

What Do You Get When You Cross Beam Search with Nucleus Sampling?

Uri Shaham Omer Levy

The Blavatnik School of Computer Science
Tel Aviv University

Abstract

We combine beam search with the probabilistic pruning technique of nucleus sampling to create two deterministic *nucleus search* algorithms for natural language generation. The first algorithm, *p*-exact search, locally prunes the next-token distribution and performs an exact search over the remaining space. The second algorithm, dynamic beam search, shrinks and expands the beam size according to the entropy of the candidate’s probability distribution. Despite the probabilistic intuition behind nucleus search, experiments on machine translation and summarization benchmarks show that both algorithms reach the same performance levels as standard beam search.

1 Introduction

The standard approach to natural language generation uses a search algorithm, guided by an autoregressive (conditional) language model, to search through the space of possible strings. Since this search space is immense, pruning techniques have been introduced to facilitate tractable text generation. Beam search (Reddy, 1977) is a deterministic algorithm that prunes the search space according to the relative *rank* of each prefix, keeping only the top *b* prefixes at every step. Although rank-based pruning has no probabilistic justification – it is mainly motivated by its ability to limit memory consumption – beam search is an effective approach for generation tasks such as machine translation and summarization. Nucleus sampling (Holtzman et al., 2020), on the other hand, is a stochastic algorithm, which prunes the bottom *percentile* of the model’s next-token distribution, thus eliminating bad candidates while retaining some degree of randomness, which is important for free-form generation. What if we were to replace beam search’s rank-based pruning mechanism (top *k*) with the probabilistic mechanism of nucleus sampling (top *p*)?

We experiment with two variants of this hypothetical *nucleus search*. The first algorithm, *p*-exact search, locally prunes the search space by retaining only the top *p* of every next-token distribution that the underlying language model produces. It then performs an exact search over the remaining space, guaranteeing the most probable sequence under the local pruning assumption. The second algorithm, *dynamic beam search*, selects the top *p* beams at each step, according to their normalized probabilities (rather than top *k*, by rank). This method can shrink or enhance the number of beams to match the current step’s low or high entropy, respectively.

We evaluate both algorithms on three conditional generation benchmarks: subword-level translation (WMT’14 EN-FR), character-level translation (IWSLT’14 DE-EN), and summarization (XSUM). While we observe that both nucleus search algorithms produce competitive results with standard beam search, we do not find any empirical advantage to our probabilistically-motivated approach.

We further analyze the algorithms by isolating the impact of dynamically expanding or shrinking the number of candidates. Experiments show that expanding the beam, even when entropy is high, tends to decrease performance. Pruning candidates, on the other hand, appears to have no adverse effects, and may even have a marginal positive effect in certain cases, which possibly cancels out with the negative effects of beam expansion.

2 Background

Natural language generation can be defined as a search problem in the space of possible sequences over a token vocabulary V , where the goal is to find an optimal sequence $Y = (y_1, \dots, y_n) \in V^*$ according to some cost function. Typical search algorithms explore this infinite space via sequence prefixes, starting with the empty sequence, and

appending one potential token y_t at a time. Search terminates by returning a sequence (or a sequences set) that ends with a special token that indicates the end of the sequence (EOS).

The cost function is based on an underlying language model that, given a prefix $Y_{<t}$, induces a probability distribution over V , which we denote $P(y_t|Y_{<t})$.¹ The probability of a sequence (or prefix) Y is computed as the product of its tokens probabilities:

$$P(Y) = \prod_t P(y_t|Y_{<t}) \quad (1)$$

In practice, it is common to use the negative log probability instead:

$$-\log P(Y) = \sum_t -\log P(y_t|Y_{<t}) \quad (2)$$

This defines a monotonic additive cost function, where appending each token y_t adds a positive cost $-\log P(y_t|Y_{<t})$ to the total cost of the sequence.

