
QUEST: Quality-Aware Metropolis-Hastings Sampling for Machine Translation

Gonçalo R. A. Faria^{1,*}, Sweta Agrawal², António Farinhas^{2,3},
Ricardo Rei⁴, José G. C. de Souza⁴, André F.T. Martins^{2,3,4,5}

¹University of Washington, ²Instituto de Telecomunicações,
³Instituto Superior Técnico, Universidade de Lisboa, ⁴Unbabel, ⁵ELLIS Unit Lisbon
gfarria@cs.washington.edu

Abstract

An important challenge in machine translation is to generate high-quality and diverse translations. Prior work has shown that the estimated likelihood from the MT model correlates poorly with translation quality. In contrast, quality evaluation metrics (such as COMET or BLEURT) exhibit high correlations with human judgments, which has motivated their use as rerankers (such as quality-aware and minimum Bayes risk decoding). However, relying on a single translation with high estimated quality increases the chances of “gaming the metric”. In this paper, we address the problem of sampling a *set* of high-quality and diverse translations. We provide a simple and effective way to avoid over-reliance on noisy quality estimates by using them as the energy function of a Gibbs distribution. Instead of looking for a mode in the distribution, we generate multiple samples from high-density areas through the Metropolis-Hastings algorithm, a simple Markov chain Monte Carlo approach. The results show that our proposed method leads to high-quality and diverse outputs across multiple language pairs (ENGLISH↔{GERMAN, RUSSIAN}) with two strong decoder-only LLMs (ALMA-7B, TOWER-7B).

1 Introduction

Machine translation (MT) is becoming increasingly more accurate and powerful, as it benefits from the many capabilities and acquired knowledge of large language models (LLMs) (Freitag et al., 2023b; Hendy et al., 2023). However, for many domains and languages the quality of translation is still not satisfactory—for example, hallucinations or critical errors are a serious issue when translating high-risk content, as in medical and legal domains (Khoong et al., 2019; Taira et al., 2021; Shen et al., 2023; Guerreiro et al., 2023b; Sanz-Valdivieso and López-Arroyo, 2023; Grimm et al., 2024). Developing procedures for sampling higher-quality translations is therefore in high demand.

It is known that the output quality of the translations generated by maximizing the model likelihood is limited because of the *inadequacy of the mode*—models tend to produce distributions over outputs that are highly peaked, favoring a single hypothesis (Eikema and Aziz, 2020; Peters and Martins, 2021; Eikema, 2024); and improving search often makes things worse (Koehn and Knowles, 2017; Stahlberg and Byrne, 2019). In general, maximizing model likelihood tends to overlook hypotheses that could be equally valid and more appropriate in certain contexts. To address this limitation, a string of work has been initiated towards “quality-aware decoding”, which leverages powerful quality estimation (QE) and evaluation metrics, such as COMET (Rei et al., 2022) or BLEURT (Yan et al., 2023), to explore and rerank a broader range of candidate hypotheses generated via sampling from the model’s distribution, selecting the best-1 (Freitag et al., 2022a; Fernandes et al., 2022b; Farinhas et al., 2023) or the best-*k* (Jinnai et al., 2024; Singhal et al., 2023) according to these learned metrics.

*Work done while at Instituto de Telecomunicações.

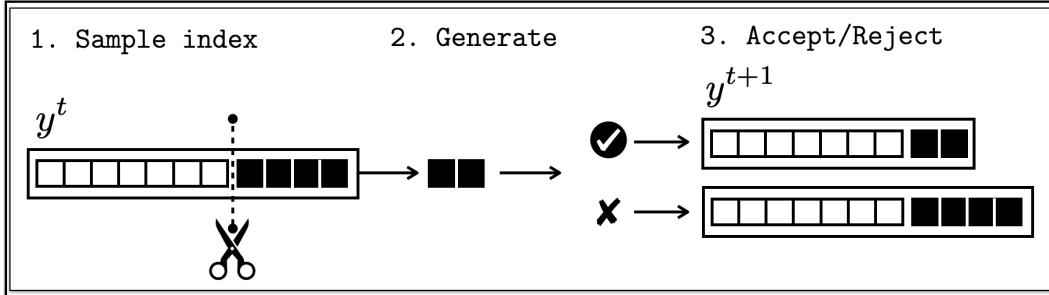


Figure 1: QUEST samples an index from the current translation (y^t), removes all elements to the right of the index, generates a new continuation, and then uses the Metropolis-Hastings acceptance criterion to decide whether to accept or reject the resulting new translation. The process continues for a fixed number of T iterations.

Despite the benefits, reranking-based methods have their own drawbacks. There is a risk of these approaches overfitting to the metrics used, potentially leading to illusory gains in quality, as the translations obtained via optimizing these metrics may not always be preferred by humans when compared to other alternatives (Fernandes et al., 2022b). Secondly, their effectiveness is limited by the quality of the initial candidate set. For instance, Vernikos and Popescu-Belis (2024) show that many high-quality translations, generated by combining translation hypotheses, are less likely to be sampled from the model’s distribution even with a very large pool size.

One potential way to remedy these issues, which we explore in this paper, is to use these automatic metric proxies as the **energy function of a Gibbs distribution**. However, sampling hypotheses from this distribution presents a difficult challenge: unlike likelihood-based sampling, “quality-based sampling” from a Gibbs distribution cannot be performed autoregressively, and it is intractable to enumerate and score all possible hypotheses. We address this challenge by proposing a simple and effective Markov chain Monte Carlo approach (MCMC) method, combining the Metropolis-Hastings algorithm with a suitable proposal distribution. Our method, **Quality-Aware Metropolis-Hastings (QUEST) Sampling**, uses a novel proposal distribution that is compatible with sentence-level evaluation metrics, common in text generation tasks like MT (Figure 1). We further note that while we focus on MT, where automatic QE metrics are more developed and robust, our proposed method is general and can be applied to other natural language processing (NLP) tasks. Furthermore, as our method is agnostic to the specific quality metric used in the Gibbs distribution, it can directly benefit from any future improvements in the metrics.

We show that QUEST sampling results in high-quality and diverse samples on multiple test beds (WMT23 ENGLISH \leftrightarrow {GERMAN, RUSSIAN}) and with multiple decoder-only LLMs (TOWER-7B, ALMA-7B). Our method generates many novel hypotheses from the high-density regions unlikely to be generated via ancestral sampling. Furthermore, with increasing chain size, average quality as measured by automatic metrics continues to improve, unlike ancestral sampling, where the candidate set quality remains unchanged even with a larger pool size.²

2 Background

2.1 Large Language Models for Machine Translation

Generating translations from autoregressive LLMs (either encoder-decoder or decoder-only) involves conditioning the language model on a prompt x , which is a sequence of tokens $(x_1, x_2, \dots, x_L) \in \mathcal{X}$ encoding the text to be translated coupled with a translation instruction (Raffel et al., 2020; Hendy et al., 2023). Let \mathcal{V} be a fixed vocabulary and $\mathcal{Y} = \mathcal{V}^*$ its Kleene closure. The joint distribution over the output translations, $y \in \mathcal{Y}$, given the prompt x , can be factorized as the product of conditional probabilities over individual tokens (y_1, y_2, \dots, y_N) , where each $y_i \in \mathcal{V}$. The probability of a

²We release the code to replicate our experiments at <https://www.questdecoding.com>.

particular translation y for a given input x can be written as

$$p_{\text{LM}}(y|x) = \prod_{i=1}^N p_{\text{LM}}(y_i|y_{<i}, x). \quad (1)$$

Here, $y_{<i} := (y_1, y_2, \dots, y_{i-1})$. During inference, a translation is generated by sampling one token at a time from the distribution $p_{\text{LM}}(y_i|y_{<i}, x)$, adjusted by a temperature, τ . The (adjusted) probability of generating a particular token y_i , given the preceding tokens $y_{<i}$ and the input x , is defined as:

$$p_{\text{LM},\tau}(y_i = v|y_{<i}, x) = \frac{p_{\text{LM}}(y_i = v|y_{<i}, x)^{1/\tau}}{\sum_{v' \in \mathcal{V}} p_{\text{LM}}(y_i = v'|y_{<i}, x)^{1/\tau}}. \quad (2)$$

Lower temperature values τ make the distribution more deterministic, favoring the most probable tokens, whereas higher values make the distribution flatter, approximating a uniform distribution.

However, multiple works have scrutinized the reliability of model likelihood as a measure for translation quality (Ott et al., 2018; Stahlberg and Byrne, 2019; Eikema and Aziz, 2020; Freitag et al., 2022a; Eikema, 2024). Instead of solely relying on the likelihood, these works advocate using an automatic translation quality metric as a utility function for minimum Bayes risk (MBR) decoding or reranking based on QE metrics. This shift not only improves the selection of translations but also facilitates the exploration of the underlying distribution learned by the models. However, it is crucial to acknowledge that the overall translation quality is still contingent on the quality of the candidate pool. We next discuss common automatic metrics used for assessing translation quality.

