

Determinantal Beam Search

Clara Meister¹ Martina Forster² Ryan Cotterell²
¹ETH Zürich ²University of Cambridge
meistecl@inf.ethz.ch martfors@ethz.ch
ryan.cotterell@inf.ethz.ch

Abstract

Beam search is a go-to strategy for decoding neural sequence models. The algorithm can naturally be viewed as a subset optimization problem, albeit one where the corresponding set function does not reflect interactions between candidates. Empirically, this leads to sets often exhibiting high overlap, e.g., strings may differ by only a single word. Yet in use-cases that call for multiple solutions, a diverse or representative set is often desired. To address this issue, we propose a reformulation of beam search, which we call determinantal beam search. Determinantal beam search has a natural relationship to determinantal point processes (DPPs), models over sets that inherently encode intra-set interactions. By posing iterations in beam search as a series of subdeterminant maximization problems, we can turn the algorithm into a diverse subset selection process. In a case study, we use the string subsequence kernel to explicitly encourage n -gram coverage in text generated from a sequence model. We observe that our algorithm offers competitive performance against other diverse set generation strategies in the context of language generation, while providing a more general approach to optimizing for diversity.

1 Introduction

The decoding of neural sequence models is a fundamental component of many tasks in NLP. Yet, many proposed decoding methods aim to produce only a single solution; further, decoding strategies that provide a set, such as beam search, admit high overlap between solutions. Such approaches fail to reflect that for many NLP tasks,¹ there can be multiple correct solutions—or that we may desire a diverse set of solutions. As it stands, standard beam search chooses items based purely on individual

¹As concrete examples, in machine translation there almost always exist multiple ways to translate a sentence; in story generation, we often seek creative language or multiple options to choose from.

scores, with no means for encoding interaction between candidates; this is the limitation which we attempt to address in this work.

We derive **determinantal beam search**, a novel generalization of beam search that casts subset selection as the subdeterminant optimization problem. Specifically, we formulate each iteration of beam search as a subdeterminant maximization problem parameterized by a positive semi-definite matrix that encodes interactions between the possible candidates; standard beam search is recovered by a specific diagonal matrix. This framing creates a natural paradigm for taking the relationships between candidates during the decoding process, and can thus assign higher scores to diversified sets; we show how this approach relates to k -determinantal point processes (DPPs). Given the wealth of research on efficient kernel computation (Rousu and Shawe-Taylor, 2005; Farhan et al., 2017) and DPP inference strategies (Li et al., 2016; Han et al., 2017; Chen et al., 2018), we find the impact on runtime to be quite reasonable in comparison to standard decoding techniques.

In a case study on neural machine translation (NMT), we demonstrate how to make use of the string subsequence kernel (Lodhi et al., 2002) to encode the notion of n -gram diversity in the language generation process, allowing us to derive an elegant diverse beam search. Under this scheme, we observe that determinantal beam search generates more diverse sets than standard beam search with minimal trade-off in terms of BLEU. We see improved performance over stochastic beam search (SBS; Kool et al., 2019), which is reported to encourage diversity, and a slight improvement over Vijayakumar et al. (2018)’s diverse beam search (DBS) while providing a more general approach to optimizing for intra-set diversity.

2 Neural Sequence Models

Neural sequence models are probability distributions $p(\mathbf{y} \mid \mathbf{x})$ over sequences \mathbf{y} in an output

space \mathcal{Y} conditioned on an input \mathbf{x} .² Here we define \mathcal{Y} as the set of all valid sequences derived from a vocabulary \mathcal{V} that are bookended by distinguished BOS and EOS tokens, indicating the beginning and end of the sequence, respectively. Typically, the sequence length is upper-bounded by some value $n_{\max} \in \mathbb{Z}_+$, which may depend on \mathbf{x} . In this work, we consider locally normalized models, i.e. where p is a probability distribution over $\bar{\mathcal{Y}} \stackrel{\text{def}}{=} \mathcal{Y} \cup \{\text{EOS}\}$ conditioned on previously generated tokens $\mathbf{y}_{<t}$. The probability of the full sequence $\mathbf{y} = \langle y_1, y_2, \dots \rangle$ is then calculated via the chain rule of probability:

$$p(\mathbf{y} | \mathbf{x}) = \prod_{t=1}^{|\mathbf{y}|} p(y_t | \mathbf{y}_{<t}, \mathbf{x}) \quad (1)$$

where $\mathbf{y}_{<1} = y_0 \stackrel{\text{def}}{=} \text{BOS}$. Our model p is typically parameterized by a neural network with weights θ . As we do not focus on the underlying model itself in this work, we omit the dependence of p on the parameters θ .

We define the **decoding problem** as the search for the highest-scoring \mathbf{y} among all sequences in \mathcal{Y} according to the model $p(\mathbf{y} | \mathbf{x})$, which is also called maximum-a-posteriori (MAP) inference:

$$\mathbf{y}^* = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \log p(\mathbf{y} | \mathbf{x}) \quad (2)$$

where the log transform of p is used by convention. We further define the **set decoding problem** as the search for a set Y^* of a specified cardinality k among all valid subsets $\{Y' \subseteq \mathcal{Y} \mid |Y'| = k\}$ that has the highest score where, by overloading, we define

$$p(Y | \mathbf{x}) \stackrel{\text{def}}{=} \prod_{\mathbf{y} \in Y} p(\mathbf{y} | \mathbf{x}) \quad (3)$$

Similarly to Eq. (2), the set-decoding problem is then defined as:

$$Y^* = \operatorname{argmax}_{\substack{Y' \subseteq \mathcal{Y}, \\ |Y'| = k}} \log p(Y' | \mathbf{x}) \quad (4)$$

However, as has been noted in the literature, there are a number of issues with both Eq. (2) and (4). First, as \mathcal{Y} may be an exponentially large (in \mathcal{V}) space and p is typically non-Markovian, we cannot efficiently search over \mathcal{Y} , much less over \mathcal{Y}^k . Second, specifically for language generation tasks, these might not be useful objectives.

² \mathbf{x} may be, e.g., a source sentence or an image.