2.1 Beam Search

In many natural language generation tasks, *beam search* (Reddy, 1977) is the algorithm of choice. It extends the simple greedy algorithm by considering k possible prefixes $\{Y_{\leq t}^i\}_{i=1}^k$ at each timestep. The beam size k is constant throughout the search, guaranteeing a limit on memory consumption.

At every step t , beam search ranks all the possible single-token extensions of the current k prefixes, and then keeps only the best k extensions according to their total cost (Equation 2). Once a prefix is appended with EOS, it is considered a complete sequence, and remains fixed as long as its cost is among the best k prefixes; if k (or more) better prefixes are found, it is discarded. The algorithm terminates when either the final token of all top k sequences is EOS, or when t exceeds the pre-defined maximum number of steps. In both cases, it returns all sequences in the beam that end with EOS.²

Assuming the models are tuned, results should improve as the beam size k increases. However, this assumption does *not* hold for contemporary

¹The underlying model is often a *conditional* language model $P(y_t|Y_{<t}, X)$, which takes an additional sequence X as part of its input. For brevity, we omit X from our notation.

²Typically, the system selects the top sequence in the set, or chooses an alternative sequence via some reranking criterion.

models; in practice, text quality deteriorates when using large values of k (Koehn and Knowles, 2017). Furthermore, decoding with exact search (Dijkstra, 1959) reveals that translation models often rank the empty string as the most probable sequence (Stahlberg and Byrne, 2019). Perhaps unintentionally, searching with small beam sizes mitigates this flaw.³

2.2 Nucleus Sampling

Deterministic search algorithms, such as beam search, try to generate the most probable sequence. This is a desirable property when we have many constraints regarding the target output, as in translation or question answering. However, tasks that require more creativity and diversity in language may benefit from *stochastic* algorithms.

Holtzman et al. (2020) show that sampling from a language model’s raw distribution P produces degenerate text, and instead, suggest to sample only from the *nucleus*, S_p : the smallest set of tokens whose sum of probabilities is larger than some hyperparameter p . Specifically, nucleus sampling prunes P by assigning zero probability to every token outside of S_p , and renormalizes the probabilities to get a new distribution P_p :

$$P_p(y|Y_{<t}) = \begin{cases} \frac{P(y|Y_{<t})}{\sum_{y' \in S_p} P(y'|Y_{<t})} & y \in S_p \\ 0 & y \notin S_p \end{cases}$$

Here, we refer to this mechanism as *tail pruning*. Sampling from P_p results in less degenerate and more human-like text than both full-distribution sampling and top- k sampling (Fan et al., 2018), which do not account for the distribution’s entropy.

3 Nucleus Search

We combine beam search with tail pruning, producing two variants of *nucleus search*: *p-exact search* and *dynamic beam search*.

3.1 p-Exact Search

Stahlberg and Byrne (2019) show that exact search (Dijkstra, 1959) often produces extremely short and even empty sequences because the underlying model assigns a non-zero probability to the EOS token at each step. We use tail pruning (Section 2.2)

³a.k.a. the “blessing” of beam search (Meister et al., 2020).

to round all near-zero probabilities (whether belonging to EOS or any other token) to zero. We apply exact search over the pruned space, guaranteeing the most probable sequence that contains only top- p tokens at each step.

Given a hyperparameter p , we apply tail pruning to the model’s predicted token distribution $P(y_t|Y_{<t})$. The pruned distribution $P_p(y_t|Y_{<t})$ assigns zero probability to all tokens in the bottom $1 - p$ of the original distribution, and renormalized probabilities for the rest. This procedure prunes the EOS token when it is unlikely, preventing empty sequences and reducing the brevity bias.

3.2 Dynamic Beam Search

Beam search keeps a fixed number (k) of prefixes according to their *rank*. When entropy is high, the difference between the k -th most probable prefix and the one ranked $k + 1$ might be minuscule, and we may want the search algorithm to consider such candidate prefixes as well. Conversely, when entropy is low, the best prefix dominates the alternatives, making them redundant.

