

Best-First Beam Search

Clara Meister[★] Ryan Cotterell^{*,★} Tim Vieira[‡]

[★]ETH Zürich ^{*}University of Cambridge [‡]Johns Hopkins University
clara.meister@inf.ethz.ch ryan.cotterell@inf.ethz.ch
tim.f.vieira@gmail.com

Abstract

Decoding for many NLP tasks requires a heuristic algorithm for approximating exact search since the full search space is often intractable if not simply too large to traverse efficiently. The default algorithm for this job is beam search—a pruned version of breadth-first search—which in practice, returns better results than exact inference due to *beneficial* search bias. In this work, we show that standard beam search is a computationally inefficient choice for many decoding tasks; specifically, when the scoring function is a monotonic function in sequence length, other search algorithms can be used to reduce the number of calls to the scoring function (e.g., a neural network), which is often the bottleneck computation. We propose best-first beam search, an algorithm that provably returns the same set of results as standard beam search, albeit in the minimum number of scoring function calls to guarantee optimality (modulo beam size). We show that best-first beam search can be used with length normalization and mutual information decoding, among other rescoring functions. Lastly, we propose a memory-reduced variant of best-first beam search, which has a similar search bias in terms of downstream performance, but runs in a fraction of the time.

1 Introduction

Beam search is a common heuristic algorithm for decoding structured predictors, e.g., neural machine translation models and transition-based parsers. Due to the widespread adoption of recurrent neural networks and other non-Markov models, traditional dynamic programming solutions, such as the Viterbi algorithm (Viterbi, 1967), are prohibitively inefficient; this makes beam search a common component of many state-of-the-art NLP systems. Despite offering no formal guarantee of finding the highest-scoring hypothesis under the

model, beam search yields impressive performance on a variety of tasks—unexpectedly providing a beneficial search bias over *exact* search for many tasks (Stahlberg and Byrne, 2019).

Within NLP, most research on beam search has focused on altering the standard log-probability scoring function to return improved results, e.g., higher BLEU scores (Wu et al., 2016; Murray and Chiang, 2018; Shu and Nakayama, 2018; Yang et al., 2018) or a more diverse set of outputs (Vijayakumar et al., 2016). However, little work has been done to speed up beam search itself. Filling this gap, this paper focuses on reformulating beam search in order to make it *faster*. We propose best-first beam search, a prioritized version of traditional beam search which is up to an order of magnitude faster in practice while still returning the same set of results. We additionally discuss an even faster heuristic version of our algorithm which further limits the number of candidate solutions, leading to a smaller memory footprint while still finding good solutions.

Concretely, we offer a novel interpretation of beam search as an agenda-based algorithm where traditional beam search is recovered by employing a length-based prioritization scheme. We prove that a specific best-first prioritization scheme, as in classic A* search (Hart et al., 1968), allows for the elimination of paths that will necessarily fall off the beam; for many scoring functions, including standard log-probability scoring, we can still guarantee the same k hypotheses as traditional beam search are returned. Indeed, our algorithm returns beam search’s top hypothesis the first time it encounters a complete hypothesis, allowing the program to stop early. Further, we discuss the application of best-first beam search to several popular scoring functions in the literature (He et al., 2016; Li et al., 2016); this demonstrates that we have a general framework for adapting a variety

of rescoring methods and alternate objectives to work with our algorithm.

Empirically, we compare best-first beam search to ordinary beam search on two NLP sequence-to-sequence tasks: neural machine translation (NMT) and abstractive summarization (AS). On NMT, we find that our algorithm achieves roughly a 30% speed-up over traditional beam search with increased gains for larger beams (e.g., $\approx 10x$ for a beam of 500). We find similar results hold for AS. Finally, we show that our memory-reduced version, which limits the number of active hypotheses, leads to additional speed-ups over best-first beam search across beam sizes while maintaining similar BLEU scores.

2 Sequence Transduction

A core operation in structured prediction models is the determination of the highest-scoring output for a given input under a learned scoring model.

$$\mathbf{y}^* \stackrel{\text{def}}{=} \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} \operatorname{score}(\mathbf{x}, \mathbf{y}) \quad (1)$$

where \mathbf{x} is an input and $\mathcal{Y}(\mathbf{x})$ is a set of well-formed outputs for the input. An important example of (1) is maximum a posteriori (MAP),

$$\mathbf{y}^{\text{MAP}} \stackrel{\text{def}}{=} \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} p(\mathbf{y} | \mathbf{x}). \quad (2)$$

Our work focuses on sequence-to-sequence transduction: predicting an output sequence given an input sequence. One such task is machine translation, wherein a source-language sentence is mapped (“transduced”) to a target-language sentence. While our exposition focuses on sequence-to-sequence prediction, our algorithms are directly applicable to any sequential structured prediction model, such as transition-based parsers (Nivre et al., 2008) and sequence taggers (McCallum et al., 2000; Lafferty et al., 2001).

Notation. Let $\mathbf{x} = \langle x_1, \dots, x_{N_x} \rangle$ be an input sequence of length N_x and, likewise, let $\mathbf{y} = \langle y_1, \dots, y_{N_y} \rangle$ be an output sequence of length N_y . Each y_t is an element of \mathcal{V} , the set of output tokens. Finally, let $\mathcal{Y}(\mathbf{x})$ be the set of all valid output sequences (i.e., complete hypotheses). For the task of language generation, which we focus on experimentally, this set is defined as

$$\mathcal{Y}(\mathbf{x}) \stackrel{\text{def}}{=} \{\text{BOS} \circ \mathbf{v} \circ \text{EOS} \mid \mathbf{v} \in \mathcal{V}^{<n_{\max}}\} \quad (3)$$

where \circ is string concatenation and $\mathcal{V}^{<n_{\max}(\mathbf{x})}$ is the set of all subsets of \mathcal{V}^* of size $< n_{\max}(\mathbf{x})$. In words, every valid sequence begins and ends with distinguished tokens (BOS and EOS, respectively).¹ Furthermore, each sequence has *at most* length $n_{\max}(\mathbf{x})$ —which is typically dependent on \mathbf{x} —a restriction we impose to ensure termination. Some applications may require a stronger coupling between $\mathcal{Y}(\mathbf{x})$ and \mathbf{x} (e.g., $|\mathbf{x}| = |\mathbf{y}|$). We drop the dependence of \mathcal{Y} and n_{\max} on \mathbf{x} when it is clear from context.

Scoring. We consider a general additively decomposable scoring model of the form

$$\operatorname{score}(\mathbf{x}, \mathbf{y}) = \sum_{t=1}^{N_y} \operatorname{score}(\mathbf{x}, \mathbf{y}_{<t} \circ y_t) \quad (4)$$

This framework covers a variety of modeling methodologies including probabilistic transducers (both globally and locally normalized) and non-probabilistic models such as maximum-margin techniques (Taskar et al., 2004). Most importantly, (4) covers MAP decoding (2) of neural sequence-to-sequence models à la Sutskever et al. (2014):²

$$\operatorname{score}_{s2s}(\mathbf{x}, \mathbf{y}_{<t} \circ y_t) = \log p(y_t | \mathbf{y}_{<t}, \mathbf{x}) \quad (5)$$

We note that (5) is the scoring function used for decoding many language generation models.

