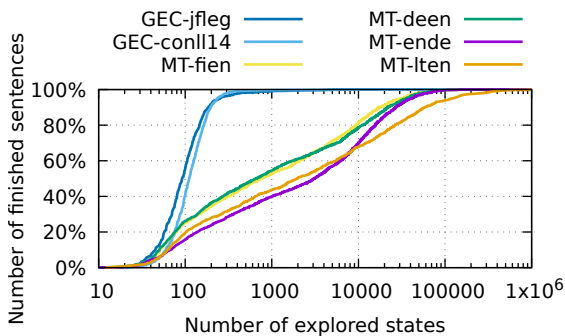


Figure 4: Number of states exact search needs to explore in order to find and verify the mode.

number of explored states, i.e. the number of forward passes through the model, which is identical to the number of recursive calls of Algorithm 1.⁶

Fig. 4 plots the fraction of sentences in the test set for which the exact search explores a certain maximum number of states to terminate. For example, exact search returned the mode for around 50% of the MT sentences after exploring no more than 1000 states. With the same computational budget, however, it was able to find the mode for nearly 100% of the GEC sentences (blue curves). For some of the MT sentences, exact search needed to explore around 100K states, or even more in the case of Lithuanian-English (orange curve).

Sentence-level uncertainty In the previous paragraph we showed that MT, a task with high intrinsic uncertainty, suffers from more beam search errors and a less tractable search space than GEC, a task with relatively low intrinsic uncertainty. Figs. 5 and 6 demonstrate that this pattern is not only present at the task-level but also at the sentence-level. First, the bar charts show that there is a general trend towards more search errors and more explored states

⁶For example, the number of explored states in standard beam search is the beam size times the target sequence length.

for longer sentences. Longer input sentences often result in higher entropy distributions (i.e. more uncertainty) since there are usually more ways to map a long sentence than a short one. We also see a pattern within each group, i.e. within a reference length interval, that shows that sentences with higher uncertainty u result in more search errors and a longer exact search runtime even when compared to other sentences with similar lengths. Table 5 lists the test set level correlation coefficients.

6.2 The Spread of Probability Mass

We argued in Sec. 4 that the ability to approximate the entire search space with a fixed set of candidates can be useful in training (Shen et al., 2016; Williams, 1992; Ranzato et al., 2015) and decoding (Kumar and Byrne, 2004; Eikema and Aziz, 2020), and proposed an exact n -best search algorithm. However, finding the exact n -best hypotheses is computationally much more expensive than finding the single-best hypothesis (mode). Therefore, to keep the runtime under control, we stopped n -best decoding after 1M explored states. Fig. 7 shows that the 1M threshold is not reached for $n = 1$ for any sentence: it was always possible to find and verify the mode. We can guarantee that the $n = 100$ best candidates returned by our algorithm are indeed the global best ones for around 90% of the MT-deen sentences (right end of the green curve in Fig. 7). The blue curves in Fig. 7 suggest that as before the GEC search space is much more tractable given that our exact n -best search algorithm was able to find the 100 global best hypotheses for all GEC sentences before reaching 1M explored states. Indeed, Fig. 8 shows that exact 100-best search terminated with fewer than 10K explored states for almost all GEC sentences while the pruning criterion in Eq. 5 is much less effective for the NMT search space (green curves in Fig. 8).

The cumulative probability mass of the set returned by exact n -best search is an upper bound for the cumulative probability mass of any hypothesis set with a cardinality of n . Despite the high number of search errors (Fig. 3), the probability mass covered by the n -best beam search hypotheses is very close to this upper bound. Fig. 9 shows that for $n = 100$ that difference is less than 0.001 for all setups except MT-fien. Since the difference in probability mass is negligible we ran our subsequent investigations of probability mass with beam search instead of exact search to save computational costs.