Dynamic beam search provides a mechanism for increasing the beam size when entropy is high, and pruning the number of prefixes when entropy is low. Let k_t be the number of viable prefixes at step t . The model predicts the next-token distribution for each prefix, creating $k_t \cdot |V|$ candidates. Each candidate Y^i is scored according to its *cumulative* probability $P(Y^i)$ (Equation 1). To determine the beam size, we first normalize the probability scores within the set of candidates, and then apply tail pruning on the normalized probability:

$$\hat{P}(Y^i) = \frac{P(Y^i)}{\sum_{j=1}^{k_t \cdot |V|} P(Y^j)}$$

As in p -exact search (Section 3.1), we use a hyperparameter p to determine the nucleus of \hat{P} , and thus the size of the next step’s beam k_{t+1} . The normalized probability $\hat{P}(Y^i)$ is only used to compute the dynamic beam; we keep the original probability $P(Y^i)$ as each prefix’s cumulative score.

4 Experiments

We compare our search algorithms to beam search on a variety of tasks, and use the same model across all settings, for each task.

4.1 Tasks

Machine Translation We evaluate on the WMT’14 EN-FR dataset (Bojar et al., 2014), using the model of Ott et al. (2018), a large Transformer (Vaswani et al., 2017) with 6 encoder and decoder layers, trained on 36M bilingual sentences, tokenized with BPE. We evaluate the generated sequences using SacreBLEU (Post, 2018), case-sensitive, with the 13a tokenizer.

Character-Level Machine Translation We train a character-level model on the IWSLT’14 DE-EN dataset (Cettolo et al., 2014), which contains approximately 172k bilingual sentences in its training set. We use the recommended settings in Fairseq (Ott et al., 2019) for a 6-layer encoder-decoder transformer. As with the subword-level dataset, performance is measured via SacreBLEU.

Summarization We evaluate on the XSUM dataset (Narayan et al., 2018). To alleviate memory issues and improve data quality, we remove examples where the source document is longer than 800 tokens (1,663 examples), or when the target is longer than one quarter of the source document (698 examples). Our cleaned version of the XSUM test set contains 8,972 document-summarization pairs. We use the large fine-tuned BART model (Lewis et al., 2020), and compute ROUGE-L (Lin and Hovy, 2003) via compare-mt (Neubig et al., 2019).

4.2 Implementation

Although both nucleus search algorithms can theoretically consume an unbounded amount of memory, our implementation caps the number of candidate prefixes by a large constant: 320 for WMT’14 and XSUM, and 160 for character-level translation.

We explore p in increments of 0.1 for both nucleus search algorithms. For beam search, we experiment with all beam sizes from 1 to 5, as well as exponentially increasing beam sizes from 5 to 320. To present a complete picture of the algorithms’ behaviors, we report results for all hyperparameter settings, rather than selecting the best configuration according to the validation set. This experiment design limits our ability to claim the superiority of one algorithm over another, but as we show in Section 5, the performance differences are so small that no such claim will be made.

Search Algo	Hyper-param (k or p)	WMT'14 EN-FR	IWSLT'14 DE-EN (Char)	XSUM
Beam	1	40.3	33.3	35.5
	2	<u>40.7</u>	<u>33.6</u>	36.2
	3	40.8	<u>33.6</u>	<u>36.4</u>
	4	40.8	<u>33.6</u>	<u>36.5</u>
	5	<u>40.6</u>	<u>33.5</u>	<u>36.5</u>
	10	<u>40.5</u>	<u>33.5</u>	36.6
	20	40.2	33.1	<u>36.4</u>
	40	39.6	27.4	36.1
	80	38.7	18.1	35.7
	160	32.2	5.3	34.3
	320	11.8	5.3	28.1
p -Exact	0.1	40.3	33.3	35.5
	0.2	40.3	33.3	35.7
	0.3	40.5	33.3	36.1
	0.4	40.5	33.4	<u>36.5</u>
	0.5	<u>40.6</u>	<u>33.5</u>	36.6
	0.6	<u>40.6</u>	<u>33.5</u>	36.6
	0.7	40.2	<u>33.6</u>	36.3
	0.8	39.2	<u>33.6</u>	35.9
	0.9	27.8	33.2	33.1
Dynamic	0.1	40.2	33.3	35.5
	0.2	40.3	33.3	35.6
	0.3	40.5	33.4	36.0
	0.4	<u>40.6</u>	33.4	36.2
	0.5	<u>40.6</u>	33.4	<u>36.5</u>
	0.6	<u>40.6</u>	33.7	<u>36.5</u>
	0.7	40.0	33.7	36.0
	0.8	38.9	<u>33.6</u>	35.4
	0.9	18.1	<u>33.1</u>	31.5