Beam search. The worst-case running time of exactly computing (1) is exponential in n_{\max} ; namely, $\mathcal{O}(|\mathcal{V}|^{n_{\max}})$.³ Beam search is a commonly used approximation to (1) in NMT and language generation tasks. It is used in many (if not most) state-of-the-art NLP systems (Wu et al., 2016; Serban et al., 2017; Edunov et al., 2018; Yang et al., 2019). Beam search may be understood as a pruned version of the classic path-search algorithm, breadth-first search (BFS), where the breadth is narrowed to the beam size k . Pseudocode is given in Alg. 1.

Although, beam search does not solve (1) exactly, it is a surprisingly useful approximation for NLP models. In many settings, beam search *outper-*

¹BOS and EOS are typically members of \mathcal{V} . Often, EOS counts towards the n_{\max} length limit while BOS does not. This is reflected in (3).

²To see why, apply \exp (an order-preserving transformation): $\exp(\operatorname{score}_{s2s}(\mathbf{x}, \mathbf{y})) = \exp\left(\sum_{t=1}^{N_y} \log p(y_t | \mathbf{y}_{<t}, \mathbf{x})\right) = \prod_{t=1}^{N_y} p(y_t | \mathbf{y}_{<t}, \mathbf{x}) = p(\mathbf{y} | \mathbf{x})$

³This can be improved if, for example, $\operatorname{score}(\cdot, \cdot)$ admits a low-order Markov factorization (Viterbi, 1967; Vieira et al., 2016). We do not discuss that setting in this paper because it limits the scoring model’s expressive power.

Algorithm 1 Standard beam search⁴

Input: \mathbf{x} : source sentence
 k : maximum beam size
 n_{max} : maximum hypothesis length
 $\text{score}(\cdot, \cdot)$: scoring function

- 1: $B_0 \leftarrow \{\langle 0, \text{BOS} \rangle\}$
- 2: **for** $t \in \{1, \dots, n_{max} - 1\}$:
- 3: $B \leftarrow \emptyset$
- 4: **for** $\langle s, \mathbf{y} \rangle \in B_{i-1}$:
- 5: **if** $\mathbf{y}.\text{last}() = \text{EOS}$:
- 6: $B.\text{add}(\langle s, \mathbf{y} \rangle)$
- 7: **continue**
- 8: **for** $y \in \mathcal{V}$:
- 9: $s \leftarrow \text{score}(\mathbf{x}, \mathbf{y} \circ y)$
- 10: $B.\text{add}(\langle s, \mathbf{y} \circ y \rangle)$
- 11: $B_i \leftarrow B.\text{top}(k)$
- 12: **return** $B.\text{max}()$

forms exact methods in terms of downstream evaluation (Koehn and Knowles, 2017; Stahlberg and Byrne, 2019). For the remainder of this paper, we will pivot our attention away from exact solutions to (1) to exact solutions to the *beam search* output.

Definition 2.1. *k-optimal hypothesis.* We say that a hypothesis is *k-optimal* if it is the top hypothesis returned by beam search with beam size k .

3 A* Beam Search

We develop a meta-algorithm that is parameterized by several choice points. Our general search algorithm for decoding (Alg. 2) takes an arbitrary prioritization function, stopping criterion, and search heuristic. With certain values of these attributes, we recover many common search algorithms: greedy search, beam search, best-first search (Dijkstra, 1959), and A* search (Hart et al., 1968). We propose an alternate prioritization function for beam search that allows for faster decoding while still returning the same *k-optimal* set of hypotheses.

⁴Often, the score function is additively decomposable in t , such as (5). Implementations can exploit this fact to make each score evaluation (line 9) $\mathcal{O}(1)$ rather than $\mathcal{O}(t)$. We did not make this implementation detail explicit in Alg. 1 or Alg. 2 for generality and simplicity.

⁵If the last token of \mathbf{y}' is the end symbol (e.g., EOS), then \mathbf{y}' is not expanded any further. One can either regard \mathbf{y}' as any other hypothesis albeit with $\mathbf{y}' \circ y_t = \mathbf{y}'$ or keep appending EOS (i.e. $\mathbf{y}' \circ y_t = \mathbf{y}' \circ \text{EOS}$) so that time step and length can be regarded as synonymous. We adopt the latter standard for comparability with subsequent algorithms.

Algorithm 2 A* beam search.^{4,5} Highlighted sections are choice points in the algorithm for which values determine the search strategy. See § 3.1 for detailed explanation.

Input: \mathbf{x} : source sentence
 n_{max} : maximum hypothesis length
 $\text{score}(\cdot, \cdot)$: scoring function
 \ominus : comparator ①
 $\text{stop}(\cdot, \cdot)$: stopping criterion ②
 k : maximum beam size ③
 $h(\cdot, \cdot)$: heuristic function ④

- 1: $Q \leftarrow \text{priority_queue}(\ominus)$
- 2: $Q.\text{push}(\langle 0, \text{BOS} \rangle)$
- 3: $\text{POPS} \leftarrow \text{counter}()$
- 4: **while not** $\text{stop}(Q)$ **and not** $Q.\text{empty}()$:
- 5: $\langle s_h, \mathbf{y} \rangle \leftarrow Q.\text{pop}()$
- 6: **if** $\text{POPS}[|\mathbf{y}|] \geq k$ **or** $|\mathbf{y}| > n_{max}$:
- 7: **continue**
- 8: $\text{POPS}[|\mathbf{y}|] \leftarrow \text{POPS}[|\mathbf{y}|] + 1$
- 9: **if** $\mathbf{y}.\text{last}() = \text{EOS}$:
- 10: $Q.\text{push}(\langle s_h, \mathbf{y} \circ \text{EOS} \rangle)$
- 11: **else:**
- 12: **for** $y \in \mathcal{V}$:
- 13: $s \leftarrow \text{score}(\mathbf{x}, \mathbf{y} \circ y)$
- 14: $s_h \leftarrow s + h(\mathbf{x}, \mathbf{y} \circ y)$
- 15: $Q.\text{push}(\langle s_h, \mathbf{y} \circ y \rangle)$
- 16: **return** $Q.\text{pop}()$ **if not** $Q.\text{empty}()$ **else null**

3.1 Choice Points of the Meta Algorithm

Here we review the components of our meta algorithm (the highlighted sections in Alg. 2) that can be varied to recover different search strategies:

- ① A priority queue Q maintains the set of active hypotheses. Elements in this set are ordered according to a generic comparator \ominus . When its $\text{peek}()$ (or $\text{pop}()$) methods are called, the first element ordered by \ominus is returned (or returned and removed).
- ② The algorithm terminates according to configurable stopping criterion based on elements in Q .
- ③ Only k paths of a given length are considered. If the algorithm has already encountered k paths of a given length, subsequent paths of that length are not evaluated. If we take $k = \infty$, we recover unpruned search algorithms.
- ④ A heuristic function $h(\mathbf{x}, \mathbf{y})$ can be used

during search to change the priority in which paths are evaluated. We note that with pruning, a heuristic may change the value of the k -optimal hypothesis (see § 4.1).