Degenerate Objective. It is important to note that the highest-probability solutions under neural sequence models are not always high-quality; specifically for tasks involving language generation, e.g., machine translation, prior work has shown the tendency for MAP decoding to lead to generic or degenerate solutions (Stahlberg and Byrne, 2019; Meister et al., 2020; Eikema and Aziz, 2020) while superior solutions assigned only slightly lower probability are often overlooked (Holtzman et al., 2020). Consequently, heuristic search methods or alternative objectives are frequently employed for decoding language generators.

2.1 Beam Search

A common heuristic to approximate the decoding problem in Eq. (2) is to sequentially choose the token y_t at each time step t that maximizes $p(y_t | \mathbf{y}_{<t}, \mathbf{x})$ until the EOS token is generated or the maximum sequence length n_{\max} is reached. This procedure is known as greedy search. Beam search is an oft-employed generalization of greedy search that returns k candidates and explores more of the search space.³ In this work, we focus on a framing of beam search as **iterative subset selection**, which allows for a remarkably concise formulation of the algorithm. Given an initial set Y_0 containing only the BOS token, we choose subsequent Y_t for $t \in \{1, \dots, n_{\max}\}$ according to the following recursion:

Standard Beam Search

$$\begin{aligned} Y_0 &\leftarrow \{\text{BOS}\} \\ Y_t &\leftarrow \operatorname{argmax}_{\substack{Y'_t \subseteq \mathcal{B}_t, \\ |Y'_t| = k}} \log p(Y'_t | Y_{t-1}, \mathbf{x}) \end{aligned} \quad (5)$$

where we are constrained to only extending candidates present in the **beam set**, which we define as

$$\mathcal{B}_t \stackrel{\text{def}}{=} \{\mathbf{y}_{<t} \circ y \mid \mathbf{y}_{<t} \in Y_{t-1} \text{ and } y \in \bar{\mathcal{V}}\} \quad (6)$$

where \circ is used to indicate string concatenations. Note that candidates in Y_{t-1} already ending in EOS are simply added directly to \mathcal{B}_t , i.e., $\text{EOS} \circ \text{EOS} = \text{EOS}$. Under this definition, we have the cardinality constraint $|\mathcal{B}_t| \leq |\bar{\mathcal{V}}| \cdot k$.

³A number of NLP tasks only take the highest-scoring element of the returned set Y while other tasks utilize the entire set of solutions.

2.2 A Determinantal Reformulation

We now introduce an alternative, equivalent notation for Eq. (5) using matrices and determinants that will shed light on the straightforward generalization of beam search that we present as the primary contribution of this paper. We define a timestep-dependent⁴ diagonal matrix $D \in \mathbb{R}^{|\mathcal{B}_t| \times |\mathcal{B}_t|}$ where we take the diagonal entry

$$D_{ii} \stackrel{\text{def}}{=} p(\mathbf{y}_{\leq t}^{(i)} | \mathbf{x}) \quad (7)$$

Here $\mathbf{y}_{\leq t}^{(i)}$ is the i^{th} candidate in \mathcal{B}_t according to a unique mapping of every element $\mathbf{y}_{\leq t} \in \mathcal{B}_t$ to an integer between 1 and $|\mathcal{B}_t|$. Furthermore, we use the notation D_{Y_t} where $Y_t \subseteq \mathcal{B}_t$, to indicate the submatrix that only contains those rows and columns corresponding to the elements of Y_t . We may now rewrite Eq. (5) as

Determinantal Standard Beam Search

$$\begin{aligned} Y_0 &\leftarrow \{\text{BOS}\} \\ Y_t &\leftarrow \operatorname{argmax}_{\substack{Y_t' \subseteq \mathcal{B}_t, \\ |Y_t'|=k}} \log \det(D_{Y_t'}) \end{aligned} \quad (8)$$

where equivalence follows from the definition of the determinant for diagonal matrices. Formally, Eq. (8) is known as the **subdeterminant maximization problem**⁵ (Klee et al., 1995; Ebrahimi et al., 2017), which—as the name suggests—refers to the problem of finding the determinant maximizing subset of a matrix. While the notation introduced in Eq. (8) may seem contrived, it allows us to perform the subsequent generalization.

3 Determinantal Beam Search

We are now in a position to ask the fundamental question of this work: What happens if we replace the diagonal matrix D with a non-diagonal matrix? This substitution allows us to account for interactions between the elements in the beam. Formally, we consider a timestep-dependent positive semi-definite (PSD) matrix $D + w \cdot K$ where the off-diagonal matrix K indicates the strength of the interactions between candidates. The non-negative weight $w \geq 0$ controls the importance of these interactions during the decoding process. In this case, the beam search recursion becomes:

⁴We have omitted the time-step dependence of D for notational brevity as it is always clear from context.

⁵Albeit with a cardinality constraint.

Full Determinantal Beam Search

$$\begin{aligned} Y_0 &\leftarrow \{\text{BOS}\} \\ Y_t &\leftarrow \operatorname{argmax}_{\substack{Y_t' \subseteq \mathcal{B}_t, \\ |Y_t'|=k}} \log \det(D_{Y_t'} + w \cdot K_{Y_t'}) \end{aligned} \quad (9)$$

Clearly, we recover beam search when $w = 0$; however, we can now select subsets based additionally on candidate interactions. That is, Eq. (9) now has an interpretation as a **diversity objective function** (Indyk et al., 2014) when K is chosen wisely. Due to the presence of the log, Eq. (9) is only well defined when the matrix $D_Y + w \cdot K_Y$ is PSD.⁶

3.1 Constructing K

One simple way to construct K is as a **Gram matrix**, where each i, j element of K is computed via a kernel function $\mathcal{K} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ that maps two items in a space \mathcal{S} to a real number. Specifically, we define $K_{ij} = \mathcal{K}(s_i, s_j)$ where $s_i, s_j \in \mathcal{S}$ are the i^{th} and j^{th} elements of \mathcal{S} , respectively. In slight abuse of notation, we overload the kernel function \mathcal{K} to take a set S such that $K = \mathcal{K}(S)$ is the kernel matrix resulting from pairwise computation over elements of S .⁷ Following from Mercer’s theorem, the matrix $K = \mathcal{K}(S)$ is necessarily PSD and, thus the matrix $D_Y + w \cdot K_Y$ is PSD for any $Y \subseteq S$.⁸

The efficient computation of kernel functions is a well-studied problem—largely due to the prevalence of kernels in various machine learning techniques. For example, dynamic programming techniques are often employed in computation of $\mathcal{K}(S)$ (Rousu and Shawe-Taylor, 2005) or approximate low-rank kernel matrices can be used in place of $\mathcal{K}(S)$ (Si et al., 2017).

