

Improving NMT Models by Retrofitting Quality Estimators into Trainable Energy Loss

Gahyun Yoo¹ Jay-Yoon Lee^{2*}

¹Seoul National University

²Graduate School of Data Science, Seoul National University

{padme0421, lee.jayyoon}@snu.ac.kr

Abstract

Reinforcement learning has shown great promise in aligning language models with human preferences in a variety of text generation tasks, including machine translation. For translation tasks, rewards can easily be obtained from quality estimation (QE) models which can generate rewards for unlabeled data. Despite its usefulness, reinforcement learning cannot exploit the gradients with respect to the QE score. We propose QE-EBM, a method of employing quality estimators as trainable loss networks that can directly backpropagate to the NMT model. We examine our method on several low and high resource target languages with English as the source language. QE-EBM outperforms strong baselines such as REINFORCE and proximal policy optimization (PPO) as well as supervised fine-tuning for all target languages, especially low-resource target languages. Most notably, for English-to-Mongolian translation, our method achieves improvements of 2.5 BLEU, 7.1 COMET-KIWI, 5.3 COMET, and 6.4 XCOMET relative to the supervised baseline.

1 Introduction

Reinforcement learning with human feedback (RLHF) has been used to successfully align language models with human preferences. By assigning a high reward to responses that exhibit certain characteristics, such as lack of harmful content, we can steer the model towards generating these types of responses in various types of text generation tasks. For machine translation, reinforcement learning has been used to improve general translation quality by using quality estimation (QE) models trained on human feedback as the reward model (He et al., 2024; Ramamurthy et al., 2023; Ramos et al., 2023).

Quality estimation is the task of assigning a translation score to a provided source and target-

side prediction pair, without referring to a gold reference. Over the last few years, quality estimation has developed rapidly, with the correlation of model ratings and human ratings of translations rising to a level at par with reference-based translation metrics (Zerva et al., 2022; Freitag et al., 2023). Quality estimators in the widely used COMET translation evaluation framework, such as XCOMET-QE and COMET-KIWI, use a cross-lingual encoder with a prediction head that is regressed on datasets composed of translation pairs and human ratings (Rei et al., 2022; Guerreiro et al., 2023).

Although reinforcement learning methods can use QE scores as rewards to produce better translations compared to vanilla fine-tuning, they fail to exploit a crucial piece of information QE scores contain: gradients. More specifically, they lose the gradients of the translation model’s parameters with respect to the QE score by treating the score as a scalar value. Energy-based training, on the other hand, can take advantage of these gradients by treating the score as a form of energy loss that can be backpropagated.

We propose QE-EBM, a method of training a translation model by using the QE score of the model-generated translations as an energy loss. Compared to RL methods, energy-based training enables more informative distillation of knowledge from the QE model to the translation model. This method is especially useful for boosting the translation quality of low resource target languages, that do not have sufficient bilingual corpora paired with English sentences. With only the translations from the translation model, the QE model can generate scores that are a rich source of knowledge the translation models can in turn exploit through the energy loss.

There are two variants of our method: QE-STATIC and QE-DYNAMIC. In QE-DYNAMIC, the energy network’s parameters are updated via

*Corresponding author

contrastive learning, while in QE-STATIC, they remain fixed. The motivation for fine-tuning the energy model in QE-DYNAMIC is so that the QE model remains relevant even as the NMT model produces increasingly more natural translations similar in quality to references.

We test our method on both high resource and low-resource target languages with English as the source language. Our methods outperform supervised fine-tuning as well as REINFORCE and proximal policy optimization (PPO), showing an increase of 2.5 BLEU, 7.1 COMET-KIWI, 5.3 COMET, and 6.4 XCOMET for English-to-Mongolian translation (Sutton et al., 1999; Schulman et al., 2017).

2 Related Work

2.1 Structured Energy Network as a Loss (SEAL)

SEAL (Lee et al., 2022) was proposed as a training framework to use energy networks as trainable loss functions for structured prediction. There are two versions of the algorithm: SEAL-STATIC and SEAL-DYNAMIC. In both versions, the task net, the neural network responsible for performing a task (i.e. prediction), is trained through a weighted sum of cross-entropy loss and energy loss. The difference lies in whether the loss net, the neural network that acts as the loss function, is updated. In SEAL-DYNAMIC, the loss net is fine-tuned with contrastive loss before each task net update so that it is better suited to predicting the energy landscape for the data samples at each step. In SEAL-STATIC, the weights of the loss net remain fixed. SEAL-DYNAMIC outperforms SPENs (Structured Prediction Energy Networks), which refers to using energy models as inference networks, as well as SEAL-STATIC in image segmentation and semantic role labeling.

