

Conditional Poisson Stochastic Beam Search

Clara Meister  Afra Amini  Tim Vieira  Ryan Cotterell 
 ETH Zürich  Johns Hopkins University  University of Cambridge
clara.meister@inf.ethz.ch amini@student.ethz.ch
tim.vieira@gmail.com ryan.cotterell@inf.ethz.ch

Abstract

Beam search is the default decoding strategy for many sequence generation tasks in NLP. The set of approximate K -best items returned by the algorithm is a useful summary of the distribution for many applications; however, the candidates typically exhibit high overlap and may give a highly biased estimate for expectations under our model. These problems can be addressed by instead using stochastic decoding strategies. In this work, we propose a new method for turning beam search into a stochastic process: Conditional Poisson stochastic beam search. Rather than taking the maximizing set at each iteration, we sample K candidates without replacement according to the conditional Poisson sampling design. We view this as a more natural alternative to Kool et al. (2019)’s stochastic beam search (SBS). Furthermore, we show how samples generated under the CPSBS design can be used to build consistent estimators and sample diverse sets from sequence models. In our experiments, we observe CPSBS produces lower variance and more efficient estimators than SBS, even showing improvements in high entropy settings.¹

1 Introduction

Many NLP tasks require the prediction of structured outputs, such as sentences or parse trees, either during decoding or as part of a training algorithm. For today’s neural architectures, beam search (Reddy, 1977) has become the decoding algorithm of choice due to its efficiency and empirical performance (Serban et al., 2017; Edunov et al., 2018; Yang et al., 2019; Meister et al., 2020b). Beam search is a deterministic method, which invites a natural question: What is the proper *stochastic* generalization of beam search? Several recent papers have investigated this question (Kool et al., 2019, 2020; Shi et al., 2020). Here we build

¹Our codebase is publically available at <https://github.com/rycolab/cpsbs>.

Operator \ Set Size	$K = 1$	$K > 1$
	argmax	Greedy Search
sample	Ancestral Sampling	Conditional Poisson Beams

Table 1: A comparison of beam-based decoding algorithms for sequence models, by solution set size and objective. The argmax and sample variants are related through annealing: As the annealing parameter of the distribution $\tau \rightarrow 0$, sampling turns into computing an argmax (see §3).

on this line of work and introduce an alternative stochastic beam search that the authors contend is a more faithful stochasticization of the original algorithm in that it recovers standard beam search as a special case. We name our algorithm conditional Poisson stochastic beam search (CPSBS) as we draw on the conditional Poisson sampling scheme (Hájek, 1964) in its construction. The relationship between CPSBS and other common decoding strategies is displayed visually in Table 1.

At every iteration, CPSBS replaces the top- K operator in the beam search algorithm with conditional Poisson sampling, resulting in a decoding strategy that generates samples-without-replacement. Importantly, annealing our sampling distribution at each time step turns local sampling into a local top- K computation and thereby recovers beam search. We subsequently show that these samples can be used to construct a statistically consistent estimator for the expected value of an arbitrary function of the output.

In our experiments with neural machine translation models, we observe that CPSBS leads to better estimates of expected BLEU and conditional model entropy than SBS and the sum-and-sample estimator (Kool et al., 2020), distinctly outperforming Monte Carlo sampling for both small sample sizes and low temperatures. Furthermore, we find that CPSBS can be used as a diverse sampling strategy.

We take these results as confirmation that CPSBS is a useful tool in the newfound arsenal of sampling strategies for neural sequence models.

2 Beam Search

In this section, we overview the necessary background on neural sequence models and beam search in order to motivate our algorithm in §3.

Neural Sequence Models. We consider locally normalized probabilistic models over sequences \mathbf{y} :

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

where \mathbf{y} is a member of a set of well-formed outputs \mathcal{Y} . In the context of language generation models, well-formed outputs are sequences of tokens $\mathbf{y} = \langle y_1, y_2, \dots \rangle$ from a vocabulary V ; all $\mathbf{y} \in \mathcal{Y}$ begin and end with special tokens BOS and EOS, respectively. We use $\mathbf{y}_{<t}$ to represent the subsequence $\langle y_1, \dots, y_{t-1} \rangle$. In this work, we consider the setting where the maximum sequence length is upper-bounded; we denote this upper bound $T > 0$. Without loss of generality, we may condition p on an input \mathbf{x} , as is necessary for machine translation and other conditional generation tasks.

Beam Search. Beam search is a commonly used search heuristic for finding an approximate solution to the following optimization problem:

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

Its most straightforward interpretation is as a pruned version of breadth-first search, where the breadth of the search is narrowed to the top- K candidates. However, here we will present beam search in a nonstandard lens (Meister et al., 2020a, 2021) in order to emphasize the connection with our stochastic generalization in §3. Specifically, we present the algorithm as iteratively finding the highest-scoring set under a specific set function.

Under this paradigm, the initial beam Y_0 contains only the BOS token. At subsequent steps $t = 1, \dots, T$, beam search selects the K highest-scoring candidates from the set $Y_{t-1} \circ V$ that we define below:²

$$Y_{t-1} \circ V \stackrel{\text{def}}{=} \{\mathbf{y} \circ y \mid \mathbf{y} \in Y_{t-1} \text{ and } y \in V\} \quad (3)$$

²Sequences already ending in EOS are not extended by $y \in V$ and are simply added to the set “as is.”

where \circ is sequence concatenation. Those candidate sets with collectively higher probability under the model p have higher score. This process continues until all $\mathbf{y} \in Y_t$ end in EOS, or $t = T$. For notational ease, we define $B_t \stackrel{\text{def}}{=} Y_{t-1} \circ V$; throughout this paper, we will assume $|B_t| = N$ and identify the elements of $B_t = \{\mathbf{y}_{\leq t}^{(1)}, \dots, \mathbf{y}_{\leq t}^{(N)}\}$ with the integers $\{1, \dots, N\}$.

We can formulate the time-step dependent set function whose argmax beam search finds as

$$Q_t(Y_t | Y_{t-1}) \stackrel{\text{def}}{\propto} \begin{cases} \prod_{n=1}^N w_n & \text{if } |Y_t| = K \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where w_n is the weight of the n^{th} element of B_t . To recover beam search, we set our weights equal to probabilities under a model p i.e. $w_n = p(\mathbf{y}_{\leq t}^{(n)})$. Note that we leave the constraint that $Y \subseteq B_t$ implicit in Eq. (4). As should be clear from notation, this set function only assigns positive scores to subsets of $Y_{t-1} \circ V$ of size exactly K and the assigned score is proportional to the product of the probability of the candidates under the model p . Putting this all together, beam search may be viewed as the following iterative process:

$$Y_0 = \{\text{BOS}\} \quad (5)$$

$$Y_t = \operatorname{argmax}_{Y'_t \subseteq B_t} Q_t(Y'_t | Y_{t-1}) \quad (6)$$

$$\text{return } Y_T \quad (7)$$

3 Conditional Poisson Stochastic Beams

Our paper capitalizes on a very simple observation: Rather than taking its argmax, we may renormalize Eq. (4) into a distribution and sample-without-replacement a size K set at each iteration:

$$Y_0 = \{\text{BOS}\} \quad (8)$$

$$Y_t \sim Q_t(\cdot | Y_{t-1}) \quad (9)$$

$$\text{return } Y_T \quad (10)$$

This recursion corresponds to performing **conditional Poisson sampling** (CPS; Hájek 1964; see App. A for overview), a common sampling-without-replacement design (Tillé, 2006),³ at every time step. Thus we term this scheme conditional Poisson stochastic beam search. CPSBS gives us a

³A sampling design is a probability distribution over sets of samples.

probability distribution over *sets* of candidates of size K , i.e., the final beam Y_T . We denote the CPSBS distribution P and we write $Y_T \sim P$ to indicate that Y_T is the stochastic beam at the end of a sampling run. We may write $P(Y_T)$ as a marginal probability, summing over all sequences of beams that could have resulted in Y_T :⁴

$$P(Y_T) = \sum_{Y_1} \cdots \sum_{Y_{T-1}} \prod_{t=1}^T Q_t(Y_t | Y_{t-1}) \quad (11)$$

Note the structural zeros of Q_t prevent any incompatible sequence of beams. We provide a theoretical analysis of the scheme in §4 and an empirical analysis in §5.