3.2 Relation to a DPPs

One interpretation of Eq. (9) is as a determinantal point process (DPP). Specifically, it is a k -DPP

⁶To see this, recall that the determinant is the product of the eigenvalues. To ensure that the determinant is strictly positive, we can simply enforce that all the eigenvalues are positive, which is necessarily the case for PSD matrices. Note that in the case where any of the eigenvalues of a submatrix are zero, we take $\log \det(\cdot) = -\infty$.

⁷In machine learning literature, the term “kernel” is often used to refer to both the function \mathcal{K} and the kernel matrix K .

⁸To see this, note that the matrix D is necessarily PSD. Since PSD matrices are closed under addition and multiplication by a positive scalar, then necessarily $D + w \cdot K$ is PSD. Lastly, any submatrix of a PSD matrix is also PSD, which makes $D_Y + w \cdot K_Y$ a PSD matrix.

(Kulesza and Taskar, 2011) in the L -ensemble parameterization where we have $L = D + w \cdot K$. This interpretation as a k -DPP gives us a very clear understanding of *why* Eq. (8) yields a diverse beam search. The diagonal entries encode quality, which tells how “good” each candidate on the beam is, while the off-diagonal entries encode how similar two elements are and, thus, how much they should be repulsed. For an overview of DPPs we refer the reader to Kulesza and Taskar (2012).

3.3 Computing Log-Determinants

Unfortunately, computing the argmax^9 in Eq. (9) is an NP-hard problem (Ko et al., 1995). However, as the subdeterminant maximization problem has many applications, there has been much research on efficient algorithms for approximating log-determinants in the context of, e.g., determinantal point processes (Gillenwater et al., 2012; Han et al., 2017).¹⁰ One such algorithm uses a first-order approximation of the log-determinant function (Han et al., 2017). The work of Chen et al. (2018) uses a greedy, iterative approach; by updating the Cholesky factorization of the matrix kernel incrementally, the algorithm reduces inference time to $\mathcal{O}(k^2|\mathcal{S}|)$ to return k candidates from set \mathcal{S} . Pseudocode for the latter approach can be found in Chen et al. (2018); pseudocode for the algorithm in log-space—since probabilistic models are often worked with in log-space for numerical stability—can be found in App. A.

3.4 Runtime Analysis

We consider the runtime of selecting k candidates at any given time step in the recursion of Eq. (9). At each time step, we must first construct the matrix K . This computation is highly dependent on the set interactions being modeled; as such, let $\mathcal{O}(c(k))$ be a runtime bound for K ’s computation when our search uses a beam size of k . Once we have constructed our matrix $D + w \cdot K$, we must next select k items. The set of hypotheses at any time step is at most $k|\bar{\mathcal{V}}|$. While as discussed in §3.3, finding the size- k subset that exactly

⁹We may also sample from the k -DPP modeled by Eq. (9) rather than taking the approximate mode; this would only require changing the inference algorithm and can be done in a similarly efficient manner (Li et al., 2016). We focus on deterministic methods in this work as we aim to find the objective maximizing set.

¹⁰As beam search is already a heuristic approach, such an approximation does not have any theoretical implications for the results of our algorithm.

maximizes Eq. (9) has exponential runtime, we assume approximate methods are employed. Using the method given by Chen et al. (2018), approximate MAP inference takes $k^3|\bar{\mathcal{V}}|$ time to return k items from a set of size $k|\bar{\mathcal{V}}|$. Thus, the runtime at each iteration of determinantal beam search under these conditions would be $\mathcal{O}(c(k) + k^3|\bar{\mathcal{V}}|)$. Note that standard beam search runs in $\mathcal{O}(k|\bar{\mathcal{V}}|\log(k|\bar{\mathcal{V}}|))$ time at each iteration. As k is generally small (≤ 20) and the impact of $c(k)$ can be made reasonable (§3.1), the practical increase in runtime is typically only moderate.

4 Case Study: Diverse Beam Search

We now consider the task of language generation, where our vocabulary $\bar{\mathcal{V}}$ is a set of words and \mathcal{Y} is the set of all valid strings derived from $\bar{\mathcal{V}}$. When the space of our kernel function $\mathcal{S} = \mathcal{B}_t$, one simple way of modeling interactions is through a string subsequence kernel (Lodhi et al., 2002).

4.1 Computing the String Kernel

The string subsequence kernel, proposed by Lodhi et al. (2002), is a function over two strings \mathbf{s} and \mathbf{t} computed as:

$$\mathcal{K}(\mathbf{s}, \mathbf{t}) = \sum_{u \in \mathcal{V}^n} \sum_{i: u = \mathbf{s}[i]} \lambda^{l(\mathbf{i})} \sum_{\mathbf{j}: u = \mathbf{t}[\mathbf{j}]} \lambda^{l(\mathbf{j})} \quad (10)$$

where \mathcal{V}^n is the set of all finite strings of length n over the alphabet \mathcal{V} ; \mathbf{i} (or \mathbf{j}) denotes a vector of indices $\mathbf{i} = (i_1, \dots, i_{|\mathbf{i}|})$ where $1 < i_1 < i_{|\mathbf{i}|} \leq |\mathbf{s}|$; $l(\mathbf{i}) \stackrel{\text{def}}{=} i_{|\mathbf{i}|} - i_1 + 1$ is the length of the substring u in \mathbf{s} ; $\lambda \in (0, 1]$ is a decay factor which serves as a penalty for gaps within a compared subsequence.