Recovering Beam Search. To recover beam search from Alg. 2, we use the choice points from Tab. 1. Explicitly, the comparator prioritizes hypotheses from earlier time steps first, but breaks ties with the hypotheses’ scores under the model. We note that while the standard algorithm for beam search does not prioritize by score within a time step, variations of the algorithm use this strategy so they can employ early-stopping strategies (Klein et al., 2017; Huang et al., 2017). Beam search terminates once either *all* hypotheses end in EOS or the queue is empty (i.e., when the k beams have been extended n_{max} time steps but none end in EOS). In the second case, no complete hypothesis is found. Finally, choosing the heuristic $h(\mathbf{x}, \mathbf{y}) = 0$ makes the algorithm a case of standard best-first search.

Note that, while standard beam search returns a set, Alg. 2 only returns the k -optimal hypothesis. This behavior is sufficient for the majority of use cases for beam search. However, if the full set of k hypotheses is desired, the stopping criterion can be changed to evaluate true only when k hypotheses are complete. Under the other beam search settings, this would guaranteeably return the same set as beam search (see § 4.1).

Recovering A*. To recover the traditional A* search algorithm we use the comparator that prioritizes hypotheses with a higher score first; ties are broken by hypothesis length. The algorithm terminates when the first item of Q contains an EOS. If we take $k = \infty$, best-first beam search recovers A*. Any admissible heuristic may be used for $h(\mathbf{x}, \mathbf{y})$.

Definition 3.1. Admissible Heuristic. A heuristic h is admissible if it never overestimates the future cost—or underestimates the future reward—of continuing down a path.

3.2 Best-First Beam Search

In its original form, A* search may traverse the entire $\mathcal{O}(|\mathcal{V}|^{n_{max}})$ graph, which as discussed earlier, is intractable for many decoding problems. While standard beam search addresses this problem by limiting the search space, it still has computational inefficiencies—namely, we *must* analyze k hypotheses of a given length (i.e., time step), regardless of how poor their scores may already be,

before considering longer hypotheses. However, prioritization by length is not strictly necessary for finding a k -optimal hypothesis. As is done in A*, we can use score as the prioritization scheme and still guarantee optimality—or k -optimality—of the paths returned by the algorithm.

We define A* beam search as the A* algorithm where breadth is limited to size k . Further, we define best-first beam search as the case of A* beam search when no heuristic is used (see Tab. 1 for algorithm settings). This formulation has two large advantages over standard beam search: (1) we gain the ability to remove paths from the queue that are guaranteed to fall off the beam and (2) we can terminate the algorithm the first time a complete hypothesis is encountered. We can therefore reduce the computation required for decoding while still returning the same set of results.

To see why the above is true, note that the standard log-probability scoring function used by many sequential structured prediction models is monotonically decreasing in t .

Definition 3.2. Monotonicity. A scoring function $\text{score}(\cdot, \cdot)$ is monotonic in t if for all \mathbf{x} , $\mathbf{y}_{<t} = \langle y_1 \dots y_{t-1} \rangle$, $y_t \in \mathcal{V}$, $1 \leq t \leq n_{max}$

$$\text{score}(\mathbf{x}, \mathbf{y}_{<t}) \geq \text{score}(\mathbf{x}, \mathbf{y}_{<t} \circ y_t)$$

Clearly, (5) is a monotonic scoring function in t because $\text{score}_{s2s} \leq 0$, i.e., score can only decrease in t . This implies we can order our search according to $\text{score}(\mathbf{x}, \mathbf{y}_{<t})$ without fear of overlooking a hypothesis whose score would increase over time. Furthermore, once k hypotheses of a given length t have been evaluated, we no longer need to consider any hypothesis where $|\mathbf{y}| < t$ since such hypotheses would necessarily fall off the beam. We can therefore remove such hypotheses from the queue and avoid wasting computational power on their evaluation. We prove this formally in § 4.1.

Another implication of the monotonicity property of score is that we may terminate best-first beam search once a hypothesis containing EOS is encountered (i.e., the end state is found). If the full set of k complete hypotheses is desired, then we simply continue until k hypotheses have reached EOS. We prove the k -optimality of these hypotheses under best-first beam search in § 4.1.

3.3 Implementation Details

Standard beam search forms a separate set of active hypotheses for each time step, i.e., each B_t is its

	Beam Search	Best-First Beam Search	A* Beam Search
①	$\langle s_h, \mathbf{y} \rangle \otimes \langle s'_h, \mathbf{y}' \rangle \iff \mathbf{y} < \mathbf{y}' $ or $(\mathbf{y} = \mathbf{y}' \text{ and } s_h \geq s'_h)$	$\langle s_h, \mathbf{y} \rangle \otimes \langle s'_h, \mathbf{y}' \rangle \iff s_h > s'_h$ or $(s_h = s'_h \text{ and } \mathbf{y} < \mathbf{y}')$	$\langle s_h, \mathbf{y} \rangle \otimes \langle s'_h, \mathbf{y}' \rangle \iff s_h > s'_h$ or $(s_h = s'_h \text{ and } \mathbf{y} < \mathbf{y}')$
②	stop(Q) \iff $\mathbf{y}.\text{last}() = \text{EOS} \quad \forall \mathbf{y} \in Q$	stop(Q) \iff $Q.\text{peek}().\text{last}() = \text{EOS}$	stop(Q) \iff $Q.\text{peek}().\text{last}() = \text{EOS}$
③	$k = \text{beam size}$	$k = \text{beam size}$	$k = \text{beam size}$
④	0	0	any admissible heuristic

	Breadth-First Search	Best-First Search	A* Search
①	$\langle s_h, \mathbf{y} \rangle \otimes \langle s'_h, \mathbf{y}' \rangle \iff \mathbf{y} < \mathbf{y}' $ or $(\mathbf{y} = \mathbf{y}' \text{ and } s_h \geq s'_h)$	$\langle s_h, \mathbf{y} \rangle \otimes \langle s'_h, \mathbf{y}' \rangle \iff s_h > s'_h$ or $(s_h = s'_h \text{ and } \mathbf{y} < \mathbf{y}')$	$\langle s_h, \mathbf{y} \rangle \otimes \langle s'_h, \mathbf{y}' \rangle \iff s_h > s'_h$ or $(s_h = s'_h \text{ and } \mathbf{y} < \mathbf{y}')$
②	stop(Q) \iff $\mathbf{y}.\text{last}() = \text{EOS} \quad \forall \mathbf{y} \in Q$	stop(Q) \iff $Q.\text{peek}().\text{last}() = \text{EOS}$	stop(Q) \iff $Q.\text{peek}().\text{last}() = \text{EOS}$
③	$k = \infty$	$k = \infty$	$k = \infty$
④	0	0	any admissible heuristic

Table 1: Values at choice points for various search algorithms. Note that any admissible heuristic may be used for variants of A* search.

own set. Once B_t has been narrowed down to the top k , the previous $B_{<t}$ can be forgotten. However in best-first beam search, since hypotheses are not evaluated in order of time step, we may need to keep B_t from several time steps at any given point.

A naive implementation of best-first beam search is to keep a single priority queue with all the active hypotheses ordered by current score. However, each push to the queue would then require $\mathcal{O}(\log(n_{max}k|\mathcal{V}|))$ time. We can reduce this runtime by instead keeping a priority queue of beams, where the priority queue is ordered by the highest-scoring hypothesis from each beam. Further, each beam can be represented by a min-max queue (Atkinson et al., 1986); this allows us to limit the size of B_t to k : we can check in $\mathcal{O}(1)$ time if a hypothesis is in the top- k before adding it to B_t .