Normalizing $Q_t(\cdot | Y_{t-1})$. At each time step t , we compute $Q_t(\cdot | Y_{t-1})$ —a distribution over subsets of size K of a base set B_t —using the CPS design. The normalizing constant for this distribution is defined as

$$Z_t \stackrel{\text{def}}{=} \sum_{\substack{Y_t \subseteq B_t, n=1 \\ |Y_t|=K}}^N \prod w_n \quad (12)$$

Despite there being exponentially many summands, we can sum over all $\binom{N}{K}$ subsets in $\mathcal{O}(NK)$ time via the following recurrence relation:⁵

$$W \binom{n}{k} = \begin{cases} 1 & \text{if } k = 0 \\ \mathbf{or } n = k \\ W \binom{n-1}{k} + w_n W \binom{n-1}{k-1} & \text{if } k \in (0, n) \\ 0 & \text{otherwise} \end{cases}$$

We give complete pseudocode in App. C. Correctness of this algorithm is shown in Kulesza and Taskar (2012). The normalizing constant can then be efficiently computed as

$$Z_t = W \binom{N}{K} \quad (13)$$

Sampling from $Q_t(\cdot | Y_{t-1})$. We can efficiently sample sets from $Q_t(\cdot | Y_{t-1})$ using the algorithm below:

- 1: $Y_t \leftarrow \emptyset \triangleright$ Initialization
- 2: **for** $n = N \dots 1$:
- 3: $k \leftarrow K - |Y_t| \triangleright$ Number of remaining elements

⁴This formulation reveals that it is wildly intractable to compute $P(Y_T)$.

⁵The reader may recognize this recurrence as the weighted generalization of Pascal’s triangle, $\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$, which is why we chose the notation $W \binom{n}{k}$.

- 4: Add the n^{th} element of B_t to Y_t with prob.

$$\frac{w_n W \binom{n-1}{k-1}}{W \binom{n}{k}}$$

- 5: **return** $Y_t \quad \triangleright$ Guaranteed to have size K

In words, the algorithm considers adding each element one at a time until K elements have been sampled. Notice that line 4 adjusts the probability of sampling item n given that $|Y_t|$ items have already been sampled, which ensures that exactly K elements are sampled at termination.

Setting w_n . The weight assigned to the n^{th} item of B_t directly affects its probability of being included in the sampled set, i.e., $\Pr(\mathbf{y}_{\leq t}^{(n)} \in Y_t)$, also termed an item’s **inclusion probability**. In this paper, we write $\pi_{Q_t}(\mathbf{y}_{\leq t}^{(n)} | Y_{t-1})$ to denote the inclusion probability under the distribution $Q_t(\cdot | Y_{t-1})$, defined as:

$$\pi_{Q_t}(\mathbf{y}_{\leq t}^{(n)} | Y_{t-1}) \stackrel{\text{def}}{=} \sum_{Y_t} Q_t(Y_t | Y_{t-1}) \mathbb{1}\{\mathbf{y}_{\leq t}^{(n)} \in Y_t\} \quad (14)$$

One strategy is to choose w_n at time step t such that $\pi_{Q_t}(\mathbf{y}_{\leq t}^{(n)} | Y_{t-1}) \approx p(\mathbf{y}_{\leq t}^{(n)})$. This choice recovers beam search when we *anneal* our chosen weights $w_n \mapsto w_n^{1/\tau}$: as the temperature parameter $\tau \rightarrow 0$, the CP distribution will assign probability 1 to the set containing the top- K elements.⁶

Finding w_n ’s that result in pre-specified inclusion probabilities is possible, but it requires solving a numerical optimization problem (Aires, 1999; Grafström, 2009). Further, in CPSBS, we will be sampling from a different distribution at each time step and it would be quite slow to solve the numerical optimization problem each iteration. Luckily, the choice of $w_n = p(\mathbf{y}_{\leq t}^{(n)}) / (1 - p(\mathbf{y}_{\leq t}^{(n)}))$ yields a good approximation to the target inclusion probabilities in both theory and practice (Hájek, 1981; Bondesson et al., 2006; Aires, 1999).

4 Statistical Estimation with Conditional Poisson Stochastic Beam Search

In this section, we discuss statistical estimation with CPSBS samples. To that end, we construct two estimators with different properties. However,

⁶In the event of ties, annealed CP will converge to a distribution that breaks ties uniformly at random.

only the second estimator provides good performance in practice, which is discussed later in §5.

4.1 The Horvitz–Thompson Estimator

We build upon the Horvitz–Thompson (HT) estimator (Horvitz and Thompson, 1952), which is a common technique for estimation from sampling-without-replacement (SWOR) schemes.

Let be $f : \mathcal{Y} \rightarrow \mathbb{R}^d$, we seek to approximate its expected value under p :

$$\mathbb{E}_{\mathbf{y} \sim p} [f(\mathbf{y})] = \sum_{\mathbf{y} \in \mathcal{Y}} p(\mathbf{y}) f(\mathbf{y}) \quad (15)$$

The Monte Carlo estimator of the above quantity is

$$G_{\text{MC}} \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M f(\mathbf{y}^{(m)}) \quad (16)$$

where $\mathbf{y}^{(m)} \stackrel{\text{i.i.d.}}{\sim} p$. However, in the special case of sampling from a finite population—which is extremely common in NLP—it can be very wasteful. For example, if a distribution is very peaked, it will sample the same item repeatedly; this could lead to inaccurate approximations for some f . As a consequence, the mean square error (MSE) of the estimator with respect to $\mathbb{E}_{\mathbf{y} \sim p} [f(\mathbf{y})]$ can be quite high for small M . Indeed, we see this empirically in §5.1.

Taking samples without replacement allows us to cover more of the support of p in our estimate of $\mathbb{E}_{\mathbf{y} \sim p} [f(\mathbf{y})]$. However, we must take into account that our samples are no longer independent, i.e., $\mathbf{y} \not\sim p$. We now define the HT estimator, using notation specifically for the case of CPSBS where $Y_T \sim P$:

$$G_{\text{HT}} \stackrel{\text{def}}{=} \sum_{\mathbf{y} \in Y_T} \frac{p(\mathbf{y})}{\pi_P(\mathbf{y})} f(\mathbf{y}) \quad (17)$$

where CPSBS’s **inclusion probability** is

$$\begin{aligned} \pi_P(\mathbf{y}) &= \sum_{Y_T} P(Y_T) \mathbb{1}\{\mathbf{y} \in Y_T\} \\ &= \sum_{Y_1} \cdots \sum_{Y_T} \prod_{t=1}^T Q_t(Y_t | Y_{t-1}) \mathbb{1}\{\mathbf{y}_{\leq t} \in Y_t\} \end{aligned} \quad (18)$$

i.e., the probability of sampling a set $Y_T \sim P$ that contains the element \mathbf{y} . In Eq. (17), the distribution π_P may be viewed as a proposal distribution in the sense of importance sampling (Owen, 2013) and $1/\pi_P(\mathbf{y})$ as the corresponding importance weight corrections. If we can exactly compute π_P , then

the HT estimator is unbiased⁷ (see App. B.1 for proof). However, the summation in Eq. (18) is intractable so we resort to estimation.

4.2 Estimating Inclusion Probabilities

In this section, we develop statistical estimators of the inclusion probabilities under conditional Poisson stochastic beam search. An important caveat: the analysis in this section only applies to the estimators of the inverse inclusion probabilities themselves. Further analysis may be undertaken to analyze the variance of the Horvitz–Thompson estimators that make use of these estimators.