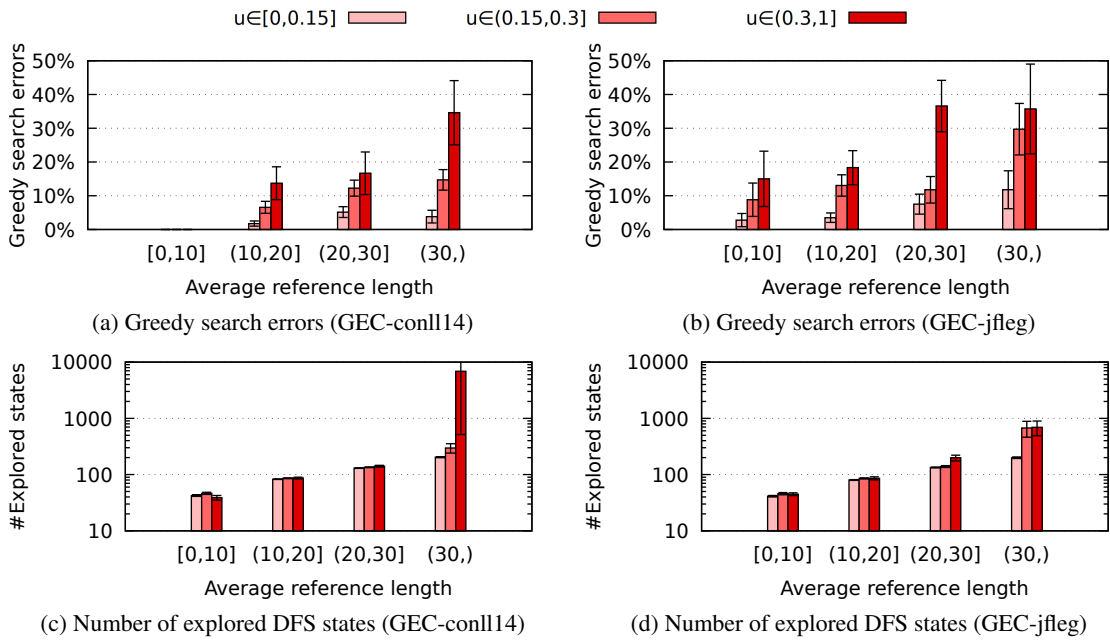


Figure 5: The impact of sentence length and uncertainty u on the number of greedy search errors and the number of explored states by exact search for GEC. The error bars show the SEM.