A potential inefficiency, which we avoid, comes from updating B_{t+1} , which we must do when evaluating a hypothesis from B_t . Since all beams are stored in a queue, there is no guarantee of the location in the queue of B_{t+1} . To avoid $\mathcal{O}(n_{max})$ lookup, we can keep a pointer to each beam, indexed by t making the lookup $\mathcal{O}(1)$. However, we acquire a $\mathcal{O}(\log n_{max})$ term to update the queue of beams as B_{t+1} may change priority.

Memory-Reduced best-first beam search. A major drawback of the A* algorithm is its memory usage, which in the worst-case is $\mathcal{O}(b^d)$ for breadth width b and maximum depth d . In the A* formulation of beam search, where the breadth width is limited to the beam size, this amounts to worst-case $\mathcal{O}(k \cdot n_{max})$ memory usage, where standard beam search has $\mathcal{O}(k)$ memory usage. While in many set-

tings the multiplicative factor may be insignificant, for neural sequence models it can be prohibitive; this is due to the large amount of memory required to store each hypothesis (e.g., prior hidden states needed to compute subsequent scores for scoring functions parameterized by neural networks).

We propose a variant of best-first beam search that limits memory usage, i.e., the queue capacity. Specifically, if we reach the chosen queue capacity, we remove the worst scoring active hypothesis from the earliest active time step. This can easily be done in $\mathcal{O}(1)$ time given our pointer to each beam.

4 Algorithm Analysis

4.1 Correctness

We show the equivalence of the top hypothesis⁶ returned by beam search and best-first beam search when $\text{score}(\cdot, \cdot)$ is monotonically decreasing in t , length-based prioritization is used and the beam size k is the same for both algorithms. Without loss of generality, we hold \mathbf{x} constant in all the following proofs.

Note that we take the terms *pop* and *push* from queue terminology. Specifically, “popping a hypothesis” refers to making it past line 7 of Alg. 2, where a hypothesis \mathbf{y} is expanded by $y_t \in \mathcal{V}$. In path search terminology, this would be equivalent to visiting a node and adding the edges from that node as potential paths to explore. Lastly, we refer

⁶best-first beam search is guaranteed to return the same set of k hypotheses as beam search. We include the proof for only the top hypothesis for simplicity. The proof for set equality follows naturally.

to the priority queue used by beam search and best-first beam search as \mathcal{Q}_{BS} and \mathcal{Q}_{A^*} , respectively.

Lemma 4.1. *best-first beam search evaluates all hypotheses of a given length t in order of their score.*

Proof. We prove the lemma by induction. The lemma holds trivially for the base case of hypotheses of length 0 because the only hypothesis of length 0 is $\langle BOS \rangle$.

Now, by the inductive hypothesis, suppose Lemma 4.1 holds for all hypotheses of length $< t$. We will show it must also hold for hypotheses of length t . Consider two competing hypotheses: $\mathbf{y} = \mathbf{y}_{<t} \circ y_t$ and $\mathbf{y}' = \mathbf{y}'_{<t} \circ y'_t$. Note that $|\mathbf{y}_{<t}| = |\mathbf{y}'_{<t}| = t - 1$. Suppose $\text{score}(\mathbf{x}, \mathbf{y}') < \text{score}(\mathbf{x}, \mathbf{y})$.

Case 1: $\text{score}(\mathbf{x}, \mathbf{y}'_{<t}) < \text{score}(\mathbf{x}, \mathbf{y}_{<t})$. Then by induction, $\mathbf{y}_{<t}$ popped first and \mathbf{y} is pushed to \mathcal{Q} before \mathbf{y}' . Since $\text{score}(\mathbf{x}, \mathbf{y}') < \text{score}(\mathbf{x}, \mathbf{y})$, \mathbf{y}' will be popped before \mathbf{y} .

Case 2: $\text{score}(\mathbf{x}, \mathbf{y}_{<t}) < \text{score}(\mathbf{x}, \mathbf{y}'_{<t})$. Then by induction, $\mathbf{y}'_{<t}$ is popped first and \mathbf{y}' is added to \mathcal{Q} before \mathbf{y} . But, since $\text{score}(\mathbf{x}, \mathbf{y}') < \text{score}(\mathbf{x}, \mathbf{y}) \leq \text{score}(\mathbf{x}, \mathbf{y}_{<t})$ by monotonicity, then $\mathbf{y}_{<t}$ will be popped before \mathbf{y}' . Consequently, \mathbf{y} will be pushed to \mathcal{Q} before \mathbf{y}' is evaluated. By the rules of the priority queue \mathbf{y} will be evaluated before \mathbf{y}' .

Case 3: $\text{score}(\mathbf{x}, \mathbf{y}') = \text{score}(\mathbf{x}, \mathbf{y})$. The lemma holds if either \mathbf{y} or \mathbf{y}' is popped first.

By the principle of induction, Lemma 4.1 holds for all $t \in \mathbb{Z}_+$. \square

Lemma 4.2. *The first hypothesis that best-first beam search pops that ends in EOS is the best hypothesis found by best-first beam search.*

Proof. Let \mathbf{y} be the first hypothesis popped by best-first beam search ending in EOS. By rules of the priority queue, no other active hypothesis has a higher score than \mathbf{y} . Additionally, by monotonicity of the scoring function, no other hypothesis can subsequently have score greater than \mathbf{y} . Therefore \mathbf{y} must be the best hypothesis found by best-first beam search. \square

Lemma 4.3. *If best-first beam search pops a hypothesis, then beam search necessarily pops that same hypothesis.*

Proof. We prove the lemma by induction on hypothesis length. The base case holds trivially: For hypotheses of length 0, both best-first beam search

and beam search must pop the $\langle BOS \rangle$ as it is the only item in the queue after initialization.

By the inductive hypothesis, suppose Lemma 4.3 holds for hypotheses of length $< t$. Suppose best-first beam search pops a hypothesis $\mathbf{y} = \mathbf{y}_{<t} \circ y_t$ of length t .

Case 1: best-first beam search pops k hypotheses of length $t - 1$ before popping \mathbf{y} , which is of length t . The sets of hypotheses of length $t - 1$ that each algorithm pops are necessarily the same by the inductive hypothesis and the fact that they have the same cardinality. If best-first beam search pops \mathbf{y} , which is of length t , then it must be in the top- k highest-scoring hypotheses of length t in \mathcal{Q}_{A^*} by the rules of the priority queue. Consequently, it must be in the top- k in \mathcal{Q}_{BS} .