4.2.1 Naïve Monte Carlo

It is not straightforward to estimate the reciprocal inclusion probabilities. Thus, we attempt to estimate the inclusion probabilities directly and take the reciprocal of this estimator. This strategy leads to a consistent, but biased, estimator.⁸ One obvious way to derive an inclusion probability estimator is the Monte Carlo estimator:

$$\widehat{\pi}_P^{\text{MC}}(\mathbf{y}) \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M \mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\} \quad (19)$$

where $Y^{(m)} \sim P$.

Proposition 4.1. *Eq. 19 has the following two properties:*

i) $\widehat{\pi}_P^{\text{MC}}$ is an unbiased estimator of π_P and

$$\mathbb{V}[\widehat{\pi}_P^{\text{MC}}] = \frac{1}{M} (\pi_P(\mathbf{y}) - \pi_P(\mathbf{y})^2) \quad (20)$$

ii) $1/\widehat{\pi}_P^{\text{MC}}$ is a consistent estimator of $1/\pi_P$ with asymptotic variance

$$\mathbb{V}_a \left[\frac{1}{\widehat{\pi}_P^{\text{MC}}(\mathbf{y})} \right] = \frac{1}{M} \left(\frac{1}{\pi_P(\mathbf{y})^3} - \frac{1}{\pi_P(\mathbf{y})^2} \right) \quad (21)$$

Here \mathbb{V}_a denotes the **asymptotic variance**, which is the variance after the number of samples M is large enough such that the central limit theorem has kicked in (Bickel and Doksum, 2015).

Proof. Proof given in App. B.2. ■

⁷Note that it is common to normalize Eq. (17) by the sum of importance weights, i.e., divide G_{HT} by the sum $\sum_{\mathbf{y} \in \tilde{Y}_T} \pi_{Q_t}(\mathbf{y})$. While this leads to a biased estimator, it can significantly reduce variance, which is often worthwhile.

⁸Since by Jensen’s inequality $\mathbb{E}[1/X] \geq 1/\mathbb{E}[X]$ for $X \in \mathbb{R}_+$, the reciprocal of an unbiased estimate of $\pi_P(\mathbf{y})$ is not an unbiased estimate of $1/\pi_P(\mathbf{y})$

Qualitatively, what this result tells us is that if we are asking about the inverse inclusion probability of a candidate with a low inclusion probability, our estimator may have very high variance. Indeed, it is unlikely that we could derive an estimator without this qualitative property due to the presence of the inverse. Moreover, the estimator given in Eq. (19) is not of practical use: If we are interested in the inverse inclusion probability of a specific candidate \mathbf{y} , then we may have to sample a very large number of beams until we eventually sample one that actually contains \mathbf{y} . In practice, what this means is that our estimate of the inclusion probability for a rare \mathbf{y} will often be zero, which we cannot invert.⁹ Instead, we pursue an importance sampling strategy for estimating $\pi_P(\mathbf{y})$, which we outline in the next section.

4.2.2 Importance Sampling

We now turn to an inclusion probability estimator that is based on importance sampling. Recall from Eq. (18) that the inclusion probability for \mathbf{y} is a massive summation over sequences of possible beams Y_1, \dots, Y_T that could have generated \mathbf{y} . Rather than computing the sum, we will estimate the sum through taking samples. Our procedure starts by generating **hindsight samples** $\tilde{Y}_1, \dots, \tilde{Y}_T$ from the following proposal distribution that is conditioned on \mathbf{y} :

$$\tilde{Q}_t(\tilde{Y}_t | \tilde{Y}_{t-1}, \mathbf{y}) \stackrel{\text{def}}{=} \frac{Q_t(\tilde{Y}_t | \tilde{Y}_{t-1})}{\pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1})} \quad (22)$$

In words, \tilde{Q}_t is Q_t conditioned on its sets Y_t containing the prefix $\mathbf{y}_{\leq t}$ (thus it is always the case that $\mathbf{y}_{\leq t} \in \tilde{Y}_t$).¹⁰ For brevity, we omit an explicit notational dependence of \tilde{Y} and \tilde{Q}_t on \mathbf{y} .

Lemma 4.1. *The joint proposal distribution¹¹ $\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) \stackrel{\text{def}}{=} \prod_{t=1}^T \tilde{Q}_t(\tilde{Y}_t | \tilde{Y}_{t-1})$ may be expressed in terms of P as follows:*

$$\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) = \frac{P(\tilde{Y}_1, \dots, \tilde{Y}_T)}{\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1})} \quad (23)$$

where we define $P(\tilde{Y}_1, \dots, \tilde{Y}_T) \stackrel{\text{def}}{=} \prod_{t=1}^T Q_t(\tilde{Y}_t | \tilde{Y}_{t-1})$ as the joint probability of the beams under

⁹One solution would be to smooth our estimates of the inclusion probabilities, adding a small ε to ensure that we do not divide by zero, but the authors find our next approach to be more methodologically sound.

¹⁰This proposal distribution can be realized through a minor modification of our algorithm in §3, where $w(\mathbf{y})$ corresponding to $\mathbf{y}_{\leq t}^{(n)}$ is placed at the beginning and added to Y_t deterministically.

¹¹We have omitted dependency on \mathbf{y} for brevity.

the original distributions Q_t . We omit that both P and \tilde{P} are conditioned on Y_0 .

Proof. See App. B.2. ■

In terms of computation, Eq. (22) makes use of the fact that the per-time-step inclusion probability $\pi_{Q_t}(\mathbf{y}_{\leq t})$ for a given Q_t can be computed efficiently with dynamic programming using the following identity:

$$\begin{aligned} \pi_{Q_t}(\mathbf{y}_{\leq t}^{(n)} | Y_{t-1}) &\stackrel{\text{def}}{=} \sum_Y Q_t(Y_t) \mathbb{1}\{\mathbf{y}_{\leq t}^{(n)} \in Y_t\} \\ &= \frac{w_n}{Z} \frac{\partial Z}{\partial w_n} \end{aligned} \quad (24)$$

For completeness, we give pseudocode in App. C. Given samples $\tilde{Y}_T^{(m)} \sim \tilde{P}$ for \tilde{P} defined in Eq. (22) with respect to a given \mathbf{y} , we propose the following unbiased estimator of inclusion probabilities:

$$\hat{\pi}_P^{\text{is}}(\mathbf{y}) \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1}^{(m)}) \quad (25)$$

where $\mathbf{y}_{\leq t}$ is a prefix of \mathbf{y} . One simple derivation of Eq. (25) is as an importance sampler. We start with the estimator given in Eq. (19) and perform the standard algebraic manipulations witnessed in importance sampling:

$$\begin{aligned} &\sum_{Y_T} P(Y_T) \mathbb{1}\{\mathbf{y} \in Y_T\} \quad (26) \\ &= \sum_{Y_T} \dots \sum_{Y_1} P(Y_1, \dots, Y_T) \mathbb{1}\{\mathbf{y} \in Y_T\} \\ &= \sum_{\tilde{Y}_T} \dots \sum_{\tilde{Y}_1} P(\tilde{Y}_1, \dots, \tilde{Y}_T) \frac{\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T)}{\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T)} \\ &= \sum_{\tilde{Y}_T} \dots \sum_{\tilde{Y}_1} \tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) \frac{P(\tilde{Y}_1, \dots, \tilde{Y}_T)}{\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T)} \\ &\stackrel{(i)}{=} \sum_{\tilde{Y}_T} \dots \sum_{\tilde{Y}_1} \tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1}) \end{aligned}$$

where equality (i) above follows from Lemma 4.1. This derivation serves as a simple proof that Eq. (25) inherits unbiasedness from Eq. (19)

Proposition 4.2. *Eq. (25) has the following two properties:*

- i) $\hat{\pi}_P^{\text{is}}$ is an unbiased estimator of π_P ;
- ii) The estimator of the inverse inclusion probabilities $1/\hat{\pi}_P^{\text{is}}(\mathbf{y})$ is consistent with the following upper bound on the asymptotic variance:

$$\mathbb{V}_a \left[\frac{1}{\widehat{\pi}_p^{\text{IS}}(\mathbf{y})} \right] \leq \frac{1}{M} \frac{r-1}{\pi_p(\mathbf{y})^2} \quad (27)$$

where we assume that the following bound:

$$\frac{\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \widetilde{Y}_{t-1})}{\pi_p(\mathbf{y})} \leq r \quad (28)$$

holds for all $\widetilde{Y}_1, \dots, \widetilde{Y}_T$.