Direct computation of Eq. (10) is exponential in $|\mathcal{V}|$, but efficient dynamic programs can be utilized: In this work, we employ the trie-based methods of Rousu and Shawe-Taylor (2005) to compute Eq. (10). Under this scheme, the computation of the kernel between two strings \mathbf{s} and \mathbf{t} is $\mathcal{O}(n \cdot M \cdot \log(\max(|\mathbf{s}|, |\mathbf{t}|)))$, where n is the chosen subsequence length (a hyperparameter) and M is the number of words that strings \mathbf{s} and \mathbf{t} have in common. Note that $|\mathbf{s}|$, and thus M , are bounded by the time step t . Further, we can reuse many of the computations between subsequent decoding rounds due to the iterative nature of both beam search and the subsequence kernel computations. Additionally, since the magnitude of Eq. (10) is influenced by the lengths of \mathbf{s} and \mathbf{t} , we normalize

the kernel as follows:

$$\mathcal{K}_{\text{norm}}(\mathbf{s}, \mathbf{t}) = \frac{\mathcal{K}(\mathbf{s}, \mathbf{t})}{\sqrt{\mathcal{K}(\mathbf{s}, \mathbf{s}) \cdot \mathcal{K}(\mathbf{t}, \mathbf{t})}} \quad (11)$$

4.2 Integration into DetBS

The string subsequence kernel gives us a straightforward method for decoding diverse sets of strings from language generators. We construct the matrix

$$D + w \cdot \mathcal{K}(\mathcal{B}_t) \quad (12)$$

using the dynamic program mentioned above to compute $\mathcal{K}(\mathcal{B}_t)$. Intuitively, we can expect the argmax—i.e., the size k set corresponding to the objective-maximizing submatrix—of $D + w \cdot \mathcal{K}(\mathcal{B}_t)$ to have higher subsequence diversity as w is increased. This is perhaps most easily seen when viewing our problem as a k -DPP: if strings $\mathbf{y}^{(i)}$ and $\mathbf{y}^{(j)}$ have high overlap, this will be reflected in the matrix $\mathcal{K}(\mathcal{B}_t)$ at position i, j . Higher values of $\mathcal{K}(\mathcal{B}_t)_{i,j} = \mathcal{K}(\mathbf{y}^{(i)}, \mathbf{y}^{(j)})$ lead to lower probability of both $\mathbf{y}^{(i)}$ and $\mathbf{y}^{(j)}$ being in the set drawn according to the k -DPP parameterized by $D + w \cdot \mathcal{K}(\mathcal{B}_t)$, which follows from the properties of DPPs outlined in §2.2. In short, higher values of $\mathcal{K}(\mathbf{y}^{(i)}, \mathbf{y}^{(j)})$ decrease the value of $\log \det(D_Y + w \cdot \mathcal{K}(\mathcal{B}_t)_Y)$ for sets Y containing both $\mathbf{y}^{(i)}$ and $\mathbf{y}^{(j)}$, which makes Y less likely to be chosen in the recursion of Eq. (9).

5 Experiments

In our experiments, we explore the use of deterministic beam search as a diverse decoding strategy for language generation.

5.1 Baselines

Various diverse decoding strategies exist in the NLP literature. We first discuss those strategies that we employ as baselines in our experiments.

Standard Beam Search. Beam search is one of the most widely used decoding algorithms in NLP, where many problems require efficient strategies for decoding solutions from structured predictors. Specifically, for language generation tasks, beam search has repeatedly proved its effectiveness at decoding state-of-the-art solutions (Wu et al., 2016; Serban et al., 2017; Edunov et al., 2018; Yang et al., 2019). We refer back to §2.1 for the algorithm.

Stochastic Beam Search. Kool et al. (2019) propose stochastic beam search (SBS), a decoding technique that samples without replacement from sequence models according to their distribution over the entire space \mathcal{Y} . For random sampling methods such as SBS, it is customary to use a sampling temperature $T > 0$ at generation time to control for the peakiness of the sampling distribution. This results in the generalized softmax:

$$\begin{aligned} p_T(y \mid \mathbf{y}_{<t}, \mathbf{x}) & \quad (13) \\ &= \frac{\exp\left(\log p(y \mid \mathbf{y}_{<t}, \mathbf{x})/T\right)}{\sum_{y' \in \mathcal{V}} \exp\left(\log p(y' \mid \mathbf{y}_{<t}, \mathbf{x})/T\right)} \end{aligned}$$

where larger T may lead to more diverse sets simply due to additional smoothing.

Diverse Beam Search. Vijayakumar et al. (2018) propose a modification to the standard beam search algorithm—which they term diverse beam search (DBS)—to alleviate lack of diversity. The algorithm further divides the beam into G groups $\mathcal{B}_t^1, \dots, \mathcal{B}_t^G$, where G is a hyperparameter of the algorithm, and optimizes for diversity between the different groups by subtracting a similarity term $\Delta(\mathbf{y}_{\leq t}, \mathcal{B}_t^g)$ from the decoding objective.¹¹ Specifically, $\Delta(\mathbf{y}_{\leq t}, \mathcal{B}_t^g)$ represents the degree of similarity between a hypothesis $\mathbf{y}_{\leq t}$ and a group of hypotheses \mathcal{B}_t^g . They find $G = k$, i.e., each group contains a single hypothesis, and the Hamming distance similarity metric lead to the best results; we use these settings in our experiments. Note that under this scheme, the solution set may have duplicates if the diversity penalty is not large enough.

Notably, under the above experimental settings, the runtimes of diverse beam search and our algorithm are the same, up to computation of the hamming loss and string kernel, respectively. However, while string kernel computations in our algorithm can be done in parallel, the diversity penalty in diverse beam search must be computed sequentially for each hypothesis, as it is based on the previously chosen groups.

5.2 Setup

We run experiments on neural machine translation (NMT) models trained on the WMT’14 (Bogjar et al., 2014) En–Fr and the WMT’19 (Barrault

¹¹The diversity term has coefficient w to determine the strength of the penalty. When this weight is 0 or sufficiently small, all groups will return the same solution(s).