Case 2: best-first beam search has popped fewer than k hypotheses of length $t - 1$ before popping \mathbf{y} . Then, all remaining hypotheses of length $t - 1$ in \mathcal{Q}_{A^*} must have $\text{score}(\mathbf{x}, \mathbf{y}'_{<t}) < \text{score}(\mathbf{x}, \mathbf{y})$ by the rules of the priority queue. By the monotonicity of the score function, all extensions of those $\mathbf{y}'_{<t}$ will also have $\text{score}(\mathbf{x}, \mathbf{y}'_{<t} \circ y'_t) < \text{score}(\mathbf{x}, \mathbf{y})$. Because none of $\mathbf{y}'_{<t} \circ y'_t$ has greater score than \mathbf{y} , \mathbf{y} must be in B_t . \square

Corollary 4.3.1. *best-first beam search will never pop more hypotheses than beam search.*

Theorem 4.4. *Once best-first beam search has popped k hypotheses of length t , hypotheses from time steps $< t$ do not need to be popped.*

Proof. This follows from Lemma 4.1. If k hypotheses of length t have been popped, then these must be the top- k hypotheses from time step t . Therefore no hypothesis from time step $< t$ that is still in \mathcal{Q}_{A^*} would be in the top- k at time step t . \square

Theorem 4.5. *Upon stopping, beam search and best-first beam search return the same hypothesis.*

Proof. Let $\tilde{\mathbf{y}}_{A^*}$ be the hypothesis best-first beam search returns. By Lemma 4.3 beam search also pops $\tilde{\mathbf{y}}_{A^*}$. Because $\tilde{\mathbf{y}}_{A^*}$ ends in EOS, we have that $\tilde{\mathbf{y}}_{A^*} \circ \text{EOS}$ will also have the same score; it follows by Lemma 4.1 that $\tilde{\mathbf{y}}_{A^*} \circ \text{EOS}$ is the highest-scoring hypothesis of length $|\tilde{\mathbf{y}}_{A^*}| + 1$ and, by monotonicity, of all subsequent lengths. Thus, when beam search reaches n_{max} , $\tilde{\mathbf{y}}_{A^*}$ remains in \mathcal{Q}_{BS} . Therefore, when beam search terminates, $\tilde{\mathbf{y}}_{A^*}$ will be the highest-scoring hypothesis that ends in EOS and thus the solution returned by beam search. \square

Non-monotonic Scoring Functions. Non-monotonic scoring functions (definition 3.2) break the assumptions of § 4.1, in which case best-first beam search is not guaranteed to return a k -optimal hypothesis. However, when the scoring function is boundable from above, we can alter the original stopping criterion (② in Alg. 2) such that k -optimality is again guaranteed.

Given our assumed restriction on the search space—namely, $|\mathbf{y}^* \in \mathcal{Y}(\mathbf{x})| \leq n_{\max}(\mathbf{x})$ —we can upper-bound the maximal score of any hypothesis under the scoring function in use. Formally, for any function score we have:

$$\begin{aligned} \text{stop}(\mathcal{Q}) \iff \\ \text{score}(\mathbf{x}, \hat{\mathbf{y}}) \geq \text{score}(\mathbf{x}, \mathbf{y}') + \mathcal{U}(\mathbf{x}, \mathbf{y}') \\ \forall \mathbf{y}' \in \mathcal{Q} \end{aligned} \quad (6)$$

where $\hat{\mathbf{y}}$ is the best complete hypothesis found so far and $\mathcal{U}(\mathbf{x}, \mathbf{y}')$ is the score function-dependent upper bound on how much the score of \mathbf{y}' can increase as \mathbf{y}' is expanded further.⁷ In this situation, best-first beam search only terminates once no other hypothesis in \mathcal{Q} can have a score greater than the best finished hypothesis. We note that Huang et al. (2017) use a similar scheme for optimal stopping with bounded length normalization. We discuss examples of non-monotonic scoring functions in § 5.

A Note on Heuristics. Our analysis shows the equivalence of beam search and best-first beam search, i.e., when $h(\mathbf{x}, \mathbf{y}) = 0$. The analysis does *not* hold for arbitrary admissible heuristics. A poor heuristic, e.g., one that grossly overestimates the future score of continuing down one path, may cause other items to be pruned from best-first beam search that otherwise would have remained on the beam in standard beam search.

4.2 Runtime

Theorem 4.6. *The runtime of best-first beam search is $\mathcal{O}(n_{\max}k(|\mathcal{V}|\log(k) + \log(n_{\max})))$*

Proof. We pop at most $n_{\max} \cdot k$ items. Each pop requires us to push $|\mathcal{V}|$ items. Each push requires $\log(k)$ time when the priority queue is implemented with a min-max heap (Atkinson et al., 1986) and incrementally pruned so that it has no more than k items. After pushing those $|\mathcal{V}|$ items, we have to perform a percolation in the

⁷For monotonic scoring functions, we have $\mathcal{U}(\mathbf{x}, \mathbf{y}') = 0$.

priority queue of priority queues which requires $\log(n_{\max})$ time. This yields $\mathcal{O}(n_{\max}k(|\mathcal{V}|\log(k) + \log(n_{\max})))$ time. \square

Theorem 4.7. *The runtime of standard beam search is $\mathcal{O}(n_{\max}k|\mathcal{V}|\log(k))$.*

Proof. The proof is the same as Theorem 4.6, but we can forgo the percolation step in the queue of queues because standard beam search proceeds in order of hypothesis length. This yields $\mathcal{O}(n_{\max}k|\mathcal{V}|\log(k))$. \square

While the theoretical bound of best-first beam search has an additional log factor compared to standard beam search, we find this to be negligible in practice. Rather, we find number of calls to score, the scoring function under our model (e.g., a neural network), is often the bottleneck operation when decoding neural networks (see § 6 for empirical evidence). In terms of this metric, the beam search algorithm makes $\mathcal{O}(kn_{\max})$ calls to score, as score is called once for each active hypothesis in B and B may evolve for n_{\max} rounds. The worst-case number of calls to score will be the same as for beam search, which follows from Lemma 4.3.

5 Scoring Functions

Even before the findings of Stahlberg and Byrne (2019), it was well known that the best scoring hypothesis with respect to the traditional likelihood objective can be far from ideal in practice (Wu et al., 2016; Murray and Chiang, 2018; Yang et al., 2018). For language generation tasks specifically, the results returned by neural models using the standard scoring function are often short and default to high frequency words (Vinyals and Le, 2015; Shen et al., 2016).

To alleviate such problems, methods that revise hypothesis scores to incorporate preferences for longer, less repetitive, or more diverse options have been introduced and are often used in practice. While most such techniques change the scoring function such that it is no longer monotonic, we can still guarantee the k -optimality of the returned hypothesis for (upper) bounded scoring functions using the methods discussed in § 4.1. In the remainder of this section, we present alternate scoring schemes adapted to work with best-first beam search. Additionally, we present several heuristics which, while breaking the k -optimality guarantee, provide another set of decoding strategies worth exploring.

Length Normalization. Length normalization is a widely-used hypothesis scoring method that aims to counteract the propensity for shorter sequences to have higher scores under neural models; this is done by normalizing scores by hypothesis length (see Murray and Chiang (2018) for more detail).