2.2 Automatic Metrics for Machine Translation

Automatic quality assessment of machine-generated translations has received considerable attention recently, resulting in metrics that attain high correlations with human judgment of translation quality (Freitag et al., 2022b, 2023b). These automatic metrics are meant to assess the quality of a translation across multiple dimensions (*e.g.* fluency, adequacy) and can provide reliable feedback when human judgments are unavailable. Among these metrics, neural learned metrics that are trained on human translation quality assessment scores or error span annotations have gained significant traction (Rei et al., 2022; Yan et al., 2023; Guerreiro et al., 2023a; Perrella et al., 2022).

For aligning automatically generated translations with human quality preferences, many works have proposed to use automatic metrics as an alternative to human feedback during MT training (Shen et al., 2016; Wieting et al., 2019; He et al., 2024; Xu et al., 2024b) or decoding Shen et al. (2004); Fernandes et al. (2022b); Freitag et al. (2022a, 2023a). Freitag et al. (2022a) show the efficacy of using reference-based neural metrics as utility functions for MBR decoding over lexical alternatives. Further, Fernandes et al. (2022a) use QE metrics like COMET-QE to select 1-best or N -best translations from a pool of candidate hypotheses. Their analysis shows that while these automatic metrics can be useful, they might not always reflect human preferences accurately. Additionally, extra care has to be taken when optimizing systems for these metrics, as the improvements might be attributable to overfitting or “gaming the metric”, rather than being genuine improvements in translation quality. Nevertheless, these metrics encode useful information about the quality of translations and can still be useful to obtain high-quality translations.

3 An MCMC-based Decoding Approach for Text Generation

Given that we have access to an automatic metric that quantifies how desirable a particular translation is, we aim to address the following questions:

Can we sample directly in proportion to their corresponding quality values? Can we use automatic metrics that already give us a reliable estimate of human-perceived quality to achieve that? Finally, can we use this process to obtain diverse high-quality samples?

To answer these questions and to address the limitations of quality-aware decoders, we propose to take an alternative approach, quality-aware *sampling*, which ensures high quality and diversity. Recognizing that naturally occurring texts can have numerous valid translations (Nida, 1964; Dreyer and Marcu, 2012; Ott et al., 2018; Mayhew et al., 2020), the ability to generate diverse translation hypotheses is paramount.

To this end, we first present background on Metropolis-Hastings in Section 3.1. Section 3.2 discusses our proposal distribution. Finally, we connect our approach to reinforcement learning with human feedback (RLHF) in Section 3.3.

3.1 Metropolis-Hastings

The problem of sampling translation hypotheses in proportion to a metric, $r(x, y)$, can be framed as sampling from the following Gibbs distribution:

$$\pi_\beta(y|x) = \frac{1}{Z_\beta(x)} \exp\left(\frac{r(y, x)}{\beta}\right), \quad (3)$$

where $Z_\beta(x) = \sum_y \exp\left(\frac{r(y, x)}{\beta}\right)$ and β is the temperature of the Gibbs distribution. We denote the corresponding unnormalized density by $\tilde{\pi}_\beta(y|x)$, which unlike $Z_\beta(x)$ can be evaluated for any (x, y) .

While several algorithms have been proposed to sample approximately from such intractable distributions (Miao et al., 2018; Berglund et al., 2015; Amini et al., 2023), we resort to Metropolis-Hastings (Metropolis et al., 1953, MH) due to its simplicity and flexibility in handling a wide range of proposal distributions. The MH algorithm generates a sequence of samples from a **target distribution**, here $\pi_\beta(y|x)$, by constructing a Markov chain (y^0, y^1, \dots, y^T) that has $\pi_\beta(y|x)$ as its equilibrium distribution. It starts from an arbitrary hypothesis y^0 . In the t^{th} iteration, it draws a new hypothesis y from a **proposal distribution** $q(y|y^t, x)$. This hypothesis is accepted with an acceptance probability $\alpha_\beta(y, y^t) \in [0, 1]$ given by:

$$\alpha_\beta(y, y^t) = \min\left\{1, \frac{\pi_\beta(y|x) q(y^t|y, x)}{\pi_\beta(y^t|x) q(y|y^t, x)}\right\}. \quad (4)$$

If the candidate y is accepted, the next state in the chain becomes $y^{t+1} = y$; if rejected, the chain stays at $y^{t+1} = y^t$. The process repeats for some number of steps T and in the end it returns the hypothesis y^T . Note that, while computing the likelihood $\pi_\beta(y|x)$ of a particular hypothesis y under the Gibbs distribution is intractable (due to the intractable partition function $Z_\beta(x)$), evaluating the acceptance criterion $\alpha_\beta(y, y^t)$ is easy, because it depends only on the likelihood ratio, in which the normalization constants cancel out, *i.e.*:

$$\frac{\pi_\beta(y|x)}{\pi_\beta(y^t|x)} = \frac{\tilde{\pi}_\beta(y|x)}{\tilde{\pi}_\beta(y^t|x)} = \exp\left(\frac{r(y, x) - r(y^t, x)}{\beta}\right). \quad (5)$$

MH converges to the unique stationary distribution $\pi_\beta(y|x)$, regardless of the initial distribution we start with, if the transition distribution of the Markov chain, *i.e.*, $p(y^t|y^{t-1}, x) = q(y^t|y^{t-1}, x)\alpha(y^t, y^{t-1})$ satisfies the *Markov chain ergodic theorem* (Neal, 2011).

This requires a suitable proposal distribution $q(y|y^t, x)$ which must be **irreducible** and **aperiodic**. To ensure irreducibility, the proposal distribution should have sufficient support, such that it is possible to transition from any state to any other state with a nonzero probability in a finite number of steps. Aperiodicity ensures that the Markov chain does not get stuck in a cyclic behavior, where it keeps returning to the same states in a fixed pattern. Additionally, due to the acceptance criterion defined in Eq. 4, the transition distribution satisfies the detailed balance condition:

$$\pi_\beta(y|x) p(y^t|y, x) = \pi_\beta(y^t|x) p(y|y^t, x). \quad (6)$$

It is trivial to show that the chain has the target distribution, $\pi_\beta(y|x)$ as its stationary distribution:

$$\pi_\beta(y|x) = \sum_{y' \in \mathcal{Y}} p(y|y', x) \pi_\beta(y'|x). \quad (7)$$

Note that MH can work with almost any proposal distribution. However, if the detailed balance conditions are far from holding, the acceptance probability will be low for transitions, making the convergence to the target distribution very slow and the approach impractical. Hence, choosing a suitable proposal distribution for the task is essential. We next describe our proposal distribution, $q(y|y^t, x)$, which overcomes these limitations and satisfies the required constraints.

3.2 Proposal distribution

Previous works (Berglund et al., 2015; Miao et al., 2018; Su et al., 2018) use proposal distributions that generate hypotheses with one-word or token-level modifications. The different formulations in their most basic form propose a Markov chain in which the state comprises the sentence $y^t \in \mathcal{Y}$ and an index variable $i^t \in \{0, \dots, |y^{t-1}|\}$. At every step, an index is sampled that determines the token to be changed, and a new token is then sampled based on the following full conditional distribution:

$$q(y_i | y_{<i}, y_{>i}) = \frac{\tilde{\pi}(y_{<i}, y_i, y_{>i})}{\sum_{y'_i \in \mathcal{V}} \tilde{\pi}(y_{<i}, y'_i, y_{>i})}, \quad (8)$$

where \mathcal{V} is created by considering the k most likely tokens (k is generally large) under $p_{\text{LM}}(y_i | y_{<i})$.

This procedure, however, has several limitations: first, it makes it difficult to handle more general combinatorial constraints, in our case more sophisticated metrics. As we only explore adjacent positions in the sentence space due to small local changes, the Markov chain risks becoming trapped in infeasible states, necessitating a very large number of steps T to converge. Second, relying solely on token-level modifications makes it exceedingly difficult to generate plausible text, making it impractical to align generations with QE metrics or general reward models. We show that this is indeed the case in Section 5.2.

We instead propose a simple and effective procedure that only requires generating a single hypothesis from the model p_{LM} and a single evaluation from the metric, $r(y, x)$, and still allows the Markov chain to converge to the target distribution. We characterize the proposal by the following procedure:

1. Given an instance y^t with length $n^t := |y^t|$, sample an index i , $i \sim q(i | n^t)$.
2. Generate a completion $y_{\geq i}$ from $p_{\text{LM}}(y_{\geq i} | y_{<i}^t, x)$.

Note that, due to the nature of our proposal distribution, y^t and y share a common substructure (prefix) before the index i , i.e., $y_{<i}^t = y_{<i}$, which implies that

$$q(y | y^t, x, i) = p_{\text{LM}}(y_{\geq i} | y_{<i}^t, x) \prod_{j < i} \delta(y_j, y_j^t), \quad (9)$$

where $\delta(y_j, y_j^t)$ is the Kronecker delta function, which assigns zero probability to prefix tokens which are different from $y_{<i}^t$ and probability one to tokens matching the prefix.