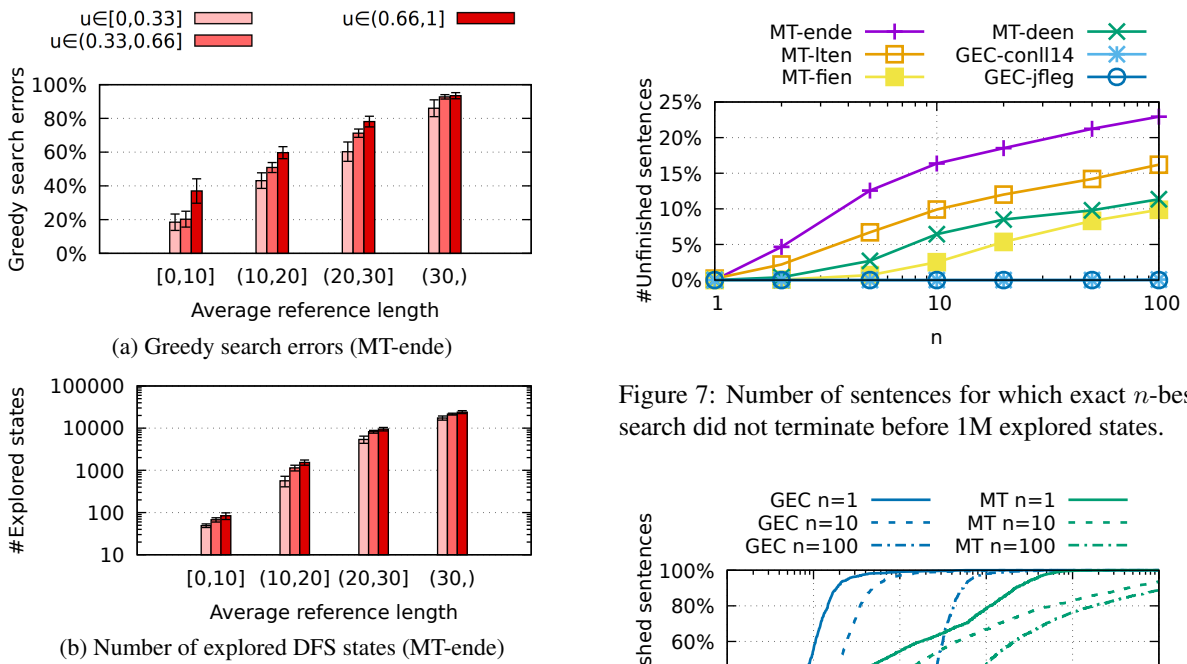


Figure 6: The impact of sentence length and uncertainty u on the number of greedy search errors and the number of explored states by exact search for MT. The error bars show the SEM.




Figure 7: Number of sentences for which exact n -best search did not terminate before 1M explored states.




Figure 8: Number of states exact n -best search needs to explore in order to terminate for GEC-jfleg and MT-deen.

Fig. 10 visualizes the difference between NMT and GEC in terms of the probability mass covered by the beam search hypotheses. We confirm the finding of Ott et al. (2018); Eikema and Aziz (2020) that the NMT distribution is rather flat: even 1000 MT candidates cover only 20% of the probability

mass on average. In GEC, however, the model assigns twice as much probability (40%) to the *single* best hypothesis on average (left end of the




Figure 9: Difference in cumulative probability mass between the global n best hypothesis set returned by exact n -best search and the n -best list returned by beam search with different beam sizes.

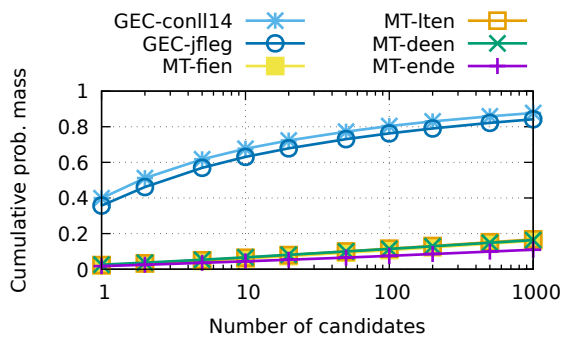


Figure 10: Average probability mass covered by the n -best list from beam search for beam sizes between 1 and 1000.

ρ between u and...	GEC		MT
	conll14	jfleg	ende
Greedy search errors	0.18	0.19	0.24
#Explored DFS states	0.20	0.18	0.19
Cumul. prob. mass	-0.23	-0.51	-0.53

Table 5: Spearman’s rank correlation coefficient ρ between the uncertainty u and the number of greedy search errors, the number of explored DFS states, and the 100-best cumulative probability mass. All correlations are significant with a p -value of less than 0.00001.

blue curves in Fig. 10). Fig. 11 provides even more insight: A beam size of 1000 covers 40% of the probability mass for nearly all sentences in the GEC test sets. Even more practical beam sizes of 10 cover more than half of the probability mass for around 75% of the GEC-conll14 sentences. The same plot looks very different for MT (Fig. 12): Covering half the probability mass is only possible for a tiny fraction of the MT sentences.

Sentence-level uncertainty In Sec. 6.1 we reported that the effects caused by intrinsic uncertainty on the ability to find the mode are visible




Figure 11: The number of sentences for which the total probability mass contained in a beam search n -best list with beam sizes of 1, 10, 100, 1000 is a certain fraction of the total probability mass (GEC).




Figure 12: The number of sentences for which the total probability mass contained in a beam search n -best list with beam sizes of 1, 10, 100, 1000 is a certain fraction of the total probability mass (MT).

at both the task- and the sentence-levels. Similarly, we can track down our observations about how uncertainty determines the probability mass of n -best lists at the sentence level. Fig. 13 shows that the cumulative probability mass in the n -best list decreases for longer sentences as the mappings of long sentences are more uncertain. Again, the trend within a group in Fig. 13 suggests that even among sentences with similar lengths, n -best lists for uncertain sentences (higher u) accumulate less probability mass. We make analogous observations for NMT (Fig. 14), although the total n -best probability mass is much smaller than for GEC.