Proof. Proof given in App. B.2. ■

Proposition 4.2 tells us that we can use Eq. (25) to construct a consistent estimator of the inverse inclusion probabilities. Moreover, assuming $\Pr(\mathbf{y} \in Y_T) > 0$, then we have that the importance sampling yields an estimate $\widehat{\pi}_p^{\text{IS}}(\mathbf{y}) > 0$, unlike the Monte Carlo estimator $\widehat{\pi}_p^{\text{MC}}(\mathbf{y})$. We further see that, to the extent that $\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \widetilde{Y}_{t-1})$ approximates $\pi_p(\mathbf{y})$, then we may expect the variance of Eq. (25) to be small—specifically in comparison to the naïve Monte Carlo estimator in Eq. (19)—which is often the case for estimators built using importance sampling techniques when a proposal distribution is chosen judiciously (Rubinstein and Kroese, 2016). Thus, given our estimator in Eq. (25), we can now construct a practically useful estimator for $\mathbb{E}_{\mathbf{y} \sim p}[f(\mathbf{y})]$ using the HT estimator in Eq. (17). In the next section, we observe that this estimator is quite efficient in the sequence model setting.

5 Experiments

We repeat the analyses performed by Kool et al. (2019), running experiments on neural machine translation (NMT) models; for reproducibility, we use the pretrained Transformer model for WMT’14 (Bojar et al., 2014) English–French made available by fairseq¹² (Ott et al., 2019). We evaluate on the En-Fr newstest2014 set, containing 3003 sentences. Further details can be found in App. D. Our implementation of CPSBS modifies the beam search algorithm from the fairseq library. We additionally consider the beam search, stochastic beam search, diverse beam search, and ancestral sampling algorithms available in fairseq.

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

5.1 Statistical Estimators for Language Generation Models

Estimators have a large number of applications in machine learning. For example, the REINFORCE algorithm (Williams, 1992) constructs an estimator for the value of the score function; minimum-Bayes risk decoding (Kumar and Byrne, 2004) uses an estimate of risk in its optimization problem. In this section, we compare estimators for sentence-level BLEU score and conditional model entropy for NMT models. Notably, NMT models that are trained to minimize cross-entropy with the empirical distribution¹³ are not peaky distributions (Ott et al., 2018a; Eikema and Aziz, 2020); thus, standard estimation techniques, e.g., Monte Carlo, should generally provide good results. However, we can vary the annealing parameter of our model in order to observe the behavior of our estimator with both high- and low-entropy distributions, making this a comprehensive case study. Here the annealed model distribution is computed as

$$p_\tau(y_t | \mathbf{y}_{<t}) \propto p(y_t | \mathbf{y}_{<t})^{\frac{1}{\tau}} \quad (29)$$

where we should expect a standard Monte Carlo estimator to provide good results at τ close to 1 when p is naturally high entropy. We test our estimator in this setting so as to give a comparison in a competitive setting. Specifically, we assess the performance of our estimator of $\mathbb{E}_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{x})}[f(\mathbf{y})]$ given in Eq. (17)—using inclusion probability estimates from Eq. (25) with $M = 1$ and with importance weight normalization—in comparison to three other approaches: Monte Carlo (MC) sampling, the sum-and-sample (SAS) estimator, and stochastic beam search (SBS).

Monte Carlo. Under the Monte Carlo sampling scheme with sample size K , we estimate the expected value of f under our model using Eq. (16) with a sample $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(K)} \stackrel{\text{i.i.d.}}{\sim} p$.

Sum and Sample. The sum-and-sample estimator (Botev et al., 2017; Liu et al., 2019; Kool et al., 2020) is an unbiased estimator that takes as input a deterministically chosen set Y of size $K - 1$ and samples an additional \mathbf{y}' from the remaining elements, $\text{supp}(p) \setminus Y$, where we obtain the set Y using beam search in our experiments. Formally,

¹³Label-smoothing (Szegedy et al., 2016) is typically also employed, which leads to even higher entropy distributions.

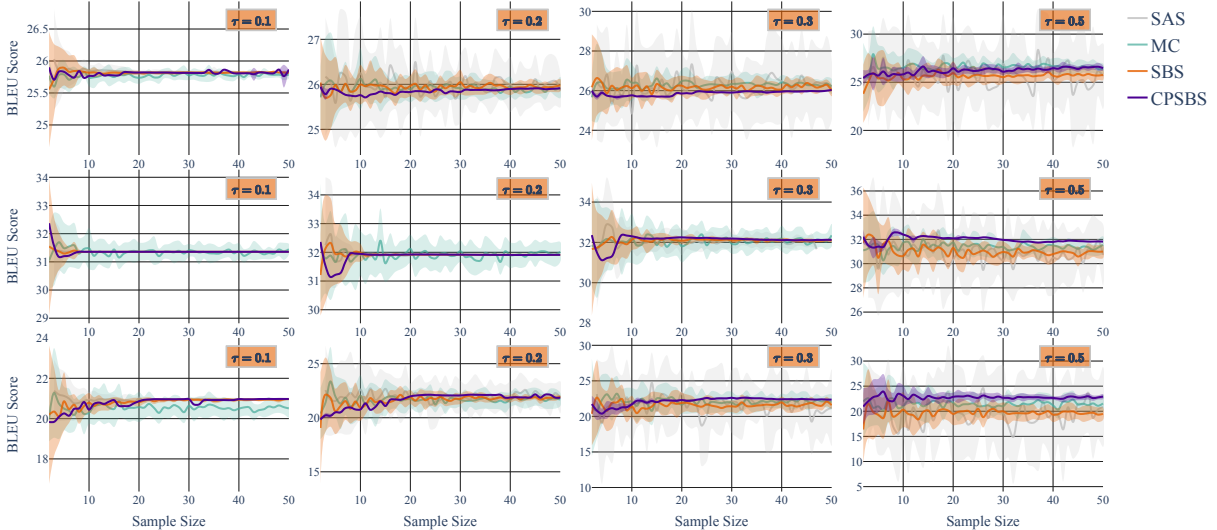


Figure 1: BLEU score estimates for three different sentences using estimators for respective decoding methods. τ indicates scaling temperature; τ values and sentences are chosen to mimic (Kool et al., 2019).

the SAS estimator can be written as:

$$G_{\text{SAS}} \stackrel{\text{def}}{=} \sum_{k=1}^{K-1} p(\mathbf{y}^{(k)}) f(\mathbf{y}^{(k)}) + \left(1 - \sum_{k=1}^{K-1} p(\mathbf{y}^{(k)})\right) f(\mathbf{y}')$$
 (30)

Stochastic Beam Search. Stochastic beam search (Kool et al., 2019, 2020) is a SWOR algorithm likewise built on top of beam search. The algorithm makes use of truncated Gumbel random variables at each iteration, resulting in a sampling design equivalent to performing the Gumbel-top- k trick (Vieira, 2014) on the distribution p . Estimators built using SBS likewise follow the Horvitz–Thompson scheme of Eq. (17); we refer the reader to the original work for inclusion probability computations. They suggest normalizing the estimator by the sum of sample inclusion probabilities to help reduce variance; we therefore likewise perform this normalization in our experiments.

To assess the error of our estimator, we compute its root MSE (RMSE) with respect to a baseline result. While computing the exact value of an expectation is typically infeasible in the sequence model setting, we can average our (unbiased) MC estimator in Eq. (16) over multiple runs to create a good baseline. Specifically, we compute our MC estimator 50 times for a large sample size ($K = 200$); variances are reported in App. D.

¹³We refer the reader to the original work (Kool et al., 2019) for equations for inclusion probability estimates.