The complete proposal when we include the index distribution $q(i | n^t)$, which depends on the previous sentence lengths $n^t := |y^t|$, is given as

$$q(y, i | y^t, x) = q(i | n^t) p_{\text{LM}}(y_{\geq i} | y_{<i}^t, x) \prod_{j < i} \delta(y_j, y_j^t). \quad (10)$$

We will use the uniform distribution as the index distribution, i.e., $q(i | n^t) = 1/n^t$ for each $i \in [n^t]$, unless otherwise stated. Algorithm 1 describes the complete sampling process.

This proposal is both **irreducible** and **aperiodic** as there is a non-zero probability of going from a particular sentence to any other sentence and back. When $i = 0$, we can generate any text from the language model, which, when using ancestral sampling, implies that we can sample any possible sequence. As our approach only requires access to the token-level log probabilities and the ability to generate *good* continuations given a prefix, it can also be used with closed LLMs through an API as long as the it provides access to the sample likelihoods. However, we limit our experiments to open-source models due to the incurred cost and lack of training details about these black-box models.

Algorithm 1 Quality-Aware Metropolis-Hastings (QUEST) Sampling

```

1: Input:  $x, p_{\text{LM}}, r, T$ 
2: Hyperparameters:  $\tau, t_{\text{burning}}, \beta$ 
3: Sample initial response  $y_0 \sim p_{\text{LM}}(\cdot | x)$ 
4:  $t \leftarrow 1$ 
5: for  $1$  to  $T$  do
6:   Sample index  $i \sim q(i | n^{t-1})$ 
7:   Sample  $y_{\geq i} \sim p_{\text{LM}}(\cdot | y_{<i}^{t-1}, x)$ 
8:    $y \leftarrow (y_{<i}^{t-1}, y_{\geq i})$ 
9:   Compute  $\alpha_{\beta}(y, y^{t-1})$  through Eq. 4
10:  Sample  $\alpha$  uniformly in  $[0, 1]$ 
11:  if  $\alpha \leq \alpha_{\beta}(y, y^{t-1})$  then
12:     $y^t \leftarrow y$ 
13:     $t \leftarrow t + 1$ 
14:  end if
15: end for
16: return  $y^{t_{\text{burning}}}, \dots, y^t$ 

```

3.3 Connections to Reinforcement Learning with Human Feedback

Reinforcement Learning with Human Feedback (RLHF) leverages human feedback to guide the learning process of complex NLP tasks (Stiennon et al., 2022; Fernandes et al., 2023; Kaufmann et al., 2023). The process is as follows. Given a language model, one generates hypotheses $y \in \mathcal{Y}$ given an input prompt x and gathers human feedback about which outputs are preferable. This preference data is then used to train a proxy reward model $r_\phi(y, x)$. Finally, reinforcement learning (RL) methods are used to optimize the original LM with respect to the reward model, following

$$\max_{\pi} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi(y|x)} \left[\frac{r_\phi(x, y)}{\beta} \right] - D_{\text{KL}} [\pi(y|x) \parallel p_{\text{LM}}(y|x)]. \quad (11)$$

Many works show that the optimal solution to the KL-constrained reward maximization objective takes the form (Peters and Schaal, 2007; Peng et al., 2019; Korbak et al., 2022b,a; Go et al., 2023):

$$\pi(y|x) = \frac{1}{Z_\beta(x)} p_{\text{LM}}(y|x) \exp \left(\frac{r(y, x)}{\beta} \right), \quad (12)$$

where $Z_\beta(x) = \sum_{y \in \mathcal{Y}} p_{\text{LM}}(y|x) \exp \left(\frac{r(x, y)}{\beta} \right)$. While we cannot sample autoregressively from this distribution, this density can be represented as a Gibbs distribution with the corresponding reward function $\tilde{r}(x, y) = \log p_{\text{LM}}(y|x) + \frac{r(x, y)}{\beta}$.

Instead of the target distribution expressed in Eq. 3, we could define the above distribution as our target when using QUEST to sample from it, avoiding optimizing the objective in Eq. 11 directly. If we formulate the acceptance criterion using this target density and our proposal distribution introduced in Eq. 10, we obtain the following:

$$\alpha_\beta(y, y^t) = \min \left\{ 1, \exp \left(\frac{r(y, x) - r(y^t, x)}{\beta} \right) \frac{q(i|n)}{q(i|n^t)} \right\}. \quad (13)$$

The full derivation is in Appendix A. Using this approach, we can align language model generations without access to the model weights, log probabilities, or RL. We provide a preliminary experimental comparison using the two target distributions (Eq. 3 and Eq. 12) using QUEST in Appendix C.4.

4 Experimental Settings

Data and Evaluation We test our approach on the WMT23 test sets (Kocmi et al., 2023) covering four language pairs, ENGLISH \leftrightarrow {GERMAN, RUSSIAN}.

We evaluate the quality and the diversity of the generated texts as follows. Suppose $\bar{\mathcal{Y}}$ is the set of hypotheses generated for the source text x with reference hypothesis y^* and $\mathcal{D} = \{(x, y^*, \bar{\mathcal{Y}})\}$ represents the evaluation set. We compute the mean quality over each set of hypotheses using xCOMET-XL (Guerreiro et al., 2023a) as $\frac{1}{|\mathcal{D}|} \sum_{(x, y^*, \bar{\mathcal{Y}}) \in \mathcal{D}} \frac{1}{|\bar{\mathcal{Y}}|} \sum_{y \in \bar{\mathcal{Y}}} \text{xCOMET-XL}(x, y, y^*)$. Similarly, we measure the mean diversity using the average pairwise BLEU (Papineni et al., 2002; Shen et al., 2019) as $1 - \frac{1}{|\mathcal{D}|} \sum_{(x, y^*, \bar{\mathcal{Y}}) \in \mathcal{D}} \frac{1}{|\bar{\mathcal{Y}}|(|\bar{\mathcal{Y}}|-1)} \sum_{(y, y') \in \bar{\mathcal{Y}}^2 \text{ s.t. } y \neq y'} \text{BLEU}(y, y')$.³

Models We use TOWER-7B (Alves et al., 2024, Unbabel/TowerInstruct-7B-v0.2) and ALMA-7B (Xu et al., 2024a, haoranxu/ALMA-7B), two strong decoder-only MT models based on LLAMA2-7B (Touvron et al., 2023) as these models achieve competitive translation quality with GPT-4 and productized models like Google Translate. Unlike ALMA-7B, TOWER-7B uses bilingual MT data as well as datasets from MT-related tasks during training. Prompts are shown in Appendix B.

Automatic Metrics for QUEST We use COMETKIWI-XL (Rei et al., 2023), a reference-free QE model built on top of XLM-R XL (Goyal et al., 2021) and trained to predict human-rated direct assessments (Graham et al., 2013). This metric showed the highest correlations with human judgments on the QE Shared Task organized by the eighth conference on Machine Translation (WMT 2023) Blain et al. (2023). We transform the normalized scores from the QE model into z -scores using a logit transformation with clamping applied to mitigate overflow.

³<https://github.com/mjpost/sacrebleu/tree/master>

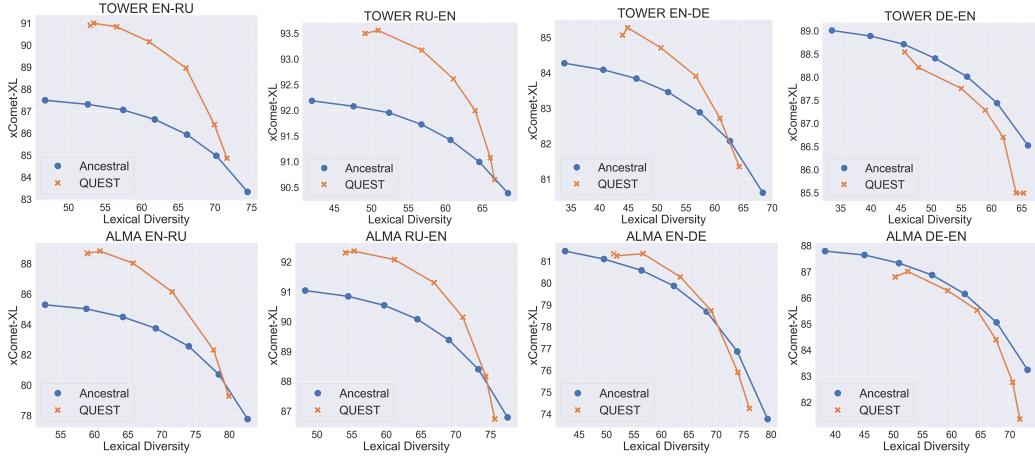


Figure 2: Average quality vs. diversity on WMT23 datasets. Different points represent different hyperparameter values. QUEST outperforms ancestral sampling in six out of eight settings.