For early stopping in beam search with length normalization, Huang et al. (2017) propose bounding the additive length reward as the minimum of a pre-determined optimal sequence length ratio r and the final sequence length $N_{\mathbf{y}}$:

$$\text{score}_{\text{LN}}(\mathbf{x}, \mathbf{y}) = \text{score}(\mathbf{x}, \mathbf{y}) + \beta \cdot \min\{r|\mathbf{x}|, N_{\mathbf{y}}\} \quad (7)$$

where β is the scaling parameter for the reward. We note, however, that the same can be done with the maximum sequence length n_{max} such that the traditional length reward used by He et al. (2016) is recovered:

$$\begin{aligned} \text{score}_{\text{LN}}(\mathbf{x}, \mathbf{y}) &= \text{score}(\mathbf{x}, \mathbf{y}) + \beta \min\{n_{\text{max}}, N_{\mathbf{y}}\} \\ &= \text{score}(\mathbf{x}, \mathbf{y}) + \beta N_{\mathbf{y}} \end{aligned} \quad (8)$$

We formally propose two methods for length normalization. We use the scoring functions in (7) or (8) with either: (1) the following heuristic:

$$h(\mathbf{x}, \mathbf{y}) = \begin{cases} 0 & \text{for } \mathbf{y}.\text{last}() = \text{EOS} \\ \beta \max\{b - |\mathbf{y}|, 0\} & \text{for } \mathbf{y}.\text{last}() \neq \text{EOS} \end{cases} \quad (9)$$

where b can be $r|\mathbf{x}|$ or n_{max} ,⁸ or (2) stopping criterion as in (6) albeit with score_{LN} and upper-bound function:

$$\mathcal{U}(\mathbf{x}, \mathbf{y}) = \beta \max\{0, b - |\mathbf{y}|\} \quad (10)$$

Despite their similarities, these two methods are not guaranteed to return the same results. While the second method will return the same k -optimal hypotheses as beam search, using a heuristic during pruned search means we can no longer guarantee the k -optimality of the results with respect to the scoring function as the heuristic may push hypotheses off of the beam. We present experimental results for both methods in § 6.

Mutual Information. Maximum mutual information decoding (Li et al., 2016) aims to alleviate the inherent preference of neural models for high-frequency tokens when using the log-probability decoding objective. Rather than choosing the hypothesis \mathbf{y} to maximize conditional probability with

⁸We enforce $r|\mathbf{x}| < n_{\text{max}}$.

respect to the input \mathbf{x} , we instead choose \mathbf{y} to maximize pointwise mutual information (PMI):

$$\text{PMI}(\mathbf{x}; \mathbf{y}) = \log \frac{p(\mathbf{x}, \mathbf{y})}{p(\mathbf{x})p(\mathbf{y})} \quad (11)$$

Note that (11) is equivalent to $\log \frac{p(\mathbf{y}|\mathbf{x})}{p(\mathbf{y})}$, which can be rewritten as $\log p(\mathbf{y} | \mathbf{x}) - \log p(\mathbf{y})$ making the objective additive and thus (11) can conform to (4).

From this last form, we can see how mutual information decoding penalizes high-frequency and generic outputs; the negative $p(\mathbf{y})$ term, as Li et al. (2016) point out, acts as an ‘‘anti-language model.’’ One unfortunate side effect of this objective is that ungrammatical and nonsensical outputs, which have probabilities close to 0 under a language model like $p(\mathbf{y})$, end up with high scores due to the second term in the score function. To address this problem, and to upper-bound the scoring function, we propose lower-bounding the language model term by a hyperparameter $1 \geq \varepsilon > 0$. We additionally use the strength hyperparameter λ employed by Li et al. (2016):

$$\begin{aligned} \text{score}_{\text{PMI}}(\mathbf{x}, \mathbf{y}) &= \log p(\mathbf{y} | \mathbf{x}) \\ &\quad - \lambda \log \max\{p(\mathbf{y}), \varepsilon\} \end{aligned} \quad (12)$$

Similarly to our methods for length normalization, we can use the scoring function in (12) either with the heuristic:

$$h(\mathbf{x}, \mathbf{y}) = \begin{cases} 0 & \text{for } \mathbf{y}.\text{last}() = \text{EOS} \\ -\lambda \log \varepsilon(n_{\text{max}} - |\mathbf{y}|) & \text{for } \mathbf{y}.\text{last}() \neq \text{EOS} \end{cases} \quad (13)$$

or with stopping criterion as in (6) albeit with $\text{score}_{\text{PMI}}$ and upper-bound function:

$$\mathcal{U}(\mathbf{x}, \mathbf{y}) = -\lambda \log \varepsilon(n_{\text{max}} - |\mathbf{y}|) \quad (14)$$

Since $-\lambda \log \varepsilon$ is the best possible score at any given time step, clearly we can bound the increase in $\text{score}_{\text{PMI}}$ by the above function. However, as with our length normalization strategy, we lose the k -optimality guarantee with the heuristic method for mutual information decoding. We present experimental results for both methods in § 6.

6 Experiments

We run our algorithm on several language-related tasks that typically use beam search for decoding: neural machine translation (NMT) and abstractive summarization (AS). Specifically, experiments are performed on IWSLT’14 De-En (Cettolo et al., 2012), WMT’17 De-En (Bojar et al., 2017),

	IWSLT'14 De-En				MTTT Fr-En			CNN-DailyMail		
	$k=5$	$k=10$	$k=100$	$k=500$	$k=10$	$k=100$	$k=500$	$k=5$	$k=10$	$k=100$
	(35.6)	(35.4)	(34.7)	(7.9)	(33.0)	(9.9)	(1.2)	(31.5)	(30.9)	(29.1)
BF beam search	93 (24%)	169 (36%)	1275 (79%)	1168 (736%)	184 (16%)	867 (138%)	885 (836%)	200 (33%)	305 (43%)	2960 (92%)
Beam search (ES)	107 (7%)	210 (9%)	2047 (12%)	7685 (27%)	196 (9%)	1310 (58%)	4182 (98%)	224 (19%)	357 (22%)	3942 (59%)
Beam search	115	229	2286	9770	214	2066	8281	266	435	5673

Table 2: Average number of calls (rounded to nearest whole digit) to score, the sequence transduction model, per generated sequence when using different decoding algorithms. Green percentages are performance improvements over standard beam search. Beam search (ES) refers to the OpenNMT early-stopping method (Klein et al., 2017). All methods provably return the same solution and thus, evaluation metrics (in dark blue) for a given beam size are identical.




Figure 1: Number of calls to scoring function score vs. total sequence generation time. Each point is a decoded sequence. Colors represent different model architectures and shapes signify the decoding algorithm used (beam sizes 3 and 10 are included for each). There is no notable difference in the overhead (time-wise) of best-first beam search and beam search.

MTTT Fr-En (Duh, 2018), and CNN-DailyMail (Hermann et al., 2015) using both Transformers (Vaswani et al., 2017) and Convolutional sequence-to-sequence models (Gehring et al., 2017).