7 Related Work

Ambiguity is one of the core challenges in MT, a fact that is supported (inter alia) by the long history of designing evaluation metrics that are robust against it (Papineni et al., 2002; Banerjee and Lavie, 2005; Sellam et al., 2020). In this work we examine the impact of ambiguity on the NMT search space, and show how it is related to various well-




Figure 13: The impact of sentence length and uncertainty u on the cumulative probability mass of the 100-best list from beam search for GEC. The error bars show the SEM.




Figure 14: The impact of sentence length and uncertainty u on the cumulative probability mass of the 100-best list from beam search for MT. The error bars show the SEM.

known issues of NMT models like the beam search curse (Koehn and Knowles, 2017), a pathology that has also been linked to the local normalization in sequence models (Soutsov and Sarawagi, 2016; Murray and Chiang, 2018) or poor model calibration (Kumar and Sarawagi, 2019).

Our work is heavily inspired by Ott et al. (2018) who analyzed different kinds of uncertainty in NMT. In particular, they found that NMT spreads out the probability mass over a large number of candidates, and connected the beam search curse with uncertainty. We confirm their results and extend their line of research along the following directions: We introduce a measure for uncertainty in multi-reference test sets, and show that the negative effects of uncertainty are visible even on the sentence level. Second, we propose an exact n -best search algorithm and demonstrate how it can be used to analyze the spread of probability mass. Third, we focus not only on MT but also on GEC.

Stahlberg and Byrne (2019) showed that beam search errors often obscure the length deficiency of the NMT modes, and reducing search errors by

using large beams exposes this model error. In this work, we found that these mechanics are limited to NMT: GEC does not suffer from the beam search curse since search errors are rare and modes are not too short. Eikema and Aziz (2020) suggested that picking a hypothesis based solely on probability is erratic because NMT spreads out the probability mass over a large set of hypotheses with similar probabilities. Therefore, alternative approaches that in addition to the probabilities incorporate MT-specific metrics such as BLEU (Papineni et al., 2002) or BLEURT (Sellam et al., 2020) have recently been in focus of research, including minimum Bayes risk decoding (Eikema and Aziz, 2020, 2021; Müller and Sennrich, 2021), Monte-Carlo tree search (Leblond et al., 2021), and energy-based (Bhattacharyya et al., 2021) or discriminatively trained (Lee et al., 2021) rerankers. Our work on how uncertainty determines the spread of probability mass is relevant to those approaches.

8 Conclusion

We identified a major culprit behind various inference-related issues in sequence-to-sequence models such as the intractability of the search space, degenerate large beam or exact search outputs and the large spread in probability mass over the output space. This factor is intrinsic uncertainty – the existence of multiple ways to correctly map an input sequence. We measured the intrinsic uncertainty of input sentences as the degree of agreement between multiple references and showed that ambiguous sentences typically result in a higher number of beam search errors and an exceedingly flat output distribution. We also find that known NMT pathologies such as the beam search curse or inadequate modes do not extend to less ambiguous tasks like GEC despite using the same neural architecture.

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.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 conference on machine translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Sumanta Bhattacharyya, Amirmohammad Rooshenas, Subhajit Naskar, Simeng Sun, Mohit Iyyer, and Andrew McCallum. 2021. [Energy-based reranking: Improving neural machine translation using energy-based models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4528–4537, Online. 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. 2018. [JAX: Composable transformations of Python+NumPy programs](#).
- Chris Brockett, William B. Dolan, and Michael Gamon. 2006. [Correcting ESL errors using phrasal SMT techniques](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 249–256, Sydney, Australia. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- 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.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- 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.
- Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural safety*, 31(2):105–112.
- 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.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Markus Dreyer and Daniel Marcu. 2012. [HyTER: Meaning-equivalent semantics for translation evaluation](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 162–171, Montréal, Canada. Association for Computational Linguistics.
- Bryan Eikema and Wilker Aziz. 2020. [Is MAP decoding all you need? the inadequacy of the mode in neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4506–4520, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Bryan Eikema and Wilker Aziz. 2021. Sampling-based minimum bayes risk decoding for neural machine translation. *arXiv preprint arXiv:2108.04718*.
- Markus Freitag, George Foster, David Grangier, and Colin Cherry. 2020. [Human-paraphrased references improve neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation*,