Decoding Configurations For ancestral sampling, we consider temperature values τ between 0.2 and 1.0, with an equally spaced interval of 0.1. For generations with QUEST, we sample from the proposal distribution using $\tau = 0.8$ and vary the parameter β of the target Gibbs distribution from the following range of values $\{0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0\}$. The number of ancestral samples and decoding steps are both set to 128. We use VLLM (Kwon et al., 2023), a high-throughput and memory-efficient inference engine for generating outputs.

Compute Comparison: Ancestral Vs QUEST We use the number of output tokens as a metric to compare the different approaches. For the sake of simplicity, let us assume that the output sentence has a fixed size of N . On average, QUEST in T steps generates $(\frac{T-1}{2} + 1)N = \frac{(T+1)N}{2}$ tokens, N tokens for the initial hypothesis, and then on average $\frac{N}{2}$ tokens for the remaining $T - 1$ steps in the chain. If we contrast against sampling T sentences using ancestral sampling, we decode $T \times N$ tokens. This suggests that for an equal number of samples generated using ancestral sampling and QUEST, the latter results in about $\frac{2T}{(T+1)} \approx 2$ times as many tokens, on average. Note, however, that the computational cost of QUEST is higher than ancestral sampling, as the hypotheses are generated sequentially and evaluated at every step. We run our experiments on NVIDIA RTX A6000 GPUs. Each ancestral sampling and QUEST for run with $T = 128$ takes, on average, 1 hour and 6 hours, respectively, for 2000 unique source texts on a single GPU. The compute bottleneck for QUEST also arises from using a large QE metric, potentially distilling this metric into smaller models could help reduce the compute time. We leave this to future work.

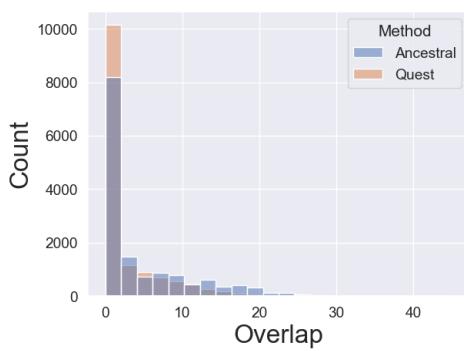
5 Results

5.1 Main Findings

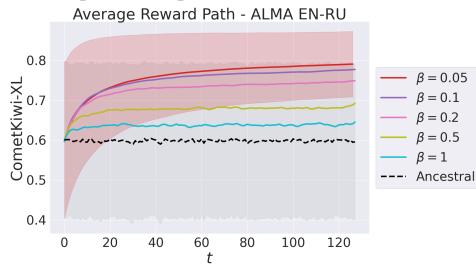
The main results of our experiments are presented in Figure 2.

QUEST results in better translation quality-diversity trade-offs. Across language directions and models, the samples generated by QUEST tend to have better or similar quality than ancestral sampling as measured by xCOMET-XL.⁴ As QUEST does not directly use the reference-based metric, xCOMET-XL, we reduce the chance of overfitting to the metric and thus the gains represent the ability of QUEST to improve translation quality more realistically. We also report Comet-22 vs diversity results in the Appendix Figure 9: the trends remain the same. The benefits of the proposed approach are more noticeable when translating from English ($EN \rightarrow \{DE, RU\}$). Specifically, for $EN \leftrightarrow RU$, our model improves xCOMET-XL by up to 2 points for both language models, showing the efficacy of QUEST over ancestral sampling.

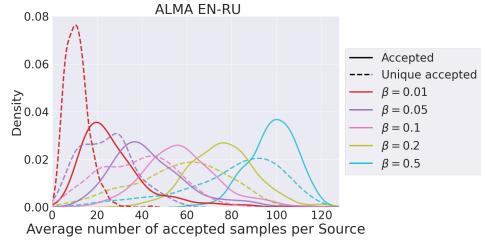
⁴We provide results on four additional language pairs in Appendix C.6.



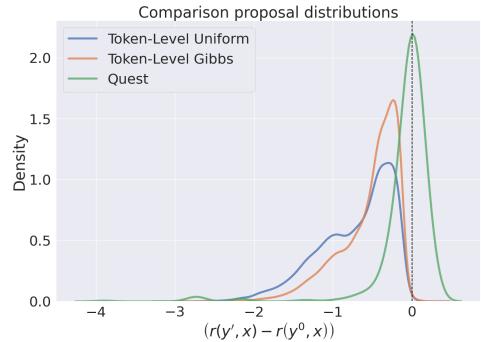
(a) Sample Overlap between ancestral and QUEST.



(b) Average quality over MCMC steps and number of ancestral samples t .



(c) Distribution of # of accepted samples: high β results in higher acceptance rates.



(d) Reward differences between an initial hypothesis and proposal hypotheses given the initial hypothesis.

Although ancestral exhibits better quality than QUEST for DE-EN, further analysis suggests that this discrepancy may stem from the constraints of the QE metric used to model the translation preferences. Notably, QUEST demonstrates significantly higher COMETKIWI-XL scores across the board (see Appendix Figure 7), suggesting that better and more robust QE models can result in improved translation quality.⁵ Furthermore, WMT23 DE-EN includes short passages that might require additional steps to reach the high-density regions. Based on a length analysis (See Appendix C.2), we observe that the quality of QUEST lags behind ancestral sampling, specifically for longer sentences. We leave the exploration of finding optimal steps for convergence and adapting the proposal distribution for document-level MT to future work.

5.2 Ablation Analysis

We present further analysis on some properties of our proposed approach using WMT23 EN→RU datasets with translations generated using ALMA-7B.

QUEST generates less-likely high-quality samples. We compute the set overlap for QUEST and ancestral samples ($T = 128$) with an independent draw of 512 ancestral samples to investigate whether our approach results in hypotheses from unexplored regions (Fig. 3a). Compared to ancestral, QUEST results in hypotheses that are not found in the larger pool ($4\times$) of ancestral samples as illustrated by a higher mass at the overlap value of 0. This demonstrates that QUEST effectively gets to regions of the output manifold that would be less likely sampled by the LM and still attain, on average, better quality and diversity.

Average reward improves over decoding steps. Figure 3b shows the average reward with increasing decoding steps and the sample size for QUEST and ancestral respectively. As ancestral results in independent samples, the mean reward estimate remains unchanged for different sample sizes. However, for QUEST, hypotheses are gradually sampled in the direction of higher-quality regions, closer to the target density, resulting in increased average reward with more steps. We can also observe that the standard deviation of the reward over all samples decreases with the increasing number of steps, suggesting that the model eventually reaches a better target distribution.

⁵Kocmi et al. (2024) report that a difference of 2.7 COMETKIWI-XL on average is statistically significant with a 98.9% accuracy. QUEST improvements are always greater than this value across the board.

High β value results in higher acceptance rates. Figure 3c shows the distribution of accepted samples when varying β , considering all (blue) versus unique (red) accepted samples. As expected, on average, the number of accepted samples is smaller than the number of steps T depending upon the rejection rate—low β values lead to higher rejection rates. Furthermore, within the accepted samples, we also observe many repeats. This can be attributed to the observation that the language model has a low entropy distribution for continuations when the sample indices lie at the end of sentences. Moreover, for probable sentences, only a few have a high reward, which could result in the Markov chain getting stuck in a particular state. Increasing the temperature of the LM or adjusting the index distribution to have less density in the last few tokens could potentially reduce the number of repeats. We leave this exploration to future work.

Our sentence-level proposal generates better candidates. For a random example sampled from the WMT23 dataset, we generate 25k hypotheses using three proposal distributions: a) token-level modification with uniform probability b) token-level modification with full posterior presented in Eq. 8, and c) our sentence-level proposal (Eq. 10) and calculate the reward differences between the original (y) and the generated hypothesis (y'): $r(y', x) - r(y, x)$. Figure 3d shows its distribution: for proposals (a) and (b), the reward for the generated hypotheses is almost always lower than the initial translation. On the other hand, our proposal results in assigning half of the probability mass to hypotheses that improve over y , leading to faster convergence over token-level alternatives.

6 Related Work

Sampling from Intractable Gibbs distributions. Several methods have been proposed to sample approximately from Gibbs distributions in text generation using autoregressive and masked-language models. Prior works (Miao et al., 2018; Zhang et al., 2020) use MH with a proposal distribution that makes token-level modifications for constrained generation tasks. Since masked language models (MLM) do not have a straightforward mechanism for sampling text, MCMC has been widely explored using variations of Gibbs sampling (Berglund et al., 2015; Su et al., 2018; Wang and Cho, 2019; Yamakoshi et al., 2022). However, Goyal et al. (2022) shows that the masked conditional distributions from MLMs result in invalid Gibbs samplers and, therefore, proposes to use MH on the masked conditionals, resulting in higher-quality text. Mireshghallah et al. (2022); Forristal et al. (2023) build on this work and use MLMs for sampling from Gibbs distributions. Some works (Kumar et al., 2022; Qin et al., 2022; Amini et al., 2023; Du et al., 2023) also adapt the Hamiltonian MCMC algorithms originally designed for high-dimensional continuous distributions for the discrete scenario (Duane et al., 1987; Neal, 2011). Furthermore, Hu et al. (2024) apply GFlowNets (Bengio et al., 2021) to fine-tune language models for solving posterior inference problems, which can be considered as sampling in proportion to an intractable Gibbs distribution. In our work, we instead aim to use MCMC for MT to sample translations in proportion to a sentence-level evaluation metric.