Probabilistic models for language generation typically have large vocabularies. In this setting, the computation of Eq. (12) is inefficient due to the large number of items in the set that are assigned very small probability under the model. We experiment with truncating this distribution such that the set of possible extensions of a sequence consist only of the highest probability tokens within the core $n\%$ of probability mass (0.99 in our experiments), similar to the process in nucleus sampling (Holtzman et al., 2020). We compare this approach to the original algorithm design in App. D and see that empirically, results are virtually unchanged; the following results use this method. We also compare the decoding time of different sampling methods in Fig. 7.

BLEU Score Estimation. BLEU (Papineni et al., 2002) is a widely used automatic evaluation metric for the quality of machine-generated translations. Estimates of BLEU score are used in minimum risk training (Shen et al., 2016) and reinforcement learning-based approaches (Ranzato et al., 2016) to machine translation. As such, accurate and low-variance estimates are critical for the algorithms’ performance. Formally, we estimate the expected value of $f(\mathbf{y}) = \text{BLEU}(\mathbf{x}, \mathbf{y})$, whose dependence on \mathbf{x} we leave implicit, under our NMT model p for reference translation \mathbf{x} . For comparison, we use the same sentences and similar annealing temperatures¹⁴ τ evaluated by Kool et al. (2019). We repeat

¹⁴Results for $\tau = 0.05$ converged rapidly for all estimators, thus not providing an interesting comparison.

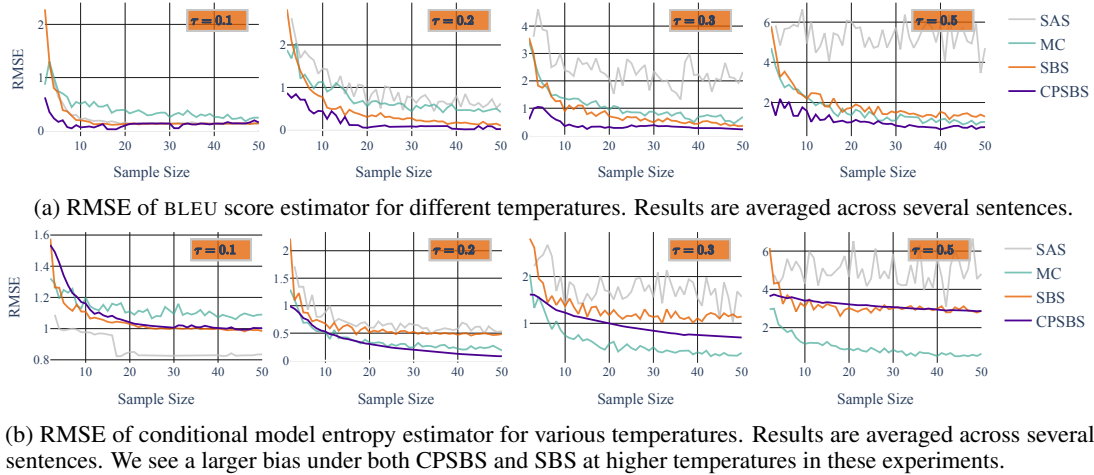


Figure 2: RMSE evaluations.

the sampling 20 times and plot the value and standard deviation (indicated by shaded region) of different estimators in Fig. 1. From Fig. 1, we can see that CPSBS has lower variance than our baseline estimators across all temperatures and data points.¹⁵ Especially in the low temperature setting, our estimator converges rapidly with minor deviation from the exact values even for small sample sizes. Additionally, in Fig. 2a we see that the RMSE is typically quite low except at higher temperatures. In such cases, we observe the effects of biasedness, similar to Kool et al. (2019)’s observations.

Conditional Entropy Estimation. We perform similar experiments for estimates of a model’s conditional entropy, i.e., $f(\mathbf{y}) = -\log p(\mathbf{y} \mid \mathbf{x})$, whose dependence on \mathbf{x} we again leave implicit. We show results in Fig. 2b, with plots of the value in App. D since results are quite similar to Fig. 1. We see further confirmation that our estimator built on CPSBS is generally quite efficient.

5.2 Diverse Sampling

We show how CPSBS can be used as a diverse set sampling design for language generation models. We generate a sample of translations $Y_T \sim P$, i.e., according to the CPSBS scheme, where weights are set as $w_n = p(\mathbf{y}_{\leq t}^{(n)}) / (1 - p(\mathbf{y}_{\leq t}^{(n)}))$ at each time step, as recommended in §3. In Fig. 3, we show the trade-off between minimum, average and maximum sentence-level BLEU score (as a quality measure) and n -gram diversity, where we define n -gram diversity D as the average fraction of unique

¹⁵The sampling distribution at $n = 1$ is not the same across strategies, hence the difference in variances even at $n = 1$.

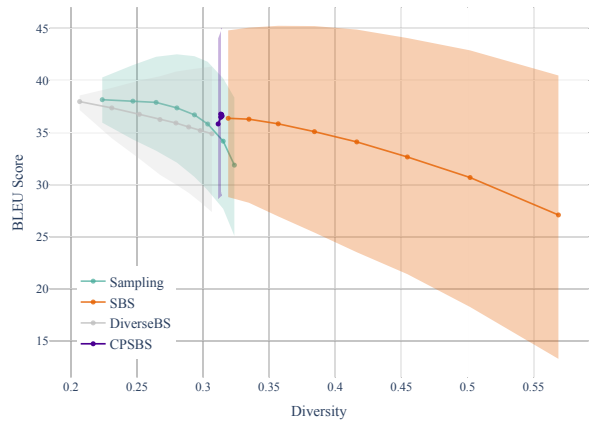


Figure 3: Average (\pm min and max) BLEU score versus diversity for sample size $k = 5$. Points correspond to different annealing temperatures $\{0.1, \dots, 0.8\}$. Results for $k = 10, 20$ show very similar trends.

vs. total n -grams for $n = 1, 2, 3, 4$ in a sentence:

$$D = \sum_{n=1}^4 \frac{\#\text{unique } n\text{-grams in } K \text{ strings}}{\# n\text{-grams in } K \text{ strings}} \quad (31)$$

Metrics are averaged across the corpus. We follow the experimental setup of Kool et al. (2019), using the newstest2014 dataset and comparing three different decoding methods: SBS, diverse beam search (DiverseBS; Vijayakumar et al., 2018) and ancestral sampling. As in their experiments, we vary the annealing temperature in the range $\{0.1, 0.2, \dots, 0.8\}$ as a means of encouraging diversity; for DiverseBS we instead vary the *strength parameter* in the same range. Interestingly, we see that temperature has virtually no effect on the diversity of the set of results returned by CPSBS. Despite this artifact, for which the authors have not

found a theoretical justification,¹⁶ the set returned by CPSBS is still overall more diverse (position on x -axis) than results returned by DiverseBS and reflect better min, max, and average BLEU in comparison to random sampling. SBS provides a better spectrum for the diversity and BLEU tradeoff; we thus recommend SBS when diverse sets are desired.

6 Conclusion

In this work, we present conditional Poisson stochastic beam search, a sampling-without-replacement strategy for sequence models. Through a simple modification to beam search, we turn this mainstay decoding algorithm into a stochastic process. We derive a low-variance, consistent estimator of inclusion probabilities under this scheme; we then present a general framework for using CPSBS to construct statistical estimators for expectations under sequence models. In our experiments, we observe a reduction in mean square error, and an increase in sample efficiency, when using our estimator in comparison to several baselines, showing the benefits of CPSBS.

References

- Nibia Aires. 1999. [Algorithms to find exact inclusion probabilities for conditional Poisson sampling and Pareto \$\pi\$ ps sampling designs](#). *Methodology and Computing in Applied Probability*, 1(4):457–469.
- Peter J. Bickel and Kjell A. Doksum. 2015. *Mathematical Statistics: Basic ideas and Selected Topics*. CRC Press.
- 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, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Lennart Bondesson, Imbi Traat, and Anders Lundqvist. 2006. [Pareto sampling versus Sampford and conditional Poisson sampling](#). *Scandinavian Journal of Statistics*, 33(4):699–720.
- Aleksandar Botev, Bowen Zheng, and David Barber. 2017. [Complementary sum sampling for likelihood](#)

[approximation in large scale classification](#). In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.