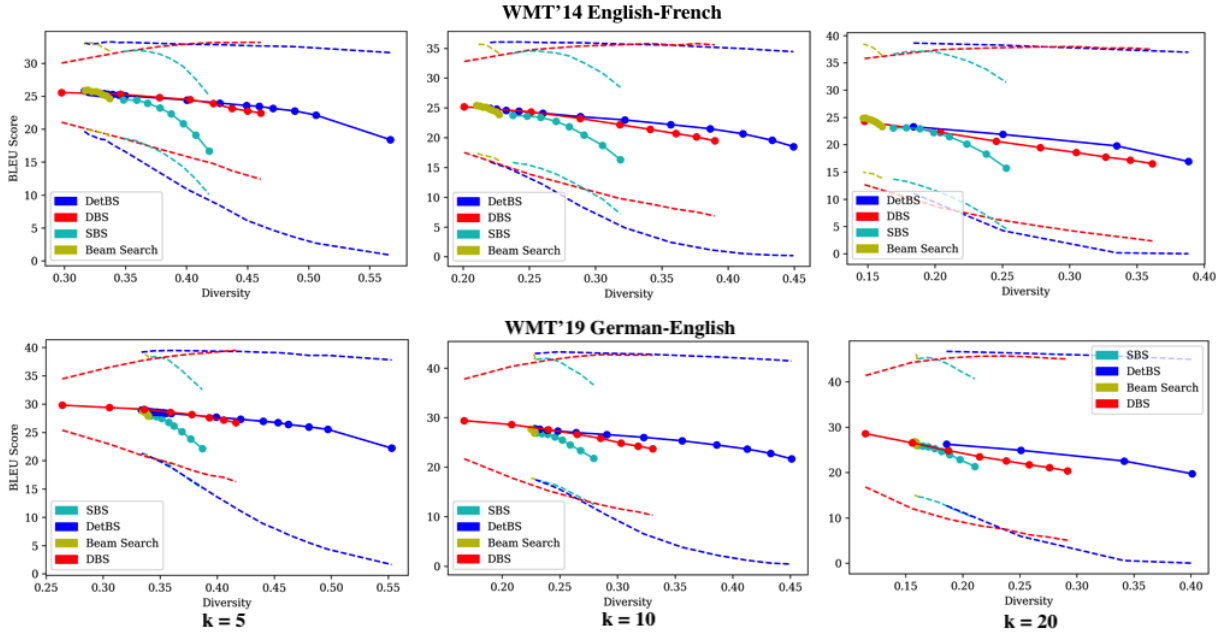


Figure 1: Averaged n -gram diversity vs. minimum, median, and maximum BLEU score for beam sizes $k = 5, 10, 20$ on WMT’14 En–Fr and WMT’19 De–En newstest using various decoding strategies. The free parameter for each strategy is either the softmax temperature or the weight of the diversity parameter (see §5.2).

et al., 2019) De–En datasets; for reproducibility, we use the pretrained models made available by fairseq¹² (Ott et al., 2019). We evaluate on the newstest set from the respective datasets, each containing 3003 sentences. Further details can be found in App. B.

For determinantal beam search (DetBS), we perform a hyperparameter search (precise details likewise in App. B) over λ and n , the decay factor and subsequence length, respectively. Search is performed for fixed $w = 0.1$ and $k = 10$ on validation sets for both languages; we omit a search over the entire space of w, k, λ, n so as to not create an unfair advantage for DetBS in comparison with the other decoding strategies, for which no hyperparameters are tuned. We use subsequence length $n = 2$ and $\lambda \in \{0.1, 0.3\}$ for De–En and En–Fr, respectively.

We decode sets of size $k \in \{5, 10, 20\}$ with each strategy, comparing sentence-level BLEU and n -gram coverage d_n averaged across $n \in \{1, 2, 3, 4\}$ in the decoded sets, where we define d_n as

$$d_n = \frac{\text{\#of unique } n\text{-grams in } k \text{ strings}}{\text{\#of } n\text{-grams in } k \text{ strings}} \quad (14)$$

¹²<https://github.com/pytorch/fairseq/tree/master/examples/translation>

¹³For each decoding strategy, we choose the diversity parameter corresponding to the most diverse set that had median BLEU 28.5 ± 0.05 .

While d_n has a more natural interpretation as *coverage* of different n -grams, the above quantity is often referred to as n -gram diversity in the literature and so we transition to this term for consistency. Following the experimental setup of Kool et al. (2019), we vary sampling temperature $T \in \{0.1, 0.2, \dots, 0.8\}$ in the case of beam search and stochastic beam search and diversity weight $w \in \{0.1, 0.2, \dots, 0.8\}$ in the case of diverse beam search. For DetBS, we observe that larger sets require a smaller diversity penalty to achieve good n -gram diversity: in Fig. 1 we show results for DetBS with the string subsequence kernel for $w \in \{0.01, 0.02, \dots, 0.1, 0.2, 0.3, 0.4\}$ for $k = 5$, $w \in \{0.01, 0.02, \dots, 0.15\}$ for $k = 10$, and $w \in \{0.01, 0.02, \dots, 0.05\}$ for $k = 20$.¹⁴ To observe how BLEU is affected by larger diversity coefficients under DetBS, we explore a finer grain of weights for DetBS in App. C.

5.3 Results

Fig. 1 shows the sentence-level BLEU score and averaged n -gram diversity on the newstest set for different decoding strategies; Tab. 2 shows explicit coverage of 1, 2, 3, 4-grams and averaged across 1, 2, 3, 4-grams for different decoding strategies when BLEU is controlled for. The 3 lines

¹⁴Recall $w = 0$ recovers standard beam search with a temperature of $T = 1$.