Diverse Decoding for Machine Translation. Variants of beam search (Cho, 2016; Vijayakumar et al., 2017; Kulikov et al., 2019; Tam, 2020) have been proposed to produce a diverse set of translations using diversity-promoting objectives. However, the increased computation cost with the model size and beam width makes it infeasible and impractical to use with LLMs and it fails to yield consistent improvement over ancestral with an increase in beam width (Stahlberg and Byrne, 2019; Eikema and Aziz, 2020; Pang et al., 2024). Quality-aware decoding approaches on the other hand are almost always used to generate a single best hypothesis. Concurrently, Jinnai et al. (2024) add diversity promoting objective to MBR decoding to generate a set of high-quality diverse candidates. However, their approach only allows for the selection of hypotheses from a predefined set. We further note that we do not *directly* promote diversity—rather, diverse translations are a side product of efficiently exploring multiple high-quality regions from the model’s distribution.

7 Conclusion

We propose a new decoding approach for MT, *Quality-Aware Metropolis-Hastings* (QUEST) sampling based on MCMC that enables the generation of hypotheses in proportion to an automatic QE metric. We present a simple and novel proposal distribution that satisfies the constraints imposed by the Metropolis-Hastings algorithm. Our experiments on four language directions and two strong decoder-only language models show the efficacy of our approach over baselines. We further show that our

approach results in samples from underexplored high-density regions and that the average quality continues to improve as the Markov chain size increases.

8 Limitations and Broad Impact

QUEST requires generating many samples to reach high-density regions sequentially from an LLM for each input prompt, which can be computationally expensive for time-critical applications. Additionally, the required number of steps may vary depending on the input and the quality of the initial hypothesis. Furthermore, our proposal distribution only modifies the sentence suffix, which becomes restrictive once the chain reaches a high-quality region. As at this point, only minor changes to the hypothesis are accepted, slowing the mixing process. Addressing these limitations and extending this approach to other NLP tasks are avenues for future work.

Furthermore, we leverage recent advances in QE methods for MT and integrate them directly in the generation process of LLMs, which can potentially reduce the errors generated by these systems. However, despite the high correlations of evaluation metrics with human judgments, they are sometimes hard to interpret and occasionally fail to detect simple mistakes such as incorrect translations of numbers or entities (Amrhein and Sennrich, 2022). In such cases, sampling from the Gibbs distributions induced by these metrics might increase the chances of sampling those erroneous translations. We believe these risks will be mitigated as better metrics are constantly being developed—our method, being agnostic to the specific quality metric, will directly benefit from it. In addition, since QUEST supports any Gibbs distribution, it can also incorporate multiple QE models or additional checks which can rule out problematic samples by assigning them a very low QE score.

Acknowledgments

We thank Ben Peters, Marcos Treviso, and Sergey Troshin for their helpful and constructive feedback on the initial versions of the paper. Part of this work was supported by the EU’s Horizon Europe Research and Innovation Actions (UTTER, contract 101070631), by the project DECOLLAGE (ERC-2022-CoG 101088763), by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (Center for Responsible AI), and by Fundação para a Ciéncia e Tecnologia through contract UIDB/50008/2020.

References

Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024. Tower: An open multilingual large language model for translation-related tasks.

Afra Amini, Li Du, and Ryan Cotterell. 2023. Structured voronoi sampling.

Chantal Amrhein and Rico Sennrich. 2022. Identifying weaknesses in machine translation metrics through minimum Bayes risk decoding: A case study for COMET. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1125–1141, Online only. Association for Computational Linguistics.

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. 2021. Flow network based generative models for non-iterative diverse candidate generation.

Mathias Berglund, Tapani Raiko, Mikko Honkala, Leo Kärrkäinen, Akos Vetek, and Juha Karhunen. 2015. Bidirectional recurrent neural networks as generative models - reconstructing gaps in time series.

Frederic Blain, Chrysoula Zerva, Ricardo Rei, Nuno M. Guerreiro, Diptesh Kanodia, José G. C. de Souza, Beatriz Silva, Tânia Vaz, Yan Jingxuan, Fatemeh Azadi, Constantin Orasan, and André Martins. 2023. Findings of the WMT 2023 shared task on quality estimation. In *Proceedings of the Eighth Conference on Machine Translation*, pages 629–653, Singapore. Association for Computational Linguistics.

Kyunghyun Cho. 2016. Noisy parallel approximate decoding for conditional recurrent language model. *arXiv preprint arXiv:1605.03835*.

Markus Dreyer and Daniel Marcu. 2012. Hyter: Meaning-equivalent semantics for translation evaluation. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 162–171.

Li Du, Afra Amini, Lucas Torroba Hennigen, Xinyan Velocity Yu, Jason Eisner, Holden Lee, and Ryan Cotterell. 2023. Principled gradient-based markov chain monte carlo for text generation.

Simon Duane, A.D. Kennedy, Brian J. Pendleton, and Duncan Roweth. 1987. Hybrid monte carlo. *Physics Letters B*, 195(2):216–222.

Bryan Eikema. 2024. The effect of generalisation on the inadequacy of the mode. In *Proceedings of the 1st Workshop on Uncertainty-Aware NLP (UncertaiNLP 2024)*, pages 87–92, St Julians, Malta. 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.

António Farinhas, José de Souza, and Andre Martins. 2023. An empirical study of translation hypothesis ensembling with large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11956–11970, Singapore. Association for Computational Linguistics.

Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022a. Quality-aware decoding for neural machine translation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1396–1412, Seattle, United States. Association for Computational Linguistics.

Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and André F. T. Martins. 2022b. Quality-aware decoding for neural machine translation.

Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José G. C. de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, and André F. T. Martins. 2023. Bridging the Gap: A Survey on Integrating (Human) Feedback for Natural Language Generation. *Transactions of the Association for Computational Linguistics*, 11:1643–1668.

Jarad Forristal, Niloofar Mireshghallah, Greg Durrett, and Taylor Berg-Kirkpatrick. 2023. A block metropolis-hastings sampler for controllable energy-based text generation.

Markus Freitag, Behrooz Ghorbani, and Patrick Fernandes. 2023a. Epsilon sampling rocks: Investigating sampling strategies for minimum Bayes risk decoding for machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9198–9209, Singapore. Association for Computational Linguistics.

Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022a. High quality rather than high model probability: Minimum Bayes risk decoding with neural metrics. *Transactions of the Association for Computational Linguistics*, 10:811–825.

Markus Freitag, Nitika Mathur, Chi-ku Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Tom Kocmi, Frederic Blain, Daniel Deutsch, Craig Stewart, Chrysoula Zerva, Sheila Castilho, Alon Lavie, and George Foster. 2023b. Results of WMT23 metrics shared task: Metrics might be guilty but references are not innocent. In *Proceedings of the Eighth Conference on Machine Translation*, pages 578–628, Singapore. Association for Computational Linguistics.

Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-ku Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022b. Results of WMT22 metrics shared task: Stop using BLEU – neural metrics are better and more robust. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Dongyoung Go, Tomasz Korbak, Germán Kruszewski, Jos Rozen, Nahyeon Ryu, and Marc Dymetman. 2023. Aligning language models with preferences through f-divergence minimization.

Kartik Goyal, Chris Dyer, and Taylor Berg-Kirkpatrick. 2022. Exposing the implicit energy networks behind masked language models via metropolis–hastings.

Naman Goyal, Jingfei Du, Myle Ott, Giri Anantharaman, and Alexis Conneau. 2021. Larger-scale transformers for multilingual masked language modeling. In *Proceedings of the 6th Workshop on Representation Learning for NLP (Rep4NLP-2021)*, pages 29–33, Online. Association for Computational Linguistics.

Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. Continuous measurement scales in human evaluation of machine translation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 33–41, Sofia, Bulgaria. Association for Computational Linguistics.

David R Grimm, Yu-Jin Lee, Katherine Hu, Longsha Liu, Omar Garcia, Karthik Balakrishnan, and Noel F Ayoub. 2024. The utility of chatgpt as a generative medical translator. *European Archives of Oto-Rhino-Laryngology*, pages 1–5.

Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2023a. xcomet: Transparent machine translation evaluation through fine-grained error detection.

Nuno Miguel Guerreiro, Duarte M Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André FT Martins. 2023b. Hallucinations in large multilingual translation models. *Transactions of the Association for Computational Linguistics*, 11.