- Martin Bücker, George Corliss, Paul Hovland, Uwe Naumann, and Boyana Norris. 2006. *Automatic Differentiation: Applications, Theory, and Implementations (Lecture Notes in Computational Science and Engineering)*. Springer-Verlag, Berlin, Heidelberg.
- 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.
- Bryan Eikema and Wilker Aziz. 2020. [Is MAP decoding all you need? The inadequacy of the mode in neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4506–4520, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Anton Grafström. 2009. [Non-rejective implementations of the Sampford sampling design](#). *Journal of Statistical Planning and Inference*, 139(6):2111–2114.
- J. Hájek. 1981. *Sampling from a Finite Population*. Statistics Series. Marcel Dekker.
- Jaroslav Hájek. 1964. [Asymptotic theory of rejective sampling with varying probabilities from a finite population](#). *Annals of Mathematical Statistics*, 35(4):1491–1523.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *Proceedings of the 7th International Conference on Learning Representations*.
- D. G. Horvitz and D. J. Thompson. 1952. [A generalization of sampling without replacement from a finite universe](#). *Journal of the American Statistical Association*, 47(260):663–685.
- 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 36th International Conference on Machine Learning*.
- Wouter Kool, Herke van Hoof, and Max Welling. 2020. [Estimating gradients for discrete random variables by sampling without replacement](#). In *Proceedings of the 8th International Conference on Learning Representations*.
- Alex Kulesza and Ben Taskar. 2012. [Determinantal point processes for machine learning](#). *Foundations and Trends® in Machine Learning*, 5(2–3):123–286.

¹⁶While scaling sampling weights by a constant should not change the distribution P , an exp transformation of weights—which is the computation performed by temperature annealing—should.

- Shankar Kumar and William Byrne. 2004. [Minimum Bayes-risk decoding for statistical machine translation](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 169–176, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Runjing Liu, Jeffrey Regier, Nilesh Tripuraneni, Michael Jordan, and Jon McAuliffe. 2019. [Rao-Blackwellized stochastic gradients for discrete distributions](#). In *Proceedings of the 36th International Conference on Machine Learning*.
- Clara Meister, Ryan Cotterell, and Tim Vieira. 2020a. [If beam search is the answer, what was the question?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2173–2185, Online. Association for Computational Linguistics.
- Clara Meister, Martina Forster, and Ryan Cotterell. 2021. [Determinantal beam search](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6551–6562, Online. Association for Computational Linguistics.
- Clara Meister, Tim Vieira, and Ryan Cotterell. 2020b. [Best-first beam search](#). *Transactions of the Association for Computational Linguistics*, 8:795–809.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018a. [Analyzing uncertainty in neural machine translation](#). In *Proceedings of the 35th International Conference on Machine Learning*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018b. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Art B. Owen. 2013. *Monte Carlo theory, methods and examples*.
- 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.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#). In *Proceedings of the 4th International Conference on Learning Representations*.
- Raj Reddy. 1977. [Speech understanding systems: A summary of results of the five-year research effort at Carnegie Mellon University](#). Technical report.
- Reuven Y. Rubinstein and Dirk P. Kroese. 2016. *Simulation and the Monte Carlo Method*, 3rd edition. Wiley Publishing.
- 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*.
- 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.
- 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*.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. 2016. [Rethinking the inception architecture for computer vision](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, Los Alamitos, CA, USA. IEEE Computer Society.
- Yves Tillé. 2006. *Sampling Algorithms*. Springer.
- Tim Vieira. 2014. [Gumbel-max trick and weighted reservoir sampling](#).
- 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 *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 1.
- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Machine Learning*, 8(3–4):229–256.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [XLNet: Generalized autoregressive pre-training 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*. Curran Associates, Inc.

A Conditional Poisson Sampling

Here we provide a brief overview of the sampling design at the core of CPSBS: conditional Poisson sampling. We consider a base set B where $|B| = N$ and we map the elements of $B = \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}\}$ to the integers $\{1, \dots, N\}$. As a warm up, we first consider (unconditional) Poisson sampling, also known as a Bernoulli point process. To sample a subset $Y \subseteq B$, we do as follows: for each element $\mathbf{y} \in B$, we flip a coin where the odds of heads is $w(\mathbf{y})$. Then, we simply take Y to be the subset of elements whose coin flips were heads. However, this sampling scheme clearly does not guarantee a sample of K items, which can cause problems in our application; sampling more than K items would make the stochastic beam search process inefficient while sampling fewer than K —or even 0—items may not leave us with a large enough set at the end of our iterative process.

If instead, we *condition* on the sets always having a prescribed size K , i.e., reject samples where $|Y| \neq K$, we arrive at the *conditional* Poisson process. Formally, the conditional Poisson distribution is defined over $Y \subseteq B$ as follows,

$$Q(Y) \stackrel{\text{def}}{\propto} \begin{cases} \prod_{\mathbf{y} \in Y} w(\mathbf{y}) & \text{if } |Y| = K \\ 0 & \text{otherwise} \end{cases} \quad (32)$$

By analyzing Eq. (32), we can see that sets with the largest product of weights are the most likely to be sampled; further, this distribution is invariant to rescaling of weights due to the size requirement. This is similar to the conditions under which beam search chooses the set of K largest weight, i.e., highest scoring, elements. Indeed, we note the extreme similarity between Eq. (4) and Eq. (32), the only difference being a dependence on a prior set. However, unlike beam search, sets with a lower weight product now have the possibility of being chosen.

B Proofs

B.1 Unbiasedness of the Horvitz–Thompson Estimator

Proposition B.1. *Given a SWOR design Q over the set $B = \{1, \dots, N\}$ with inclusion probabilities $\pi(n)$, the Horvitz–Thompson estimator (Eq. (17)) gives us an unbiased estimator of $\mathbb{E}_{n \sim p} f(n)$, where $f : B \rightarrow \mathbb{R}^d$ is a function whose expectation under p we seek to approximate.*

Proof.

$$\mathbb{E}_{Y \sim Q} [G_{\text{HT}}] = \mathbb{E}_{Y \sim Q} \sum_{n=1}^N \frac{p(n)}{\pi(n)} f(n) \quad (33a)$$

$$= \mathbb{E}_{Y \sim Q} \sum_{n \in B} \frac{p(n)}{\pi(n)} \mathbb{1}\{n \in Y\} f(n) \quad (33b)$$

$$= \sum_{n \in B} \frac{p(n)}{\pi(n)} f(n) \mathbb{E}_{Y \sim Q} \mathbb{1}\{n \in Y\} \quad (33c)$$

$$= \sum_{n \in B} \frac{p(n)}{\pi(n)} f(n) \pi(n) \quad (33d)$$

$$= \sum_{n \in B} p(n) f(n) \quad (33e)$$

$$= \mathbb{E}_{n \sim p} f(n) \quad (33f)$$

■

B.2 Proofs of Expected Values and Variances of Inclusion Probability Estimators

Lemma 4.1. *The joint proposal distribution¹⁷ $\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) \stackrel{\text{def}}{=} \prod_{t=1}^T \tilde{Q}_t(\tilde{Y}_t | \tilde{Y}_{t-1})$ may be expressed in terms of P as follows:*

$$\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) = \frac{P(\tilde{Y}_1, \dots, \tilde{Y}_T)}{\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1})} \quad (23)$$

where we define $P(\tilde{Y}_1, \dots, \tilde{Y}_T) \stackrel{\text{def}}{=} \prod_{t=1}^T Q_t(\tilde{Y}_t | \tilde{Y}_{t-1})$ as the joint probability of the beams under the original distributions Q_t . We omit that both P and \tilde{P} are conditioned on Y_0 .