For reproducibility, we use the data pre-processing scripts provided by fairseq (Ott et al., 2019) and follow their methods for training sequence transduction models. Hyperparameter are set in accordance with previous works. Specifically, on IWSLT’14 and MTTT tasks, we follow the recommended Transformer settings for IWSLT’14 in fairseq,⁹ which are based on Vaswani et al. (2017) and Gehring et al. (2017). Hyperparameters for models trained on the WMT task are set following version 3 of the Tensor2Tensor toolkit (Vaswani et al., 2018). We use byte-pair encoding (BPE; Sennrich et al. 2016) for all languages. Vocabulary sizes for WMT and IWSLT’14 are set from rec-

⁹<https://github.com/pytorch/fairseq/tree/master/examples/translation>

ommendations for the respective tasks in fairseq; for the MTTT tasks, vocabulary sizes are tuned on models with standard label smoothing regularization. Similarly, the CNN/DailyMail dataset is pre-processed and uses BPE following the same steps as (Lewis et al., 2019). Hyperparameters are the same as for their model fine-tuned on CNN/DailyMail. Details are available on the fairseq website.¹⁰

We use BLEU (Papineni et al., 2002) (evaluated using SacreBLEU (Post, 2018)) for MT metrics and ROUGE-L (Lin, 2004) for abstractive summarization metrics. We build our decoding framework in SGNMT.¹¹

6.1 Running Time

In Tab. 2, we report values as the average number of calls to the scoring function per input; we do not use wall-clock time as this is heavily dependent on hardware. See Fig. 1 for empirical justification of the correlation between calls to the scoring function and runtime on the hardware our experiments were run on. For reference, in our experiments, the scoring function took on average > 99% of the total computation time, even with larger beam sizes, when overhead of the search algorithm is most significant.

We find that best-first (BF) beam search leads to significant speed-ups over both traditional beam search and beam search with early stopping, with a performance increase¹² of $\approx 8x$ for a beam size of 500. We likewise find that best-first beam search offers speed-ups over early stopping methods that are *not* guaranteed to return the same results as standard beam search (see Tab. 3).

¹⁰<https://github.com/pytorch/fairseq/blob/master/examples/bart/README.cnn.md>

¹¹<https://github.com/ucam-smt/sgnmt>

¹²Performance increase is defined as $(\text{old} - \text{new})/\text{new}$

IWSLT'14 De-En				
k	method	search error	BLEU	# calls
10	shrinking	0%	35.4	229 (0%)
	early	0%	35.4	225 (2%)
	BF BS	-	35.4	169 (36%)
100	shrinking	31.7%	13.2	2278 (0%)
	early	31.7%	13.2	1738 (31%)
	BF BS	-	34.7	1275 (79%)
WMT'17 De-En				
10	shrinking	0%	28.6	260 (0%)
	early	0%	28.6	252 (3%)
	BF BS	-	28.6	230 (12%)
100	shrinking	1.7%	26.4	2587 (0%)
	early	1.7%	26.4	2402 (8%)
	BF BS	-	26.9	2046 (26%)

Table 3: BLEU, search error, and average number of calls to score for different stopping criterion. “shrinking” refers to the shrinking beam method of Bahdanau et al. (2015) and “early” refers to the stopping criterion of Huang et al. (2017). Note that neither method is guaranteed to return the same result as standard beam search. Search error and performance increases are with respect to standard beam search.

6.2 Length Normalization

We experiment with both forms of length normalization presented in § 5 and provide results in Tab. 4. We find that both methods, i.e., changing the stopping criterion and using a heuristic during search, provide improvements over baseline BLEU scores albeit with different hyperparameter settings; increases are similar to improvements reported by Murray and Chiang (2018). Notably, using a heuristic causes a large percentage of search errors with respect to standard beam search using the same scoring function. However, the difference in results appears to be *beneficial* in terms of BLEU.

6.3 Mutual Information

We train a language model on the IWSLT dataset and use it to calculate $p(\mathbf{y})$ from (12) as marginalizing over \mathbf{y} is intractable (see Li et al. (2016) for further justification). We run experiments using both of the methods discussed in § 5 and present results in Tab. 5. We find that both methods provide results of equivalent BLEU score compared with the baseline output, i.e., results obtained with the unbounded PMI objective and beam search. Again, despite the high search error rate demonstrated by the heuristic method, evaluation metrics are still comparable.

IWSLT'14 De-En						
	k	β	b	# calls	search error	BLEU
Heuristic	5	0.8	$ \mathbf{x} $	115 (0%)	40.6%	33.9+0.3
	10	1.2	$ \mathbf{x} $	229 (0%)	54.7%	33.8+0.5
Stopping Criterion	5	0.5	n_{max}	73 (58%)	-	33.7+0.1
	10	0.5	n_{max}	130 (76%)	-	33.7+0.4
MTTT Fr-En						
Heuristic	5	0.8	$.7 \mathbf{x} $	100 (8%)	16.2%	33.5+0.2
	10	1.0	$.7 \mathbf{x} $	196 (9%)	25.2%	33.6+0.6
Stopping Criterion	5	1.0	n_{max}	65 (66%)	-	34.1+0.8
	10	1.2	n_{max}	88 (143%)	-	34.1+1.1

Table 4: BLEU search error, and average number of calls to score for output obtained with length normalization scoring function on the IWSLT'14 De-En and MTTT Fr-En test sets. Increase in BLEU is over baseline with no length normalization. Search error and performance increases are with respect to standard beam search decoding using the same scoring function.

6.4 Memory Usage

We conduct a set of experiments where we limit total queue capacity to $k \cdot \gamma$ for $\gamma \in \{1, \dots, n_{max}\}$, as described in § 3.3, and report the BLEU score of the resulting set of hypotheses.

As shown in Tab. 6, we find that restricting the queue capacity does not harm output quality and additionally, leads to even greater runtime performance increase. For example, runtime for decoding of IWSLT'14 with a beam size of 10 can be improved by $> 3x$ while returning results with better evaluation metrics. We find that improvements are even more pronounced for larger beam sizes. Across beam widths and tasks, we find that search error (with respect to standard beam search) is quite low for $\gamma = 5$. Additionally, for smaller γ , the change in BLEU score demonstrates that search error in this context does not necessarily hurt the quality of results.

7 Related Work

Our work is most similar to that of Zhou and Hansen (2005), who propose beam stack search. However, they are focused on exact inference and still evaluate hypotheses in breadth-first order. Additionally, their algorithm requires $\mathcal{O}(n_{max}k)$ memory; while best-first beam search has the same requirements, we introduce effective methods for reducing them, namely memory-reduced best-first beam search.

Huang et al. (2017) propose and prove the optimality of an early-stopping criterion for beam

	k	ε	β	# calls	search error	BLEU
Baseline	5	-	.05	115	-	33.2
	10	-	.05	229	-	33.0
Heuristic	5	.02	.05	129 (0%)	42.7%	33.2
	10	.02	.05	256 (0%)	42.7%	33.0
Stopping Criterion	5	3e-4	.05	114 (1%)	29.2%	33.2
	10	5e-5	.05	224 (2%)	26.6%	33.0

Table 5: BLEU scores with mutual information scoring function on IWSLT’14 De-En. Baseline is PMI decoding with *unbounded* $p(y)$, i.e., $\varepsilon = 0$. Search error is with respect to beam search decoding of baseline with same β .