Zhiwei He, Xing Wang, Wenxiang Jiao, Zhuosheng Zhang, Rui Wang, Shuming Shi, and Zhaopeng Tu. 2024. Improving machine translation with human feedback: An exploration of quality estimation as a reward model. *arXiv preprint arXiv:2401.12873*.

Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. How good are gpt models at machine translation? a comprehensive evaluation. *arXiv preprint arXiv:2302.09210*.

Edward J. Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, and Nikolay Malkin. 2024. Amortizing intractable inference in large language models.

Yuu Jinnai, Ukyo Honda, Tetsuro Morimura, and Peinan Zhang. 2024. Generating diverse and high-quality texts by minimum bayes risk decoding. *arXiv preprint arXiv:2401.05054*.

Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. 2023. A survey of reinforcement learning from human feedback. *arXiv preprint arXiv:2312.14925*.

Elaine C Khoong, Eric Steinbrook, Cortlyn Brown, and Alicia Fernandez. 2019. Assessing the use of google translate for spanish and chinese translations of emergency department discharge instructions. *JAMA internal medicine*, 179(4):580–582.

Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Philipp Koehn, Benjamin Marie, Christof Monz, Makoto Morishita, Kenton Murray, Makoto Nagata, Toshiaki Nakazawa, Martin Popel, Maja Popović, and Mariya Shmatova. 2023. Findings of the 2023 conference on machine translation (WMT23): LLMs are here but not quite there yet. In *Proceedings of the Eighth Conference on Machine Translation*, pages 1–42, Singapore. Association for Computational Linguistics.

Tom Kocmi, Vilém Zouhar, Christian Federmann, and Matt Post. 2024. Navigating the metrics maze: Reconciling score magnitudes and accuracies. *arXiv preprint arXiv:2401.06760*.

Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.

Tomasz Korbak, Hady Elsahar, Germán Kruszewski, and Marc Dymetman. 2022a. On reinforcement learning and distribution matching for fine-tuning language models with no catastrophic forgetting.

Tomasz Korbak, Ethan Perez, and Christopher L Buckley. 2022b. RL with kl penalties is better viewed as bayesian inference.

Ilia Kulikov, Alexander Miller, Kyunghyun Cho, and Jason Weston. 2019. Importance of search and evaluation strategies in neural dialogue modeling. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 76–87, Tokyo, Japan. Association for Computational Linguistics.

Sachin Kumar, Biswajit Paria, and Yulia Tsvetkov. 2022. Gradient-based constrained sampling from language models.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Stephen Mayhew, Clinton Bicknell, Chris Brust, Bill McDowell, Will Monroe, and Burr Settles. 2020. Simultaneous translation and paraphrase for language education. In *Proceedings of the Fourth Workshop on Neural Generation and Translation*, pages 232–243, Online. Association for Computational Linguistics.

Nicholas C. Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, and A. H. Teller. 1953. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092.

Ning Miao, Hao Zhou, Lili Mou, Rui Yan, and Lei Li. 2018. CGMH: constrained sentence generation by metropolis-hastings sampling. *CoRR*, abs/1811.10996.

Fatemehsadat Mireshghallah, Kartik Goyal, and Taylor Berg-Kirkpatrick. 2022. Mix and match: Learning-free controllable text generationusing energy language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 401–415, Dublin, Ireland. Association for Computational Linguistics.

Radford M. Neal. 2011. Probabilistic inference using markov chain monte carlo methods.

Eugene Albert Nida. 1964. *Toward a science of translating: with special reference to principles and procedures involved in Bible translating*. Brill Archive.

Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018. Analyzing uncertainty in neural machine translation. In *International Conference on Machine Learning*, pages 3956–3965. PMLR.

Jianhui Pang, Fanghua Ye, Longyue Wang, Dian Yu, Derek F Wong, Shuming Shi, and Zhaopeng Tu. 2024. Salute the classic: Revisiting challenges of machine translation in the age of large language models. *arXiv preprint arXiv:2401.08350*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. 2019. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning.

Stefano Perrella, Lorenzo Proietti, Alessandro Scirè, Niccolò Campolungo, and Roberto Navigli. 2022. MaTESe: Machine translation evaluation as a sequence tagging problem. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 569–577, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ben Peters and André F. T. Martins. 2021. Smoothing and shrinking the sparse Seq2Seq search space. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2642–2654, Online. Association for Computational Linguistics.

Jan Peters and Stefan Schaal. 2007. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th International Conference on Machine Learning*, ICML '07, page 745–750, New York, NY, USA. Association for Computing Machinery.

Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. 2022. Cold decoding: Energy-based constrained text generation with langevin dynamics.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.

Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. COMET-22: Unbabel-IST 2022 submission for the metrics shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ricardo Rei, Nuno M. Guerreiro, JosÃ© Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, José G. C. de Souza, and André Martins. 2023. Scaling up CometKiwi: Unbabel-IST 2023 submission for the quality estimation shared task. In *Proceedings of the Eighth Conference on Machine Translation*, pages 841–848, Singapore. Association for Computational Linguistics.

Lucía Sanz-Valdivieso and Belén López-Arroyo. 2023. Google translate vs. chatgpt: Can non-language professionals trust them for specialized translation? In *International Conference Human-informed Translation and Interpreting Technology (HiT-IT 2023)*, pages 97–107.

Libin Shen, Anoop Sarkar, and Franz Josef Och. 2004. Discriminative reranking for 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 177–184, Boston, Massachusetts, USA. Association for Computational Linguistics.

Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Minimum risk training for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1683–1692, Berlin, Germany. Association for Computational Linguistics.

Tianxiao Shen, Myle Ott, Michael Auli, and Marc’Aurelio Ranzato. 2019. Mixture models for diverse machine translation: Tricks of the trade. In *International conference on machine learning*, pages 5719–5728. PMLR.

Yiqiu Shen, Laura Heacock, Jonathan Elias, Keith D Hentel, Beatriu Reig, George Shih, and Linda Moy. 2023. Chatgpt and other large language models are double-edged swords.

Prasann Singhal, Jiacheng Xu, Xi Ye, and Greg Durrett. 2023. Eel: Efficiently encoding lattices for reranking.

Felix Stahlberg and Bill Byrne. 2019. On NMT search errors and model errors: Cat got your tongue? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3356–3362, Hong Kong, China. Association for Computational Linguistics.

Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize from human feedback. In *NeurIPS*.

Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2022. Learning to summarize from human feedback.

Jinyue Su, Jiacheng Xu, Xipeng Qiu, and Xuanjing Huang. 2018. Incorporating discriminator in sentence generation: a gibbs sampling method.

Breena R Taira, Vanessa Kreger, Aristides Orue, and Lisa C Diamond. 2021. A pragmatic assessment of google translate for emergency department instructions. *Journal of General Internal Medicine*, 36(11):3361–3365.

Yik-Cheung Tam. 2020. Cluster-based beam search for pointer-generator chatbot grounded by knowledge. *Computer Speech & Language*, 64:101094.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Christian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madijan Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models.

Giorgos Vernikos and Andrei Popescu-Belis. 2024. Don't rank, combine! combining machine translation hypotheses using quality estimation. *arXiv preprint arXiv:2401.06688*.

Ashwin K Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2017. Diverse beam search: Decoding diverse solutions from neural sequence models.

Alex Wang and Kyunghyun Cho. 2019. BERT has a mouth, and it must speak: BERT as a markov random field language model. *CoRR*, abs/1902.04094.

John Wieting, Taylor Berg-Kirkpatrick, Kevin Gimpel, and Graham Neubig. 2019. Beyond BLEU:training neural machine translation with semantic similarity. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4344–4355, Florence, Italy. Association for Computational Linguistics.

Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. 2024a. A paradigm shift in machine translation: Boosting translation performance of large language models.

Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024b. Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation. *arXiv preprint arXiv:2401.08417*.

Takateru Yamakoshi, Thomas Griffiths, and Robert Hawkins. 2022. Probing BERT's priors with serial reproduction chains. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3977–3992, Dublin, Ireland. Association for Computational Linguistics.

Yiming Yan, Tao Wang, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Mingxuan Wang. 2023. BLEURT has universal translations: An analysis of automatic metrics by minimum risk training. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5428–5443, Toronto, Canada. Association for Computational Linguistics.

Maosen Zhang, Nan Jiang, Lei Li, and Yexiang Xue. 2020. Language generation via combinatorial constraint satisfaction: A tree search enhanced Monte-Carlo approach. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1286–1298, Online. Association for Computational Linguistics.