Proof. Consider the probability of sampling $\tilde{Y}_1, \dots, \tilde{Y}_T$ according to \tilde{P} . Algebraic manipulation reveals:

$$\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T) = \frac{Q_t(\tilde{Y}_1 | Y_0)}{\pi_{Q_t}(\mathbf{y}_{\leq 1} | Y_0)} \dots \frac{Q_t(\tilde{Y}_T | \tilde{Y}_{T-1})}{\pi_{Q_t}(\mathbf{y}_{\leq T} | \tilde{Y}_{T-1})} \quad (34a)$$

$$= \frac{P(\tilde{Y}_1, \dots, \tilde{Y}_T)}{\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1})} \quad (34b)$$

which proves the identity. ■

Proposition 4.1. *Eq. 19 has the following two properties:*

i) $\hat{\pi}_P^{\text{mc}}$ is an unbiased estimator of π_P and

$$\mathbb{V}[\hat{\pi}_P^{\text{mc}}] = \frac{1}{M} (\pi_P(\mathbf{y}) - \pi_P(\mathbf{y})^2) \quad (20)$$

ii) $1/\hat{\pi}_P^{\text{mc}}$ is a consistent estimator of $1/\pi_P$ with asymptotic variance

$$\mathbb{V}_a \left[\frac{1}{\hat{\pi}_P^{\text{mc}}(\mathbf{y})} \right] = \frac{1}{M} \left(\frac{1}{\pi_P(\mathbf{y})^3} - \frac{1}{\pi_P(\mathbf{y})^2} \right) \quad (21)$$

Here \mathbb{V}_a denotes the **asymptotic variance**, which is the variance after the number of samples M is large enough such that the central limit theorem has kicked in ([Bickel and Docksum, 2015](#)).

Proof. i) The estimator is easily shown to be unbiased:

$$\mathbb{E} \hat{\pi}_P^{\text{mc}}(\mathbf{y}) \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M \mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\} = \pi_P(\mathbf{y}) \quad (35)$$

and its variance may be derived as follows:

$$\mathbb{V}[\hat{\pi}_P^{\text{mc}}(\mathbf{y})] \stackrel{\text{def}}{=} \mathbb{V} \left[\frac{1}{M} \sum_{m=1}^M \mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\} \right] \quad (36a)$$

$$= \frac{1}{M} \mathbb{V} \left[\mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\} \right] \quad (36b)$$

$$= \frac{1}{M} \left(\mathbb{E} \left(\mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\}^2 \right) - \mathbb{E} \left(\mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\} \right)^2 \right) \quad (36c)$$

$$= \frac{1}{M} (\pi_P(\mathbf{y}) - \pi_P(\mathbf{y})^2) \quad (36d)$$

ii) By the strong law of large numbers, we have

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\} = \pi_P(\mathbf{y}) \quad (37)$$

¹⁷We have omitted dependency on \mathbf{y} for brevity.

Since $1/x$ is continuous, we may appeal to the continuous mapping theorem to achieve consistency:

$$\lim_{M \rightarrow \infty} \frac{1}{\frac{1}{M} \sum_{m=1}^M \mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\}} = \frac{1}{\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \mathbb{1}\{\mathbf{y} \in Y_T^{(m)}\}} = \frac{1}{\pi_P(\mathbf{y})} \quad (38)$$

We can compute the asymptotic variance by the delta rule:

$$\mathbb{V}_a \left[\frac{1}{\widehat{\pi}_P^{\text{mc}}(\mathbf{y})} \right] = \frac{1}{M} \frac{\mathbb{V}[\widehat{\pi}_P^{\text{is}}(\mathbf{y})]}{\pi_P(\mathbf{y})^4} \quad (\text{apply the delta rule}) \quad (39a)$$

$$= \frac{1}{M} \frac{\pi_P(\mathbf{y}) - \pi_P(\mathbf{y})^2}{\pi_P(\mathbf{y})^4} \quad (\text{plugging in the variance computed above}) \quad (39b)$$

$$= \frac{1}{M} \left(\frac{1}{\pi_P(\mathbf{y})^3} - \frac{1}{\pi_P(\mathbf{y})^2} \right) \quad (39c)$$

■

Proposition 4.2. *Eq. (25) has the following two properties:*

- i) $\widehat{\pi}_P^{\text{is}}$ is an unbiased estimator of π_P ;
- ii) *The estimator of the inverse inclusion probabilities $1/\widehat{\pi}_P^{\text{is}}(\mathbf{y})$ is consistent with the following upper bound on the asymptotic variance:*

$$\mathbb{V}_a \left[\frac{1}{\widehat{\pi}_P^{\text{is}}(\mathbf{y})} \right] \leq \frac{1}{M} \frac{r-1}{\pi_P(\mathbf{y})^2} \quad (27)$$

where we assume that the following bound:

$$\frac{\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \widetilde{Y}_{t-1})}{\pi_P(\mathbf{y})} \leq r \quad (28)$$

holds for all $\widetilde{Y}_1, \dots, \widetilde{Y}_T$.

Proof. i) We first prove that the estimator of the inclusion probabilities is unbiased through the following manipulation:

$$\mathbb{E}[\widehat{\pi}_P^{\text{is}}(\mathbf{y})] = \mathbb{E} \left[\frac{1}{M} \sum_{m=1}^M \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \widetilde{Y}_{t-1}^{(m)}) \right] \quad (40a)$$

$$= \sum_{\widetilde{Y}_1, \dots, \widetilde{Y}_T} \widetilde{P}(\widetilde{Y}_1, \dots, \widetilde{Y}_T) \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \widetilde{Y}_{t-1}) \quad (40b)$$

$$= \sum_{\widetilde{Y}_1, \dots, \widetilde{Y}_T} \widetilde{P}(\widetilde{Y}_1, \dots, \widetilde{Y}_T) \frac{P(\widetilde{Y}_1, \dots, \widetilde{Y}_T)}{P(\widetilde{Y}_1, \dots, \widetilde{Y}_T)} \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \widetilde{Y}_{t-1}) \quad (40c)$$

$$= \sum_{\widetilde{Y}_1, \dots, \widetilde{Y}_T} P(\widetilde{Y}_1, \dots, \widetilde{Y}_T) \frac{\widetilde{P}(\widetilde{Y}_1, \dots, \widetilde{Y}_T)}{\widetilde{P}(\widetilde{Y}_1, \dots, \widetilde{Y}_T)} \quad (\text{Lemma 4.1}) \quad (40d)$$

$$= \sum_{\widetilde{Y}_1, \dots, \widetilde{Y}_T} P(\widetilde{Y}_1, \dots, \widetilde{Y}_T) \quad (40e)$$

$$= \sum_{Y_1, \dots, Y_T} P(Y_1, \dots, Y_T) \mathbb{1}\{\mathbf{y} \in Y_T\} \quad (\text{definition of } \widetilde{Y}_T) \quad (40f)$$

$$= \pi_P(\mathbf{y}) \quad (40g)$$

ii) To show consistency, we appeal to the strong law of large number and the continuous mapping theorem. By the strong law of large numbers, we have that

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1}^{(m)}) = \pi_P(\mathbf{y}) \quad (41)$$

Since $1/x$ is continuous, we have

$$\begin{aligned} \lim_{M \rightarrow \infty} \frac{1}{\frac{1}{M} \sum_{m=1}^M \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1}^{(m)})} &= \frac{1}{\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1}^{(m)})} \\ &= \frac{1}{\pi_P(\mathbf{y})} \end{aligned} \quad (42a)$$

which shows consistency. Now, we derive a bound on the asymptotic variance of the inverse inclusion probabilities: Suppose,

$$\frac{\prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1})}{\pi_P(\mathbf{y})} \leq r, \quad \forall \tilde{Y}_1, \dots, \tilde{Y}_T \quad (43)$$