Our work can be viewed as an adaptation of SEAL to the task of translation, retaining the energy loss and iterative update algorithm while adding the following contributions: 1) we extend it to semi-supervised learning by incorporating unlabeled data, and 2) as the energy loss net, we plug in a quality estimation model which has already been fine-tuned with human annotations of translation pairs. This guides the model towards human preferences with minimal additional training and no extra data about human preference.

2.2 Energy Based Models for Text Generation

Existing research on using energy based models for text generation has formulated energy models as *residual* energy based models (Deng et al., 2020; He et al., 2021). Low energy samples are drawn in two steps: sampling many generations from a frozen language model, then importance sampling through the energy model. Similarly, Bhattacharyya et al. (2021) proposes training an energy based model with a margin based loss proportional to the difference in BLEU and using it to rerank generations during inference.

Our method has an advantage over these works in inference speed and distillation. These residual EBM methods introduces latency during inference, and do not pass the energy model’s knowledge to the base language model. In contrast, we distill the knowledge of the energy model to the base NMT model during training so that we can achieve improvements even with simple greedy generation. QE-EBM is also orthogonal to these methods in that we can also use the energy network to rerank samples during inference, although we report results of greedy generation for the evaluation sets, since our method does not require sampling from the energy model during inference.

2.3 Using Quality Estimators in NMT Training

With the rapid development of quality estimation models, a number of recent works have incorporated quality estimators into the training of translation models. Ramos et al. (2023) improves translation abilities of T5-based models by applying PPO with COMET-QE as the reward model, finding that filtering training data based on the reward model is important. Using COMET-QE also, He et al. (2024) improves the translation capabilities of Llama and NLLB through RAFT (Reward Ranked Fine-tuning), which involves generating multiple candidates, ranking them with the reward model, and learning from the best sample (Dong et al., 2023; Touvron et al., 2023; Team et al., 2022).

Gulcehre et al. (2023) introduces an efficient grow-batch offline reinforcement learning algorithm consisting of alternating Grow and Improve steps. During the Grow step, samples are generated offline, and are added to the supervised batch after applying a filter based on the reward model.

During the Improve step, the translation model is fine-tuned on the grown batch.

The aforementioned works of incorporating quality estimators in NMT training use a form of reinforcement learning. To the best of our knowledge, we are the first to use a quality estimation model as an energy based model and update the NMT model with gradients with respect to the score to enhance machine translation performance.

3 Method

A neural machine translation system models the conditional distribution $P_\phi(\mathbf{y}|\mathbf{x})$ of a target sentence \mathbf{y} given a source sentence \mathbf{x} . NMT models are trained to maximize this probability by minimizing the cross-entropy loss function,

$$L_{CE}(\mathbf{x}, \mathbf{y}) = - \sum_{t=1}^{T_y} \log P_\phi(\mathbf{y}|\mathbf{x}) \quad (1)$$

Energy-based models (EBMs) are parameterized models that output a scalar value for each input (Lecun et al., 2006). In this paper, we investigate the use of quality estimators as EBMs, specifically COMET-KIWI (Rei et al., 2022). Since COMET-KIWI outputs a scalar value that represents the quality of translation for each pair of source and target sentence, the model can be represented as an energy function $E_\theta(\mathbf{x}, \mathbf{y}) = -s(\mathbf{x}, \mathbf{y})$ where s is the COMET-KIWI score.

To leverage the power of EBMs during training, we train the NMT model in a multi-task setup, where the loss consists of a standard cross-entropy loss and the energy term (called energy loss hereafter). For each batch consisting of B_l labeled samples $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{B_l}$ and B_u unlabeled samples $\{\mathbf{x}_{u_i}\}_{i=1}^{B_u}$ with N translations $\{\hat{\mathbf{y}}_{u_i}^{(j)}\}_{j=1}^N$ sampled from the NMT model for each unlabeled source sentence \mathbf{x}_{u_i} , the loss for the NMT model can be expressed as

$$L_{NMT} = \alpha \cdot \frac{1}{B_l} \sum_{i=1}^{B_l} L_{CE}(\mathbf{x}_i, \mathbf{y}_i) + \beta \cdot \frac{1}{B_u} \sum_{i=1}^{B_u} \frac{1}{N} \sum_{j=1}^N E_\theta(\mathbf{x}_{u_i}, \hat{\mathbf{y}}_{u_i}^{(j)}) \quad (2)$$

α is a hyper-parameter that controls the magnitude of the cross entropy gradient, and is decreased throughout training. β is a hyper-parameter that controls the magnitude of the energy loss gradient,

and is increased throughout training. More details about hyperparameters can be found in Appendix A.