Table 1: Scores of different algorithms and settings on various generation tasks. **Bold** numbers indicate the highest result on the task, and underlined numbers indicate that the result is within 0.2 points of the top score.

5 Results

Main Result Table 1 shows the performance of each search algorithm across the different tasks.⁴ In line with previously reported trends (Koehn and Knowles, 2017), we observe that increasing the beam size beyond $k = 10$ can severely degrade performance. On the other hand, the probabilistic search algorithms appear to be more stable, with most hyperparameter settings achieving relatively high performance metrics until $p = 0.9$, where substantial performance degradation is evident.

Despite their increased stability, there appears to be no significant advantage to either p -exact search or dynamic beam search over the original beam search. In fact, the performance differences between the best settings of each algorithm are always under 0.2 BLEU/ROUGE, and often zero.

⁴This table shows performance without reranking (length normalization), to study the core algorithm. Appendix A contains the results with reranking, showing similar trends.

Search Algorithm		$\max(i) \leq 5$	$\max(i) > 5$
Beam	$k = 5$	42.2	32.9
Dynamic Beam	$p = 0.6$	42.3	32.2
#Examples		2618	385

Table 2: Performance on two subsets of WMT'14 EN-FR: (1) examples where dynamic beam search only selects prefixes from the top-5 options ($\max(i) \leq 5$), and (2) examples where the output of dynamic beam search contains at least one prefix that ranked 6 or worse ($\max(i) > 5$).

We find this trend counter-intuitive, since we originally assumed that expanding and trimming the beam based on entropy would benefit language generation. We further test these assumptions individually.

Expanded Beams We compare the performance of static beam search ($k = 5$) and dynamic beam search ($p = 0.6$) on two subsets of the translation task's test set:⁵ (1) examples where dynamic beam search always selects from its top 5 prefixes, and (2) the complement, where every generated output contains at least one prefix that was ranked 6th or worse. Table 2 shows that in those cases where dynamic beam search actually uses the expanded beam, i.e. it chooses prefixes that rank lower than 5, it performs *worse* than static top-5 beam search by 0.7 BLEU. This subset accounts for only 13% of examples – which are probably harder for the model, given the 10-point difference in BLEU – while the majority 87% of cases are always composed from the top 5 (or less) prefixes.

Trimmed Beams We isolate the effect of probabilistic trimming by applying a $k = 5$ cap on the number of active beams, for both nucleus search variations. Table 3 shows that p -exact and dynamic beam trimming strategies have no negative effects, and may have a marginal positive effect.

6 Related Work

As a standard decoding strategy, there is a significant body of literature on beam search. Recently, there has been more focus on the empty string problem (Stahlberg and Byrne, 2019), and the fact that increasing the beam size beyond a small constant typically hurts performance. Meister et al. (2020) show that beam search optimize for sequences that

⁵We select $p = 0.6$ since it is the maximal value that achieved the top score on the WMT'14 EN-FR benchmark.