A Connection to RLHF derivation

If we take the target distribution to have the following form:

$$\pi(y|x) = \frac{1}{Z(x)} p_{LM}(y|x) \exp\left(\frac{r(x,y)}{\beta}\right), \quad (14)$$

then the likelihood ratio becomes :

$$\frac{\pi(y|x)}{\pi(y^t|x)} = \exp\left(\frac{r(x,y) - r(x,y^t)}{\beta}\right) \frac{p_{LM}(y_{\geq i}|y_{<i}^t, x)}{p_{LM}(y_{\geq i}^t|y_{<i}^t, x)}. \quad (15)$$

Note that the target log ratio contains the inverse of the probability of returning, *i.e.*:

$$\frac{q(y^t|y, x, i)}{q(y|y^t, x, i)} = \frac{p_{LM}(y_{\geq i}^t|y_{<i}^t, x)}{p_{LM}(y_{\geq i}|y_{<i}^t, x)}, \quad (16)$$

this allows simplifying the criterion as follows:

$$\alpha_\beta(y, y^t) = \min \left\{ 1, \exp\left(\frac{r(x,y) - r(x,y^t)}{\beta}\right) \frac{q(i|n)}{q(i|n^t)} \right\}. \quad (17)$$

B Prompts used for MT

For both the ALMA-7B and TOWER-7B we adhere to the prompt format recommended for translation in the original papers as shown below:

ALMA-7B:

Translate this sentence from {source language} to {target language}.
{source language} Source: {source sentence}.
{target language} Translation:

TOWER-7B

<|im_start|>user
Translate the following {source_language} source text to {target_language}:
{source_language}: {source_sentence}
{target_language}: <|im_end|>
<|im_start|>assistant

C Additional Results

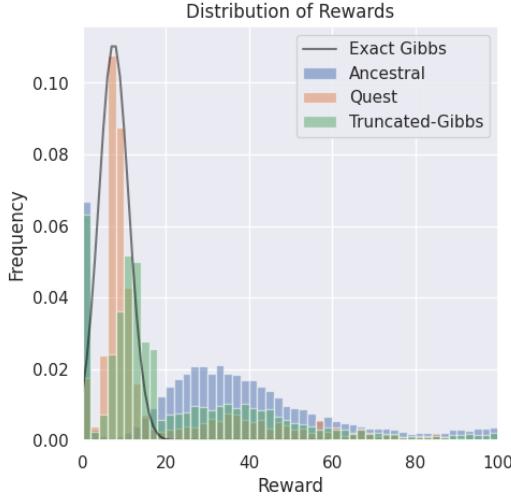
C.1 Toy problem: Validating the Approach

To validate that our approach works, we test QUEST on a controlled setting, a toy summarization problem, where the ground truth reward is known. We consider the following reward function:

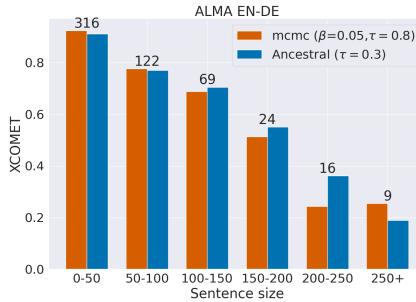
$$r(x, y) := \text{logit}(\text{clamp}(\mathcal{N}(|y|; \mu = 7.5, \sigma^2 = 3.75^2))), \quad (18)$$

where $\text{clamp}(x) = \max(10^{-2}, \min(x, 1 - 10^{-2}))$.

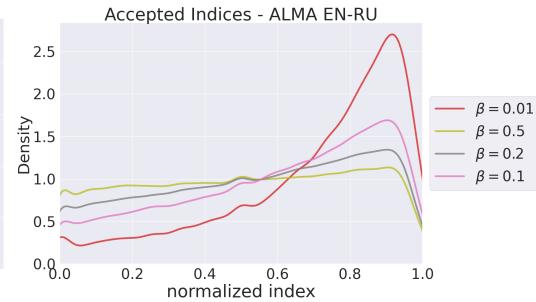
We use GPT2 (Radford et al., 2019) fine-tuned on the Reddit dataset Stiennon et al. (2020) to generate summaries for 128 examples randomly sampled from the test split. We compare the distribution of rewards obtained by different sampling techniques (Ancestral, QUEST and Truncated-Gibbs) to the ground truth reward in Figure 4a. For Truncated-Gibbs sampling, we estimate the partition function using the ancestral samples and resample them based on the probability distribution induced by the rewards. Unlike ancestral sampling which results in samples that are far from the high-reward regions, our sampling strategy, QUEST, results in the best approximation of the ground truth distribution. This simplified reward task serves as an useful example to show the efficacy of our approach over alternatives.



(a) Distribution of rewards for sampled hypotheses for the toy summarization problem.



(a) xCOMET-XL by source length.



(b) Distribution over Accepted indices.

C.2 MCMC results in better outputs for shorter segments.

We show the xCOMET-XL scores on WMT23 English-German pairs bucketed by the length of the source text. In general, the quality of samples degrades with increased source length when decoding via MCMC as shown in Figure 5a. One possible factor could be the uniform index-selection exploration strategy where the proposal might require more steps as the sentence length increases. Another factor could be the limitations of the reward function itself for selecting high-quality translations for longer sequences. As these quality estimation metrics have not been trained on longer sequences, due to a distribution shift, they might not be representative of true quality for these output lengths. We believe the latter option could be the more significant factor as the reward optimized still improves regardless of the sentence length (see Figure 7).

C.3 Accepted proposals are skewed towards the end of the sentence.

Figure 5b illustrates how the distribution of accepted indices changes with decreasing β : lower β values tighten acceptance criteria, favoring states with higher rewards and a greater likelihood of returning to the previous state. Consequently, transitions tend to yield minor alterations, typically at the end of the sentence. Altering the beginning necessitates sampling small indices values, risking a complete sentence change. This plot highlights the current limitations of our proposal, especially the downstream implications on the diversity of prefixes we get for a particular motivating the use of parallel chains and potential avenues for improvement for future work.

C.4 Comparing RLHF-QUEST vs QUEST

We compare QUEST with the alternative discussed in Section 3.3 where the target distribution in Eq. 3 is replaced with Eq. 12, referred to as “RLHF-QUEST”. We compare sampling using the two target distributions in Figure 6. Our results indicate that both approaches result in comparable translation quality, with QUEST resulting in slightly higher diversity on average than RLHF-QUEST. We leave the detailed exploration of these two methods to future work.

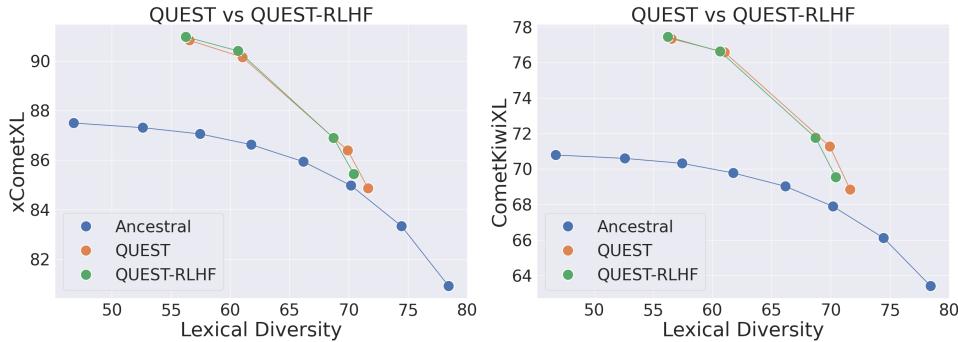


Figure 6: Average Quality by xCOMET-XL (left) and COMETKIWI-XL on English-Russian dataset using TOWER-7B

C.5 QE Results

We report the quality as measured by the metric used in sampling, *i.e.*, COMETKIWI-XL in Figure 7. QUEST results in higher quality across the board compared to ancestral sampling.

C.6 Additional Language Pairs

We further expanded our evaluation to four additional language pairs, as shown in Figure 8: WMT23 English-Chinese (EN→ZH, high-resource), WMT23 English-Czech (EN→CS, medium-resource), WMT22 English-Icelandic (EN→IS), and Icelandic-English (IS→EN, low-resource), using ALMA-7B. We did not include TOWER in these experiments because it was trained on the WMT22 test sets. Across all language pairs, QUEST consistently outperforms ancestral sampling, providing better quality-diversity trade-offs.

C.7 COMET22 Results

We report the quality as measured by COMET22 (Rei et al., 2022) in Figure 9. QUEST results in higher quality across the board compared to ancestral sampling.