This hybrid loss function is one key difference between QE-EBM and previous attempts to incorporate QE into translation model fine-tuning described in Section 2.3. The supervised fine-tuning and feedback training occur simultaneously, reducing the need to move back and forth between the two stages or decide on arbitrary hyperparameters such as the epoch number of each stage in each cycle.

To update the NMT model with the energy loss, we replace the regular softmax operation for the output logits from the last hidden layer of the NMT model’s decoder with the straight through estimator (STE) (Bengio et al., 2013), following Tu et al. (2020) which demonstrated success using STE to distill knowledge from an autoregressive energy network to a non-autoregressive inference network.

The exact algorithm for QE-DYNAMIC is given in Algorithm 1. The energy network’s parameters θ are updated using the labeled batch and K translations sampled for each source sentence in the batch, before each update of the NMT model’s parameters ϕ . The noise contrastive estimation (NCE) loss used to train the energy network in QE-DYNAMIC is given in Equation (3) (Ma and Collins, 2018).

$$L_{E-NCE} = -\frac{1}{B_l} \sum_{i=1}^{B_l} [\log \sigma(\bar{s}(\mathbf{x}_i, \mathbf{y}_i)) + \sum_{j=1}^K \log(1 - \sigma(\bar{s}(\mathbf{x}_i, \hat{\mathbf{y}}_{p_i}^{(j)})))] \quad (3)$$

where $\bar{s}(\mathbf{x}, \mathbf{y}) = -E_\theta(\mathbf{x}, \mathbf{y}) - \log P_\phi(\mathbf{y}|\mathbf{x})$, and $\hat{\mathbf{y}}_{p_i}^{(j)} \sim P_\phi(\cdot|\mathbf{x}_i)$. σ is the sigmoid function.

Through this loss, the energy model learns to prefer the gold labels over the NMT model generations.

To improve the performance of QE-STATIC and QE-DYNAMIC, we apply two data sampling techniques to select labeled and unlabeled training batches.

- **Filter:** Labeled data is filtered based on COMET-KIWI score, before training begins. A labeled pair with higher COMET-KIWI score consists of a source and target sentence with higher alignment, and is likely to pro-

Algorithm 1 QE-DYNAMIC Algorithm

Require: $\{(\mathbf{x}_{l_i}, \mathbf{y}_{l_i})\}_{i=1}^{B_l}, \{\mathbf{x}_{u_i}\}_{i=1}^{B_u}$: Training Batch**Require:** θ : QE model parameters**Require:** ϕ : NMT model parameters**for** $i = 1 \rightarrow B_l$ **do** $S_{l_i} \leftarrow \{\hat{\mathbf{y}}_{p_i}^{(j)} | \hat{\mathbf{y}}_{p_i}^{(j)} \sim P_\phi(\cdot | \mathbf{x}_{l_i})\}_{j=1}^K$ **end for** $\theta \leftarrow \theta - \nabla_\theta L_{E-NCE}$ **for** $i = 1 \rightarrow B_u$ **do** $S_{u_i} \leftarrow \{\hat{\mathbf{y}}_{u_i}^{(j)} | \hat{\mathbf{y}}_{u_i}^{(j)} \sim P_\phi(\cdot | \mathbf{x}_{u_i})\}_{j=1}^N$ **end for** $\phi \leftarrow \phi - \nabla_\phi L_{NMT}$

▷ Step1: sample K translations for each \mathbf{x}_{l_i}

▷ Step2: update energy model

▷ Step3: sample N translations for \mathbf{x}_{u_i} ▷ Step4: update NMT model

vide more accurate guidance for the generation model.

- **NN:** We select the unlabeled batch to be paired with each labeled batch by retrieving the nearest neighbor based on embedding cosine similarity between the source sentences. This is expected to reduce the mismatch in gradient direction for the model parameters that could occur as a result of training simultaneously with labeled and unlabeled data.