Search Algo	Hyper-param (k or p)	WMT'14 EN-FR	IWSLT'14 DE-EN (Char)	XSUM
Beam	1	40.3	33.3	35.5
	2	40.7	<u>33.6</u>	36.2
	3	<u>40.8</u>	<u>33.6</u>	<u>36.4</u>
	4	<u>40.8</u>	<u>33.6</u>	<u>36.5</u>
	5	40.6	33.5	<u>36.5</u>
p -Exact ($k = 5$)	0.1	40.3	33.3	35.5
	0.2	40.3	33.3	35.7
	0.3	40.5	33.3	36.1
	0.4	40.6	33.4	<u>36.4</u>
	0.5	<u>40.8</u>	33.5	36.6
	0.6	41.0	<u>33.6</u>	36.6
	0.7	<u>40.9</u>	<u>33.7</u>	36.6
	0.8	<u>40.9</u>	33.8	<u>36.5</u>
	0.9	<u>40.8</u>	33.8	<u>36.5</u>
Dynamic ($k = 5$)	0.1	40.2	33.3	35.5
	0.2	40.3	33.3	35.6
	0.3	40.5	33.4	36.0
	0.4	40.6	33.4	36.2
	0.5	40.6	33.4	<u>36.4</u>
	0.6	<u>40.8</u>	<u>33.7</u>	<u>36.5</u>
	0.7	40.7	<u>33.7</u>	36.6
	0.8	40.7	<u>33.6</u>	36.6
	0.9	40.6	33.5	<u>36.5</u>

Table 3: Scores of different algorithms and settings on various generation tasks, when limiting the beam size to a maximum of 5. **Bold** numbers indicate the highest result on the task, and underlined numbers indicate that the result is within 0.2 points of the top score.

distribute information uniformly, and therefore, using small beam sizes allows it to overcome the empty string problem. Shi et al. (2020) train models with multiple different EOS tokens based on their positions, instead of a single universal EOS token. Peters and Martins (2021) replace the softmax function with the sparse entmax transformation (Peters et al., 2019) that can assign absolute zero probability to tokens. This method has a similar effect to our p -exact search, but requires training the model with entmax, while our contribution only modifies the search algorithm.

Massarelli et al. (2020) also propose a combination of beam search and sampling methods, but with a different method and a different goal. They focus on free-form text generation, addressing two problems – repetition and hallucination – by sampling the first few tokens, and then switching over to beam search. Freitag and Al-Onaizan (2017) explore how using a small fixed beam size, pruned further according to the relative or absolute distance from the top scored candidate, can increase decoding speed. In this work, we focus on the quality of the generated text, comparing the use

of a fixed beam size to tail pruning, an established method that keeps candidates according to the nucleus of the distribution.

7 Conclusion

Language models predict a distribution over their vocabulary, yet beam search only utilizes the rank of different candidates, not their actual probability scores. A natural assumption is that searching the space of prefixes with a constant number of options is not optimal. We hypothesize that using the probability scores to dynamically determine the number of candidates may benefit natural language generation. We test our hypothesis by introducing two nucleus search algorithms, which incorporate probabilistic tail pruning (Holtzman et al., 2020) with beam search, but find that they perform on par with the baseline beam search algorithm when its beam size is restricted to a small constant.

Acknowledgements

This work was supported by the Tel Aviv University Data Science Center, the Blavatnik Fund, the Alon Scholarship, and Intel Corporation. We would like to thank Ari Holtzman, Jonathan Berant, Ori Yoran, Lior Vassertail, and Yuval Kirstain for their valuable feedback.

References

- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Christof Monz, Matt Post, and Lucia Specia, editors. 2014. *Proceedings of the Ninth Workshop on Statistical Machine Translation*. Association for Computational Linguistics, Baltimore, Maryland, USA.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. 2014. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *Proceedings of the International Workshop on Spoken Language Translation, Hanoi, Vietnam*, volume 57.
- Edsger W Dijkstra. 1959. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. *Hierarchical neural story generation*. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.