Source Sentence	
Zum Abschluss wurde eine Tombola verlost. Die Wahrheit zu sagen ist aber kein Verbrechen.	
Beam Search ($T = 0.6$)	
<ul style="list-style-type: none"> • <i>A raffle was held to close the event.</i> • <i>A raffle was held to conclude the event.</i> • <i>A raffle was held at the end.</i> • <i>At the end a raffle was held.</i> • <i>A raffle was held to close the draw.</i> 	<ul style="list-style-type: none"> • <i>But telling the truth is not a crime.</i> • <i>Telling the truth is not a crime.</i> • <i>But telling the truth isn't a crime.</i> • <i>But telling the truth is no crime.</i> • <i>However, telling the truth is not a crime.</i>
Diverse Beam Search ($w = 0.4$)	
<ul style="list-style-type: none"> • <i>To conclude, a raffle was held.</i> • <i>A raffle was held to close the event.</i> • <i>A raffle was held to close the event.</i> • <i>A raffle was held to close the event.</i> • <i>At the end of the event, a raffle was held.</i> 	<ul style="list-style-type: none"> • <i>But telling the truth is not a crime.</i> • <i>But telling the truth is not a crime.</i> • <i>But telling the truth is not a crime.</i> • <i>But telling the truth is not a crime.</i> • <i>Telling the truth, however, is not a crime.</i>
Determinantal Beam Search ($w = 0.12$)	
<ul style="list-style-type: none"> • <i>Finally, a raffle was held.</i> • <i>A raffle was held at the end.</i> • <i>At the end a raffle was held.</i> • <i>To conclude, a raffle was held.</i> • <i>A raffle was held to close the event.</i> 	<ul style="list-style-type: none"> • <i>But telling the truth is not a crime.</i> • <i>But telling the truth isn't a crime.</i> • <i>Telling the truth is not a crime.</i> • <i>However, telling the truth is not a crime.</i> • <i>But to tell the truth is not a crime.</i>

Table 1: Generated translations for given source sentence from WMT’19 De–En dataset using different decoding strategies. All use a beam size of five.¹³ We see large overlap in Beam Search results while DBS actually returns several of the same results. In comparison, DetBS turns qualitatively diverse results even for simple sentences.

per decoding strategy in Fig. 1 represent the minimum, median, and maximum sentence-level BLEU score out of the k translation options, averaged across the corpus. We consider median BLEU to be the best metric of set text-quality, as a good diverse decoding algorithm should not completely sacrifice BLEU for the sake of diversity. The plots are analogous to those in Kool et al. (2019).

On both datasets and across different set sizes, results indicate that DetBS generates diverse sets of strings while maintaining high median and maximum BLEU scores. We see similar or higher n -gram diversity in comparison to DBS for the same median BLEU and a notably better n -gram diversity vs. BLEU trade-off than standard beam search and SBS. Further, the highest quality translation (shown by max BLEU) does not appear to be sacrificed when the diversity parameter is increased for DetBS. In contrast, there is a notable drop-off for generation strategies in which diversity is controlled for using temperature. We show samples of generated text in Tab. 1.

6 Related Work

Our work is built upon much of the subset optimization literature in machine learning. We base our algorithm off the subdeterminant maximization problem (Agarwal et al., 2004), which

has been used to find core sets—a concept originating in computational geometry concerning the existence of a small, representative set of core items—in data summarization problems (Mirza-soleiman et al., 2013), nearest neighbor search (Abbar et al., 2013) and streaming algorithms (Indyk et al., 2014) *inter alia*. Informally, we can connect our problem to the notion of decoding a core set from sequence models. To the best of our knowledge, our work is the first to use this concept when decoding sequence models.

Wang and Chan (2019) incorporate DPPs into a reinforcement learning objective to optimize for diverse text when training image captioning models. We optimize for diversity during decoding, rather than training, which makes our methods applicable with out-of-the-box models and allows us to avoid highly hyperparameter-sensitive techniques, like minimum-risk training or reinforcement learning-based algorithms, while achieving the same goal. While the application of our methods at training times is an interesting research direction, we foresee technical challenges corresponding to such approaches that may outweigh their benefits.

As a decoding method, our work is closest

¹³In the case that not all strategies had such a set, we instead bounded BLEU by the lowest of the median BLEU across decoding strategies.

	De-En						En-Fr					
	BLEU	1-gram	2-gram	3-gram	4-gram	avg.	BLEU	1-gram	2-gram	3-gram	4-gram	avg.
DetBS	26.24	0.11	0.17	0.21	0.26	0.19	23.29	0.11	0.17	0.21	0.25	0.18
DBS	26.52	0.09	0.14	0.18	0.22	0.16	24.26	0.09	0.14	0.17	0.2	0.15
Beam Search	26.15	0.08	0.14	0.19	0.23	0.16	23.29	0.09	0.14	0.18	0.22	0.16
SBS	25.95	0.08	0.14	0.19	0.24	0.16	23.08	0.1	0.16	0.21	0.25	0.18

Table 2: We report the coverage (as defined in Eq. (14)) of 1, 2, 3, 4-grams and averaged across 1, 2, 3, 4-grams as well as median BLEU for $k = 20$ on the `newstest` dataset. For each decoding strategy, we report metrics on the generated set that has highest (average) d_n , where we set the constraint that median BLEU for the set is still within 1 point of the highest median BLEU (across decoding strategies and diversity parameters).¹⁵

to that of Vijayakumar et al. (2018), who propose a variation of beam search (described in §5.3). However, their algorithm lacks theoretical motivation and is not guaranteed to provide a non-overlapping set; the same solution may appear multiple times in the decoded set if the diversity penalty is not large enough, as shown in Tab. 2. Additionally, groups at each time step t must be processed in order since the score of all hypotheses considered for group $g + 1$ depend on hypotheses in groups $1, \dots, g$, which creates a large bottleneck under the recommended settings of $G = k$.

Random sampling strategies for decoding neural sequence models have received much attention in recent years. While techniques such as stochastic beam search and the UniqueRandomizer (Shi et al., 2020) are convenient for creating statistical estimators and have uses in reinforcement learning techniques due to their clear probabilistic interpretation, there are no diversity guarantees for the set of generated sequences.

Tam (2020) likewise adapts beam search, proposing a k -means clustering version that clusters solutions by averaged word embeddings. As there lacks an interpretation of distance between averaged word embeddings though, it is unclear if the method can explicitly optimize for any tangible notion of coverage or diversity.