- pages 1183–1192, Online. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Shigeeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyang Jiang, Masao Someki, Nelson Enrique Yalta Soplín, Ryuichi Yamamoto, Xiaofei Wang, et al. 2019. A comparative study on Transformer vs RNN in speech applications. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 449–456. IEEE.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- 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.
- Aviral Kumar and Sunita Sarawagi. 2019. Calibration of encoder decoder models for neural machine translation. *arXiv preprint arXiv:1903.00802*.
- Shankar Kumar and William Byrne. 2004. [Minimum Bayes-risk decoding for statistical machine translation](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 169–176, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Rémi Leblond, Jean-Baptiste Alayrac, Laurent Sifre, Miruna Pîslar, Jean-Baptiste Lespiau, Ioannis Antonoglou, Karen Simonyan, and Oriol Vinyals. 2021. Machine translation decoding beyond beam search. *arXiv preprint arXiv:2104.05336*.
- Ann Lee, Michael Auli, and Marc’Aurelio Ranzato. 2021. [Discriminative reranking for neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7250–7264, Online. Association for Computational Linguistics.
- Jared Lichtarge, Chris Alberti, and Shankar Kumar. 2020. [Data weighted training strategies for grammatical error correction](#). *Transactions of the Association for Computational Linguistics*, 8:634–646.
- 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*.
- Mathias Müller and Rico Sennrich. 2021. [Understanding the properties of minimum Bayes risk decoding in neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 259–272, Online. Association for Computational Linguistics.
- Kenton Murray and David Chiang. 2018. [Correcting length bias in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 212–223, Brussels, Belgium. Association for Computational Linguistics.
- Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. 2015. [Ground truth for grammatical error correction metrics](#). 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 588–593, Beijing, China. Association for Computational Linguistics.
- 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.
- 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.
- Sebastian Padó, Daniel Cer, Michel Galley, Dan Jurafsky, and Christopher D Manning. 2009. Measuring machine translation quality as semantic equivalence: A metric based on entailment features. *Machine Translation*, 23(2-3):181–193.
- 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.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, and Aliaksei Severyn. 2021. [A simple recipe for multilingual grammatical error correction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 702–707, Online. Association for Computational Linguistics.
- Alla Rozovskaya and Dan Roth. 2010. [Annotating ESL errors: Challenges and rewards](#). In *Proceedings of the NAACL HLT 2010 Fifth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 28–36, Los Angeles, California. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. [Minimum risk training for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1683–1692, Berlin, Germany. Association for Computational Linguistics.
- Pavel Soutsov and Sunita Sarawagi. 2016. [Length bias in encoder decoder models and a case for global conditioning](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1516–1525, Austin, Texas. Association for Computational Linguistics.
- Felix Stahlberg and Bill Byrne. 2019. [On NMT search errors and model errors: Cat got your tongue?](#) 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 3356–3362, Hong Kong, China. Association for Computational Linguistics.
- Felix Stahlberg, Adrià de Gispert, Eva Hasler, and Bill Byrne. 2017. [Neural machine translation by minimising the Bayes-risk with respect to syntactic translation lattices](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 362–368, Valencia, Spain. Association for Computational Linguistics.
- Felix Stahlberg and Shankar Kumar. 2021. [Synthetic data generation for grammatical error correction with tagged corruption models](#). In *Proceedings of the 16th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 37–47, Online. Association for Computational Linguistics.
- Joel Tetreault and Martin Chodorow. 2008. [Native judgments of non-native usage: Experiments in preposition error detection](#). In *Coling 2008: Proceedings of the workshop on Human Judgements in Computational Linguistics*, pages 24–32, Manchester, UK. Coling 2008 Organizing Committee.
- 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.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.
- Yingce Xia, Xu Tan, Fei Tian, Fei Gao, Di He, Weicong Chen, Yang Fan, Linyuan Gong, Yichong Leng, Renqian Luo, Yiren Wang, Lijun Wu, Jinhua Zhu, Tao Qin, and Tie-Yan Liu. 2019. [Microsoft Research Asia’s systems for WMT19](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 424–433, Florence, Italy. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [XLNet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

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*.