- Markus Freitag and Yaser Al-Onaizan. 2017. [Beam search strategies for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 56–60, Vancouver. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). *International Conference on Learning Representations*.
- Sébastien Jean, Orhan Firat, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. [Montreal neural machine translation systems for WMT’15](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 134–140, Lisbon, Portugal. Association for Computational Linguistics.
- 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.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [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.
- Chin-Yew Lin and Eduard Hovy. 2003. [Automatic evaluation of summaries using n-gram co-occurrence statistics](#). In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 150–157.
- Luca Massarelli, Fabio Petroni, Aleksandra Piktus, Myle Ott, Tim Rocktäschel, Vassilis Plachouras, Fabrizio Silvestri, and Sebastian Riedel. 2020. [How decoding strategies affect the verifiability of generated text](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 223–235, Online. Association for Computational Linguistics.
- Clara Meister, Ryan Cotterell, and Tim Vieira. 2020. [If beam search is the answer, what was the question?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2173–2185, 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.
- 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.
- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. [compare-mt: A tool for holistic comparison of language generation systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 35–41, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Ben Peters and André F. T. Martins. 2021. [Smoothing and shrinking the sparse Seq2Seq search space](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2642–2654, Online. Association for Computational Linguistics.
- Ben Peters, Vlad Niculae, and André F. T. Martins. 2019. [Sparse sequence-to-sequence models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1504–1519, Florence, Italy. 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.
- D. Raj Reddy. 1977. [Speech understanding systems: A summary of results of the five-year research effort at carnegie-mellon university](#).
- Xing Shi, Yijun Xiao, and Kevin Knight. 2020. [Why neural machine translation prefers empty outputs](#).
- 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.

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 *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

A Results with Reranking

When presenting our main results (Section 5), we follow related work (Peters and Martins, 2021) and focus on the outputs generated using the algorithms themselves, without reranking. For completeness, we also present the results of applying length normalization (Jean et al., 2015; Murray and Chiang, 2018), i.e. reranking the set of sequences produced by beam search according to their average log-probability, rather than their cumulative log-probability:

$$\text{score}(Y) = \frac{1}{n} \sum_{t=1}^n -\log P(y_t|Y_{<t})$$

Table 4 shows that length normalization improves stability, and slightly increases performance overall. However, it does *not* increase the performance gap between the different algorithms, with respect to the results in Section 5 (without reranking); all three variants produce text that scores within 0.2 BLEU/ROUGE from the best performing setting in every task.

Search Algo	Hyper-param (k or p)	WMT'14 EN-FR	IWSLT'14 DE-EN (Char)	XSUM
Beam	1	40.3	33.3	35.5
	2	40.8	33.8	36.3
	3	41.1	<u>34.0</u>	<u>36.4</u>
	4	41.1	<u>34.1</u>	<u>36.5</u>
	5	<u>41.0</u>	<u>34.1</u>	36.6
	10	<u>41.0</u>	34.2	36.6
	20	<u>41.0</u>	34.2	<u>36.5</u>
	40	40.6	34.2	<u>36.4</u>
	80	40.1	34.2	<u>36.3</u>
	160	39.4	34.2	36.2
	320	38.3	34.2	36.2
p-Exact	0.1	40.3	33.3	35.5
	0.2	40.3	33.3	35.6
	0.3	40.5	33.4	36.0
	0.4	40.7	33.4	36.2
	0.5	<u>41.0</u>	33.6	<u>36.4</u>
	0.6	41.1	33.7	36.3
	0.7	<u>41.0</u>	<u>34.0</u>	36.3
	0.8	40.3	<u>34.1</u>	36.2
	0.9	38.8	<u>34.1</u>	36.1
Dynamic	0.1	40.2	33.3	35.5
	0.2	40.3	33.3	35.6
	0.3	40.5	33.4	36.0
	0.4	40.6	33.4	36.2
	0.5	40.8	33.4	<u>36.4</u>
	0.6	<u>41.0</u>	33.8	<u>36.5</u>
	0.7	<u>41.0</u>	<u>34.0</u>	36.3
	0.8	40.6	<u>34.1</u>	36.2
	0.9	38.6	34.2	36.2

Table 4: The performance of different decoding algorithms and hyperparameter settings on various conditional generation tasks with *length normalization* (*reranking*). **Bold** numbers indicate the highest result on the task, and underlined numbers indicate that the result is within 0.2 points of the top score.