7 Conclusion

We propose determinantal beam search (DetBS): a new way of framing beam search that allows us to optimize set generation for diversity and coverage rather than simply individual scores. Formally, we redefine beam search as an iterative subdeterminant maximization problems where we select the approximately maximizing set according to the PSD matrix parameterizing our score function. This gives us the ability to encode the notion of intra-set diversity into the beam

search optimization problem. We discuss and experiment with efficient methods for inference and kernel computation that make DetBS an efficient decoding strategy in practice. We use DetBS in the context of language generation, where we explicitly encourage n -gram coverage through the string subsequence kernel. In our NMT experiments, we find DetBS generates much more diverse sets of strings than standard beam search and stochastic beam search with a small trade-off in median BLEU. We observe competitive performance compared with diverse beam search.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful feedback and recommendations.

Ethical Considerations

While language generation can be used for malicious purposes, e.g., to propagate misinformation or offensive text, we do not foresee any specific ethical concerns with the techniques in this work.

References

- Sofiane Abbar, Sihem Amer-Yahia, Piotr Indyk, Sepideh Mahabadi, and Kasturi R. Varadarajan. 2013. [Diverse near neighbor problem](#). In *Proceedings of the Twenty-Ninth Annual Symposium on Computational Geometry*. Association for Computing Machinery.
- Pankaj K. Agarwal, Sariel Har-Peled, and Kasturi R. Varadarajan. 2004. [Approximating extent measures of points](#). *Journal of the Association for Computing Machinery*, 51(4):606–635.
- 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](#). In *Proceedings of the Fourth Conference on*

- Machine Translation*, pages 1–61. Association for Computational Linguistics.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58. Association for Computational Linguistics.
- Laming Chen, Guoxin Zhang, and Eric Zhou. 2018. [Fast greedy map inference for determinantal point process to improve recommendation diversity](#). In *Advances in Neural Information Processing Systems*, pages 5622–5633.
- J. B. Ebrahimi, D. Straszak, and N. K. Vishnoi. 2017. [Subdeterminant maximization via nonconvex relaxations and anti-concentration](#). In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1020–1031. IEEE Computer Society.
- 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. 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. International Committee on Computational Linguistics.
- Muhammad Farhan, Juvaria Tariq, Arif Zaman, Mudassir Shabbir, and Imdad Ullah Khan. 2017. [Efficient approximation algorithms for strings kernel based sequence classification](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 6935–6945. Curran Associates, Inc.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1243–1252.
- Jennifer Gillenwater, Alex Kulesza, and Ben Taskar. 2012. [Near-optimal MAP inference for determinantal point processes](#). In *Advances in Neural Information Processing Systems*, volume 25, pages 2735–2743. Curran Associates, Inc.
- Insu Han, Prabhanjan Kambadur, Kyoungsoo Park, and Jinwoo Shin. 2017. [Faster greedy MAP inference for determinantal point processes](#). volume 70 of *Proceedings of Machine Learning Research*, pages 1384–1393.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *Proceedings of the International Conference on Learning Representations*.
- Piotr Indyk, Sepideh Mahabadi, Mohammad Mahdian, and Vahab S. Mirrokni. 2014. [Composable core-sets for diversity and coverage maximization](#). In *Proceedings of the Thirty-Third Association for Computing Machinery SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, page 100–108. Association for Computing Machinery.
- V. Klee, P. Gritzmann, and D. Larman. 1995. [Largest \$j\$ -simplices in \$n\$ -polytopes](#). *Discrete and Computational Geometry*, 13(3-4):477–516.
- Chun-Wa Ko, Jon Lee, and Maurice Queyranne. 1995. [An exact algorithm for maximum entropy sampling](#). *Operations Research*, 43(4):684–691.
- Wouter Kool, Herke Van Hoof, and Max Welling. 2019. [Stochastic beams and where to find them: The Gumbel-top- \$k\$ trick for sampling sequences without replacement](#). In *Proceedings of the International Conference on Machine Learning*, pages 3499–3508.
- Alex Kulesza and Ben Taskar. 2011. [\$k\$ -DPPs: fixed-size determinantal point processes](#). In *Proceedings of the 28th International Conference on Machine Learning*, pages 1193–1200.
- Alex Kulesza and Ben Taskar. 2012. *Determinantal Point Processes for Machine Learning*. Now Publishers Inc.
- Chengtao Li, Stefanie Jegelka, and Suvrit Sra. 2016. [Efficient sampling for \$k\$ -determinantal point processes](#). volume 51 of *Proceedings of Machine Learning Research*, pages 1328–1337.
- Zhifei Li and Jason Eisner. 2009. [First- and second-order expectation semirings with applications to minimum-risk training on translation forests](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 40–51. Association for Computational Linguistics.
- Huma Lodhi, Craig Saunders, John Shawe-Taylor, Nello Cristianini, and Chris Watkins. 2002. [Text classification using string kernels](#). *Journal of Machine Learning Research*, 2:419–444.
- 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. Association for Computational Linguistics.
- Baharan Mirzasoleiman, Amin Karbasi, Rik Sarkar, and Andreas Krause. 2013. [Distributed submodular maximization: Identifying representative elements in massive data](#). In *Advances in Neural Information Processing Systems*, volume 26, pages 2049–2057.

- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. [Facebook FAIR’s WMT19 news translation task submission](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 314–319. 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. Association for Computational Linguistics.
- Juho Rousu and John Shawe-Taylor. 2005. [Efficient computation of gapped substring kernels on large alphabets](#). *Journal of Machine Learning Research*, 6:1323–1344.
- Iulian Vlad Serban, Tim Klinger, Gerald Tesauro, Karthik Talamadupula, Bowen Zhou, Yoshua Bengio, and Aaron Courville. 2017. [Multiresolution recurrent neural networks: An application to dialogue response generation](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, page 3288–3294. AAAI Press.
- Kensen Shi, David Bieber, and Charles Sutton. 2020. [Incremental sampling without replacement for sequence models](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 8785–8795. PMLR.
- Si Si, Cho-Jui Hsieh, and Inderjit S. Dhillon. 2017. [Memory efficient kernel approximation](#). *Journal of Machine Learning Research*, 18(20):1–32.
- 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. Association for Computational Linguistics.
- Yik-Cheung Tam. 2020. [Cluster-based beam search for pointer-generator chatbot grounded by knowledge](#). *Computer Speech and Language*, 64:101094.
- Ashwin Vijayakumar, Michael Cogswell, Ramprasaath Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2018. [Diverse beam search for improved description of complex scenes](#). In *AAAI Conference on Artificial Intelligence*, pages 7371–7379.
- Qingzhong Wang and Antoni B. Chan. 2019. [Towards diverse and accurate image captions via reinforcing determinantal point process](#). *CoRR*, abs/1908.04919.
- 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, Lukasz 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](#).
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. [XLNet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32.