These data sampling techniques introduce some computational overhead, but it is negligible compared to the overall training time. A breakdown of the computation time for data sampling techniques and a single training step for each algorithm is given in Appendix C.

4 Experiments

4.1 Models

NMT Model For the NMT model, we use MBART, a Transformer encoder-decoder model with roughly 610M parameters (Tang et al., 2021; Vaswani et al., 2017). We use the weights that are pretrained only on monolingual data and not finetuned on any parallel data, allowing us to compare regular supervised finetuning and energy-based training.

QE Model The COMET-KIWI quality estimation model consists of a feedforward estimator on top of a cross-lingual encoder. The weights of the COMET-KIWI encoder were initialized with the InfoXLM weights and finetuned along with the feed-forward estimator on translations paired with human ratings for the WMT22 Metrics Task (Chi et al., 2021). COMET-KIWI ranked first among

the reference-free metrics, and seventh among all the metrics in terms of correlation with human evaluation (Zerva et al., 2022).

4.2 Datasets

We use the IWSLT2017 English- $\{\text{German, Chinese}\}$ and ML50 English- $\{\text{Bengali, Azerbaijani, Mongolian, Marathi, Kazakh}\}$ translation datasets (Cettolo et al., 2017; Tang et al., 2020). In this paper, we focus on only English to X directions. We filter the training sets by total size and sentence length. For high resource language pairs, we select 50K sentence pairs randomly, and filter out sentences with more than 50 sub-word tokens in either the source sentence or reference translation. For low-resource language pairs, we only apply length-based filtering. We will refer to this processed training set as *preprocessed data pool* for convenience. We construct the labeled subset by randomly selecting 1/5 of the sentence pairs (source + target sentence) in the *preprocessed data pool*. For the unlabeled subset, we use all of the source sentences in the *preprocessed data pool*.

4.3 Baselines

To the best of our knowledge, there are no previous works that use quality estimators as trainable loss functions to train translation models. Therefore, in addition to the supervised baseline, we compare our method against REINFORCE and PPO as these reinforcement learning methods can also utilize quality estimators as a reward. Both RL baselines are trained in a multi-task setup as in QE-EBM, with the same cross-entropy loss for the labeled batch and a separate loss using the QE score for the unlabeled batch. The REINFORCE baseline uses the vanilla policy gradient. The reward

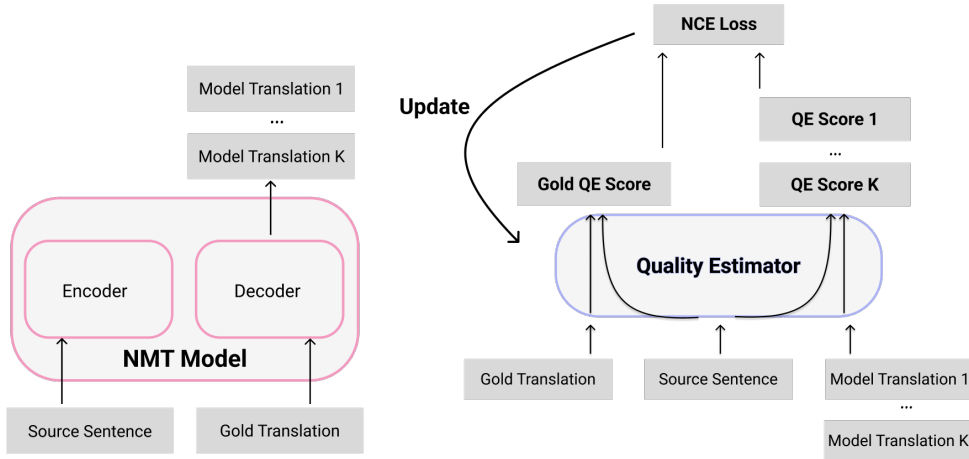


Figure 1: Diagram of QE model update process.