We start with the variance of importance sampling. This is a standard result (Rubinstein and Kroese, 2016). Then we proceed with algebraic manipulation integrating the assumption above:

$$\sum_{\tilde{Y}_1, \dots, \tilde{Y}_T} \frac{\mathbb{1}\{\mathbf{y} \in \tilde{Y}_T\}^2 P(\tilde{Y}_1, \dots, \tilde{Y}_T)^2}{\tilde{P}(\tilde{Y}_1, \dots, \tilde{Y}_T)} - \pi_P(\mathbf{y})^2 \quad (44a)$$

$$= \sum_{\tilde{Y}_1, \dots, \tilde{Y}_T} P(\tilde{Y}_1, \dots, \tilde{Y}_T) \prod_{t=1}^T \pi_{Q_t}(\mathbf{y}_{\leq t} | \tilde{Y}_{t-1}) - \pi_P(\mathbf{y})^2 \quad (44b)$$

$$\leq \sum_{\tilde{Y}_1, \dots, \tilde{Y}_T} P(\tilde{Y}_1, \dots, \tilde{Y}_T) \pi_P(\mathbf{y}) r - \pi_P(\mathbf{y})^2 \quad (44c)$$

$$= \pi_P(\mathbf{y}) \pi_P(\mathbf{y}) r - \pi_P(\mathbf{y})^2 \quad (44d)$$

$$= \pi_P(\mathbf{y})^2 r - \pi_P(\mathbf{y})^2 \quad (44e)$$

$$= (r - 1) \pi_P(\mathbf{y})^2 \quad (44f)$$

We can compute the asymptotic variance by the delta rule:

$$\mathbb{V}_a \left[\frac{1}{\hat{\pi}_P^{\text{is}}(\mathbf{y})} \right] = \frac{1}{M} \frac{\mathbb{V}[\hat{\pi}_P^{\text{is}}(\mathbf{y})]}{\pi_P(\mathbf{y})^4} \quad (\text{apply the delta rule}) \quad (45a)$$

$$\leq \frac{1}{M} \frac{(r - 1) \pi_P(\mathbf{y})^2}{\pi_P(\mathbf{y})^4} \quad (\text{plugging in the above bound}) \quad (45b)$$

$$= \frac{1}{M} \frac{(r - 1)}{\pi_P(\mathbf{y})^2} \quad (45c)$$

which proves the result. ■

C Pseudocode

Algorithm 1 Dynamic program algorithm for Z

Input: K : Size of subset

w_1, \dots, w_N : weights for each element of the base set

Output: W : elementary symmetric polynomials of w_1, \dots, w_N ; $Z = W_{N,K}$

```
1:  $W \leftarrow \mathbf{0}^{(K+1) \times (N+1)}$ 
2:  $W_{0,:} = 0$ ;  $W_{:,0} = 1$ 
3: for  $n = 1, \dots, N$  :
4:   for  $k = 1, \dots, K$  :
5:      $W_{n,k} \leftarrow W_{n-1,k} + w_n W_{n-1,k-1}$ 
6: return  $W$ 
```

Algorithm 2 Dynamic program for calculating inclusion probabilities π

Input: K : Size of subset

w_1, \dots, w_N : weights for each element of the base set

```
1: Run Alg. 1 to compute  $W$ 
2:  $\triangleright$  The code below was derived by manually apply algorithmic differentiation (Bücker et al., 2006) to Alg. 1.
3:  $\dot{W} \leftarrow \mathbf{0}^{(K+1) \times (N+1)}$   $\triangleright$  Initialize adjoints
4:  $\dot{w} \leftarrow \mathbf{0}^N$ 
5:  $\dot{W}_{N,K} = 1$   $\triangleright$  Initialize output value to 1
6: for  $n = N, \dots, 1$  :
7:   for  $k = K, \dots, 1$  :
8:      $\dot{w}_n += \dot{W}_{n,k} W_{n-1,k-1}$ 
9:      $\dot{W}_{n-1,k-1} += \dot{W}_{n,k} w_n$ 
10:     $\dot{W}_{n-1,k} += \dot{W}_{n,k}$ 
11:  $\triangleright$  Apply Eq. (24)
12:  $\pi \leftarrow \mathbf{0}^N$ 
13: for  $n = 1 \dots N$  :
14:    $\pi_n \leftarrow \frac{w_n}{Z} \dot{w}_n$ 
return  $\pi$ 
```

D Experimental Setup and Additional Results

We use a Transformer-based model trained according to Ott et al. (2018b) on the WMT’14 English-French dataset.¹⁸ We use the pre-trained model checkpoints made available by fairseq.¹⁹ Data preprocessing steps, model hyperparameters and baseline performances can be found in the original work and on the fairseq website. All evaluations are performed on the wmt14.v2.en-fr.newstest2014 version of the newstest data set. We show additional results using the setup in §5 in Figs. 4 to 6. We provide an empirical runtime analysis in Fig. 7. Table 2 shows the variance of baseline estimator value for the three sentences used in RMSE experiments.

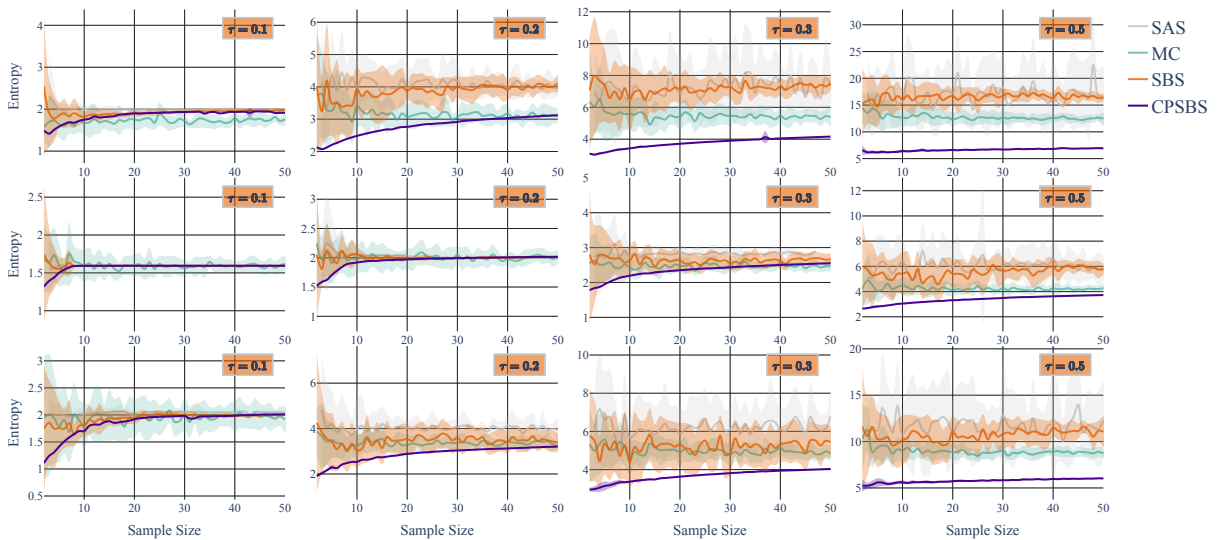


Figure 4: Entropy estimates for three different sentences using estimators for respective decoding methods. τ indicates scaling temperature. Values are chosen to mimic (Kool et al., 2019).




Figure 5: BLEU score estimates using unnormalized versions of SBS and CPSBS estimators.

	BLEU Estimator				Entropy Estimator			
	Temperature				Temperature			
	0.1	0.2	0.3	0.5	0.1	0.2	0.3	0.5
Sentence# 1500	0.00	0.00	0.03	0.08	0.00	0.01	0.07	0.10
Sentence# 2000	0.01	0.04	0.04	0.08	0.00	0.00	0.00	0.02
Sentence# 2500	0.07	0.09	0.25	0.84	0.00	0.01	0.03	0.04

Table 2: Variance of baseline estimator (MC for $k = 200$ in 50 iterations) for the three sentences.

¹⁸available at <http://statmt.org/wmt14/translation-task.html>

¹⁹<https://github.com/pytorch/fairseq/tree/master/examples/translation>




Figure 6: BLEU score estimates for CPSBS both with and without truncation of the sampling distribution. We see that our estimator with truncation provides virtually the same results.




Figure 7: A comparison between decoding time of different sampling methods. The y-axis shows the average decoding time of the three sentences as before. The x-axis shows the number of samples taken for each sentence. All methods are tested on CPU.