IWSLT’14 De-En				
k	γ	search error	BLEU	# calls
5	2	22.7%	35.7+0.1	43.8 (163%)
	5	4.4 %	35.8+0.2	79.8 (44%)
	n_{max}	-	35.6	93.0 (24%)
10	2	22.6%	35.7+0.3	48.4 (374%)
	5	4.5%	35.6+0.2	126.9 (81%)
	n_{max}	-	35.4	169.0 (36%)
WMT’17 De-En				
5	2	29.0%	29.7+0.2	77.5 (75%)
	5	1.2%	29.5+0.0	115.8 (12%)
	n_{max}	-	29.5	118.8 (10%)
10	2	36.6%	29.5+0.2	97.3 (165%)
	5	2.6%	29.3+0.0	230.0 (12%)
	n_{max}	-	29.3	230.2 (12%)

Table 6: BLEU scores and the number of calls to score on the IWSLT’14 De-En validation set and WMT’17 De-En test set with queue size restricted to $n_{max} \cdot k$. Note that $\gamma = n_{max}$ is the standard best-first beam search algorithm. Performance increases are over standard beam search. Search error is with respect to beam search with same beam width.

search. The authors find in practice though that reduction in computation from their algorithm was generally not significant. We build on this work and introduce additional methods for avoiding unnecessary computation. Our method leads to better performance, as shown in Tab. 2.

Klein and Manning (2003) use A* for PCFG parsing; however, they use the un-pruned version for exact search which is not applicable for NMT or AS as the memory requirements of the algorithm are far too large for these tasks. Subsequently, Pauls and Klein (2009) provide a method for pruning this search algorithm, albeit using a threshold rather than explicitly limiting the state space. Huang et al. (2012) also adapt A* for a k -best

decoding algorithm. While their methods differ notably from ours, they likewise employ pruning techniques that allow for substantial speedups.

Stahlberg and Byrne (2019) create an exact inference algorithm for decoding and use it to analyze the output of neural NMT models. While they likewise employ the monotonicity of the scoring function to make their method tractable, they do not focus on speed or mimicking the results of standard beam search.

8 Conclusion

We propose best-first beam search, an algorithm that allows for faster decoding while still guaranteeing k -optimality. We provide results on several sequence-to-sequence transduction tasks that show the speed-ups our algorithm provides over standard beam search for decoding neural models. We adapt several popular alternate scoring functions to best-first beam search and provide a framework that can be used to adapt other scoring methods such as coverage normalization (Wu et al., 2016) or diverse beam search (Vijayakumar et al., 2016). We also provide a memory-reduced version of our algorithm, which returns competitive results in a fraction of the time needed for standard beam search.

References

- M. D. Atkinson, J.-R. Sack, N. Santoro, and T. Strothotte. 1986. *Min-max heaps and generalized priority queues*. *Commun. ACM*, 29(10):996–1000.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. *Neural machine translation by jointly learning to align and translate*. In *3rd International Conference on Learning Representations*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. *Findings of the 2017 conference on machine translation (wmt17)*. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 169–214, Copenhagen, Denmark. Association for Computational Linguistics.
- Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. *Wit³: Web inventory of transcribed and translated talks*. In *Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT)*, pages 261–268.

- Edsger W. Dijkstra. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271.
- Kevin Duh. 2018. The multitarget TED talks task. <http://www.cs.jhu.edu/~kevinduh/a/multitarget-tedtalks/>.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. **Understanding back-translation at scale**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, and Yann Dauphin. 2017. **A convolutional encoder model for neural machine translation**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 123–135, Vancouver, Canada. Association for Computational Linguistics.
- Peter E Hart, Nils J Nilsson, and Bertram Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107.
- Wei He, Zhongjun He, Hua Wu, and Haifeng Wang. 2016. Improved neural machine translation with SMT features. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI’16*, page 151–157. AAAI Press.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. **Teaching machines to read and comprehend**. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 28, pages 1693–1701. Curran Associates, Inc.
- Liang Huang, Kai Zhao, and Mingbo Ma. 2017. **When to finish? optimal beam search for neural text generation (modulo beam size)**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2134–2139, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhiheng Huang, Yi Chang, Bo Long, Jean-Francois Crespo, Anlei Dong, Sathiya Keerthi, and Su-Lin Wu. 2012. **Iterative Viterbi A* algorithm for k-best sequential decoding**. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 611–619, Jeju Island, Korea. Association for Computational Linguistics.
- Dan Klein and Christopher D. Manning. 2003. **A* parsing: Fast exact Viterbi parse selection**. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 119–126.
- G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush. 2017. **OpenNMT: Open-Source Toolkit for Neural Machine Translation**. *ArXiv e-prints*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. **OpenNMT: Open-source toolkit for neural machine translation**. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. 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.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. **Conditional random fields: Probabilistic models for segmenting and labeling sequence data**. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML ’01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. **Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *arXiv:1910.13461*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. **A diversity-promoting objective function for neural conversation models**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. 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.
- Andrew McCallum, Dayne Freitag, and Fernando C. N. Pereira. 2000. Maximum entropy Markov models for information extraction and segmentation. In *Proceedings of the Seventeenth International Conference on Machine Learning, ICML ’00*, page 591–598, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- 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, Belgium, Brussels. Association for Computational Linguistics.
- Joakim Nivre, Igor M. Boguslavsky, and Leonid L. Iomdin. 2008. **Parsing the SynTagRus treebank of Russian**. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling*

- 2008), pages 641–648, Manchester, UK. Coling 2008 Organizing Committee.
- 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 NAACL-HLT 2019: Demonstrations*, pages 48–53.
- 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 on Association for Computational Linguistics*, pages 311–318.
- Adam Pauls and Dan Klein. 2009. **Hierarchical search for parsing**. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 557–565, Boulder, Colorado. 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.
- 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*, pages 1715–1725.
- Iulian Serban, Tim Klinger, Gerald Tesauro, Kartik Talamadupula, Bowen Zhou, Yoshua Bengio, and Aaron Courville. 2017. **Multiresolution recurrent neural networks: An application to dialogue response generation**.
- 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.
- Raphael Shu and Hideki Nakayama. 2018. **Improving beam search by removing monotonic constraint for neural machine translation**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 339–344, Melbourne, Australia. 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 3354–3360, Hong Kong, China. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. **Sequence to sequence learning with neural networks**. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Ben Taskar, Carlos Guestrin, and Daphne Koller. 2004. **Max-margin markov networks**. In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 25–32. MIT Press.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, François Chollet, Aidan N. Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. 2018. **Tensor2tensor for neural machine translation**. *CoRR*.
- 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, pages 5998–6008.
- Tim Vieira, Ryan Cotterell, and Jason Eisner. 2016. **Speed-accuracy tradeoffs in tagging with variable-order CRFs and structured sparsity**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David J. Crandall, and Dhruv Batra. 2016. **Diverse beam search: Decoding diverse solutions from neural sequence models**. *CoRR*, abs/1610.02424.
- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. *ArXiv*, abs/1506.05869.
- Andrew Viterbi. 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Gregory S. Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *ArXiv*, abs/1609.08144.
- Yilin Yang, Liang Huang, and Mingbo Ma. 2018. **Breaking the beam search curse: A study of (re-)scoring methods and stopping criteria for neural machine translation**. In *Proceedings of the 2018*

Conference on Empirical Methods in Natural Language Processing, pages 3054–3059, Brussels, Belgium. 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 H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5753–5763. Curran Associates, Inc.

Rong Zhou and Eric A. Hansen. 2005. Beam-stack search: Integrating backtracking with beam search. In *Proceedings of the Fifteenth International Conference on International Conference on Automated Planning and Scheduling*, ICAPS'05, page 90–98. AAAI Press.