A Log-Space Computations

Algorithm 1 Fast Greedy MAP Inference with log-space parameterization (Chen et al., 2018). We transform computations according to (Li and Eisner, 2009).

Input: L : log of PSD matrix
 k : desired set size

```

1: function GREEDY_MAP_INFERENCE()
2:    $\mathbf{c}_i = [], d_i = L_{ii}, \mathbf{s}_i = []$ 
3:    $j = \operatorname{argmax}_{i \in \mathcal{S}} d_i$ 
4:    $Y_g = \{j\}$ 
5:   while  $|Y_g| \neq k$ :
6:     for  $i \in \mathcal{S} \setminus Y_g$ :
7:        $s, \log\_inner \leftarrow \operatorname{LOGSUMEXP}(\mathbf{c}_j, \mathbf{c}_i, \mathbf{s}_j, \mathbf{s}_i)$ 
8:        $\operatorname{FUNC} \leftarrow \operatorname{LOG\_ADD}$  if  $s < 0$  else  $\operatorname{LOG\_MINUS}$ 
9:       if  $L_{ji} > \log\_inner$ :
10:         $s \leftarrow 1$ 
11:         $e_i = \operatorname{FUNC}(L_{ji}, \log\_inner) - 0.5 \cdot d_j$ 
12:       else
13:         $s \leftarrow -s$ 
14:         $e_i = \operatorname{FUNC}(\log\_inner, L_{ji}) - 0.5 \cdot d_j$ 
15:        $\mathbf{c}_i = [\mathbf{c}_i \ e_i], d_i = d_i - 2 \cdot e_i, \mathbf{s}_i = [\mathbf{s}_i \ s]$ 
16:      $j = \operatorname{argmax}_{i \in \mathcal{S} \setminus Y_g} d_i, Y_g = Y_g \cup \{j\}$ 
17:   return  $Y_g$ 
18: function LOGSUMEXP( $\mathbf{c}_1, \mathbf{c}_2, \mathbf{s}_1, \mathbf{s}_2$ )
19:    $s, \log\_inner \leftarrow 1, -\infty$ 
20:   for  $\langle c_1, c_2, s_1, s_2 \rangle \in \langle \mathbf{c}_1, \mathbf{c}_2, \mathbf{s}_1, \mathbf{s}_2 \rangle$ :
21:      $s' \leftarrow s_1 \cdot s_2$ 
22:      $\operatorname{FUNC} \leftarrow \operatorname{LOG\_ADD}$  if  $s == s'$  else  $\operatorname{LOG\_MINUS}$ 
23:     if  $\log\_inner > c_1 + c_2$ :
24:        $\log\_inner \leftarrow \operatorname{FUNC}(\log\_inner, c_1 + c_2)$ 
25:     else
26:        $s = s'$ 
27:      $\log\_inner \leftarrow \operatorname{FUNC}(c_1 + c_2, \log\_inner)$ 
return  $s, \log\_inner$ 

```

B Experimental Setup

Hyperparameters. As we use the string subsequence kernel of section §4 in DetBS, there are a number of hyperparameters that can be adjusted beyond the diversity weight w : the decay factor λ indicates the degree to which interior gaps are penalized and subsequence length n indicates the length of the considered substrings u . For each language, we perform a search over these two hyperparameters for set size $k = 10$ and diversity coefficient $w = 0.1$ on validation sets. We use a grid search over $n = [2, 3, 4, 5, 6, 7, 8]$ and $\lambda = [0.1, 0.3, 0.5, 0.7, 1.0]$. We choose the configuration that yields the highest (average n -gram diversity)*BLEU, using this configuration in all subsequent experiments. While there may be better performing hyperparameters under different k and w , we omit searching over the entire space to create a fairer comparison with the other decoding strategies.

Dataset	n	decay (λ)
WMT'14 En-Fr	2	0.3
WMT'19 De-En	2	0.1

Table 3: Best performing configurations found in search over string subsequence kernel parameters.

Interestingly, larger values of n did not improve performance, and were more computationally expensive; small values of n and decay λ appear to offer the best BLEU vs. n -gram diversity trade-off.

Dataset and Model Statistics We use a convolutional sequence-to-sequence model trained according to Gehring et al. (2017) on the WMT'14 En-Fr dataset.¹⁶ Data preprocessing steps, model hyperparameters and baseline performances can be found in their work. We use the pre-trained model checkpoints made available by fairseq at <https://github.com/pytorch/fairseq/tree/master/examples/translation>. We use a Transformer-based model trained according to Ng et al. (2019) on the WMT'19 De-En dataset.¹⁷ Likewise, data preprocessing steps, model hyperparameters and baseline performances can be found in Ng et al. (2019). We similarly use the pretrained model checkpoints made available by fairseq.

C Additional Results




Figure 2: n -gram diversity vs. minimum, median and maximum BLEU score for beam sizes $k = 5, 10, 20$ on WMT'19 De-En newstest using a larger range of the diversity weight w .

¹⁶available at <http://statmt.org/wmt14/translation-task.html>

¹⁷available at <http://www.statmt.org/wmt19/translation-task.html>




Figure 3: n -gram diversity vs. minimum, median and maximum BLEU score for beam sizes $k = 5, 10, 20$ on WMT'14 En-Fr newstest using a larger range of the diversity weight w .