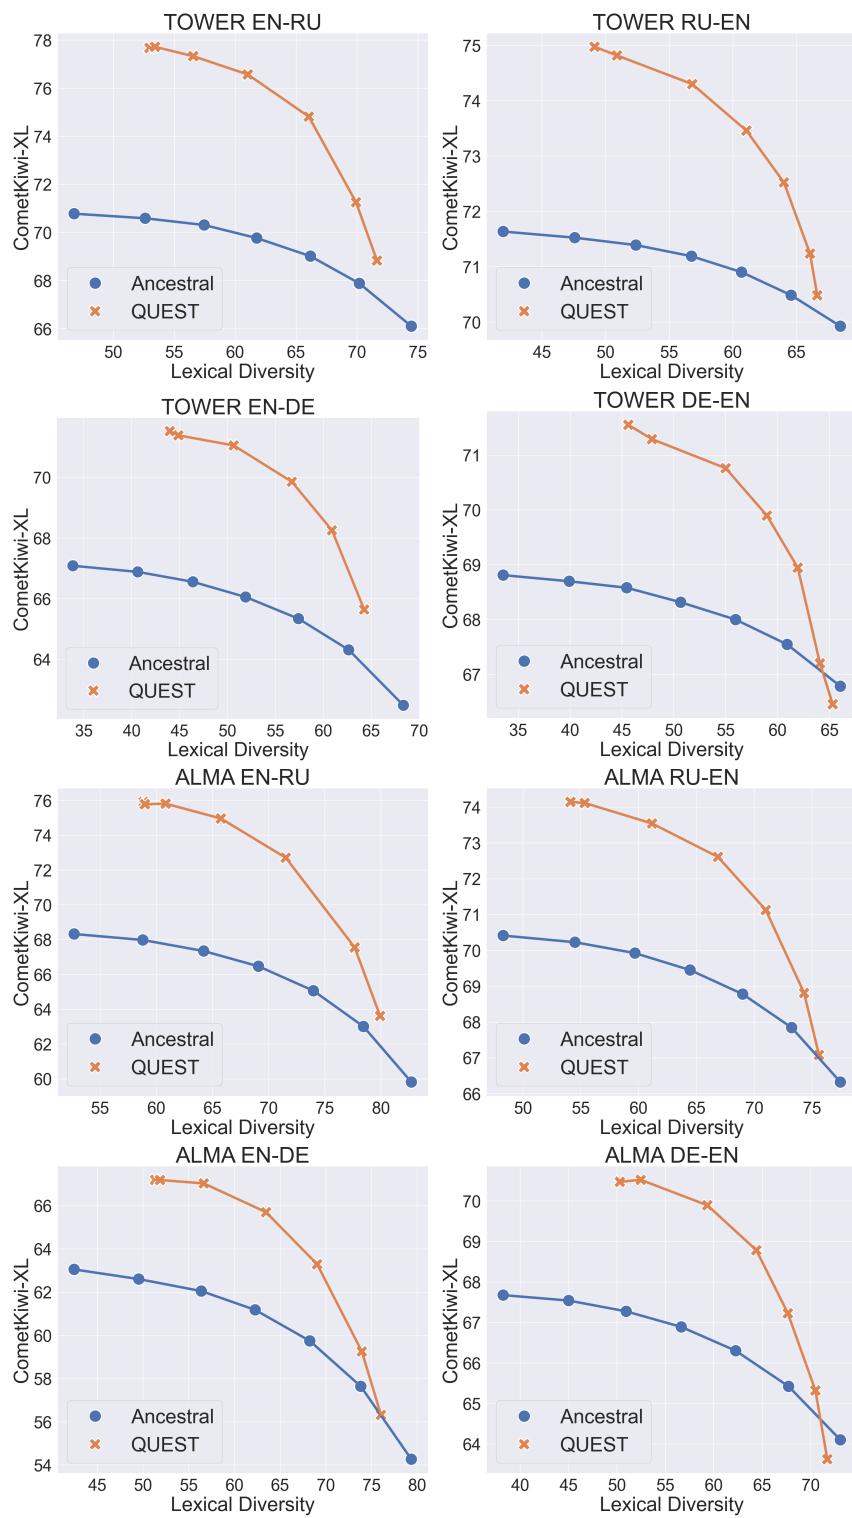


Figure 7: Average quality (COMETKIWI-XL) vs. diversity (PAIRWISE-BLEU) on WMT23 datasets. Different points represent different hyperparameter values.

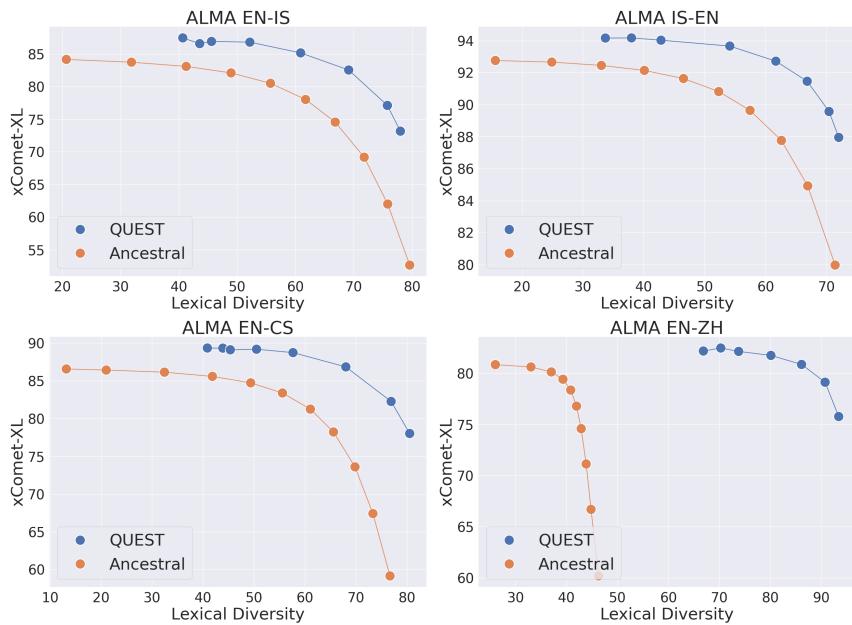


Figure 8: Average quality (xCOMET-XL) vs. diversity (PAIRWISE-BLEU) on additional LPs from WMT23 and WMT22. Different points represent different hyperparameter values.

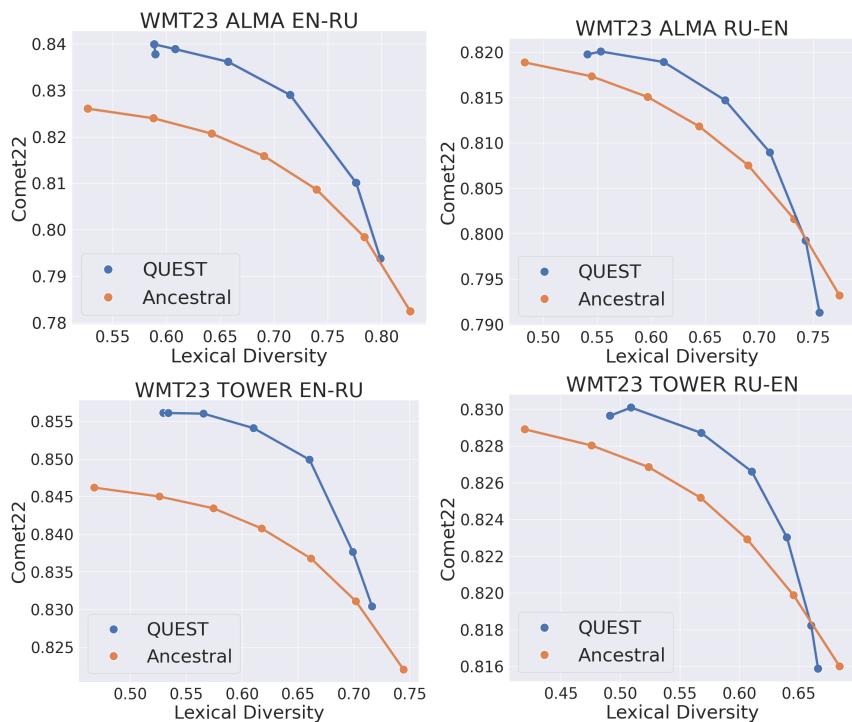


Figure 9: Average quality (COMET22) vs. diversity (PAIRWISE-BLEU) on EN→RU and RU→EN. Different points represent different hyperparameter values.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: We present the main contributions and results in the last two paragraphs of the appendix which are supported by experiments detailed in Section 4 and findings presented in Section 5.2.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We present some of the limitations of our work in the ablation analysis presented in Section 5.2, additionally, we list other limitations in Section 8.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide all necessary details to reproduce the experiments in Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We use publicly available models and datasets in our experiments. We release the code to replicate our experiments at <https://www.questdecoding.com/>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: See Section 4. Our proposed method works at inference time. We report the test data used and all the hyperparameters in the beginning of Section 4 in dedicated paragraphs.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: We show all our results aggregated over a sample size of 500-2000 source sentences across multiple language pairs and models. Furthermore, we report statistical significance for results using the optimized QE metric in Section 5 and Appendix C.5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We discuss the compute used and run time in a dedicated paragraph in Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: Our research conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: Please refer to Section 8.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work poses no such risks. We utilize existing models and resources to improve translation quality.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We provide all details with citations for the WMT23 dataset and the models (ALMA-7b, Tower-7b) in dedicated paragraphs in Section 4.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We release the code at <https://www.questdecoding.com/>.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not conduct any human study.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not conduct any human study.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.