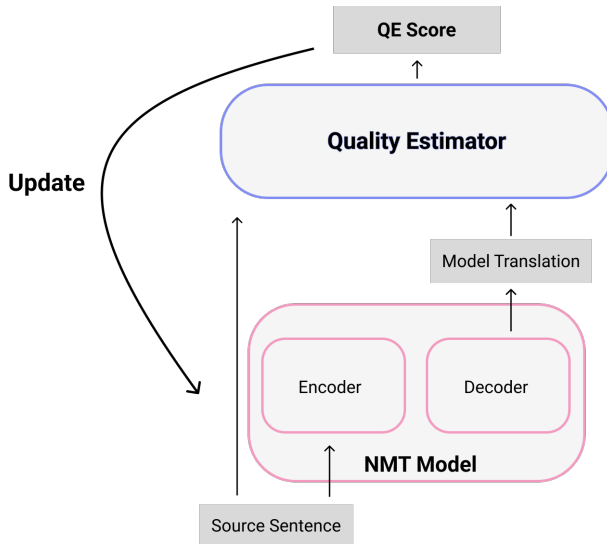


Figure 2: Diagram of NMT model update process with unlabeled samples.

is normalized through (1) scaling the reward between the running max and running min of the current epoch and (2) subtracting the running average of the current epoch. The PPO implementation is taken from TRL (von Werra et al., 2020). For a fair comparison across different approaches, we ensured an identical number of samples per source sentence in each learning algorithm.

4.4 Training and Evaluation

We use an ADAM optimizer to update the energy model parameters and NMT model parameters (Kingma and Ba, 2017). In QE-DYNAMIC, for both the NMT model and QE model, we train adapters instead of the full parameters (Pfeiffer et al., 2020). In QE-STATIC, we do not attach adapters to the QE model; instead we distill di-

rectly from the base model. All models were trained for 10 epochs, with early stopping based on COMET-KIWI score on the validation set. Evaluation was conducted using the metrics SacreBLEU, COMET-KIWI, COMET, and XCOMET (Post, 2018; Guerreiro et al., 2023). COMET is the reference-based counterpart to COMET-KIWI, and XCOMET is an explainable metric that provides error spans and categories as well as sentence-level scores. More experimental details can be found in Appendix B.

4.5 Results

We report results for QE-STATIC and QE-DYNAMIC along with the supervised baseline for low resource target languages (Bengali, Azerbaijani, Mongolian, Marathi, Kazakh) in Table 1. While the better-performing method varies among the two QE- variants, these methods exhibit consistent improvements over the baseline. For Mongolian, our method achieves improvements of 2.5 BLEU, 7.1 COMET-KIWI, 5.3 COMET, and 6.4 XCOMET relative to the supervised baseline.

We report results for QE-STATIC and QE-DYNAMIC along with the supervised baseline for high resource target languages (German and Chinese) in Table 2. For both languages, QE-DYNAMIC outperforms QE-Static and the supervised baseline in all four metrics. The amount of improvement over supervised fine-tuning is larger for lower resource languages.

Translation samples for Bengali, Marathi, German, Chinese are given in Table 3. Samples for Mongolian, Azerbaijani, and Kazakh are given in Appendix D.

Language Pair		EN-BN			
Data Size		4K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	5.76	65.46	65.72	39.64	
QE-STATIC	6.22	66.66	65.86	42.33	
QE-DYNAMIC	6.04	67.14	67.06	41.43	
Language Pair		EN-AZ			
Data Size		5K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	5.61	68.17	71.95	59.34	
QE-STATIC	6.51	73.63	75.17	63.03	
QE-DYNAMIC	6.78	73.65	75.10	62.96	
Language Pair		EN-MN			
Data Size		7K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	4.16	62.97	70.80	55.26	
QE-STATIC	6.62	70.09	76.10	61.71	
QE-DYNAMIC	6.64	69.83	75.85	61.44	
Language Pair		EN-MR			
Data Size		9K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	6.43	54.61	57.64	31.98	
QE-STATIC	6.22	57.56	58.37	33.16	
QE-DYNAMIC	6.45	58.44	58.45	33.26	
Language Pair		EN-KA			
Data Size		12K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	9.93	72.47	73.92	64.64	
QE-STATIC	9.93	73.48	74.48	65.19	
QE-DYNAMIC	9.65	72.97	74.27	64.73	

Table 1: Main results for low resource language pairs (**Bengali, Azerbaijani, Mongolian, Marathi, Kazakh**). All numbers represent the average of three experiments with different seeds. EBM methods were run with both labeled data filtering and adjacent unlabeled batch retrieval. The best result for each metric in each dataset is highlighted in bold.

5 Analysis

In this section, we first present more comprehensive experimental results for four languages (Bengali, Marathi, German, Chinese) in Tables 4, 5 and 6. In addition to the supervised baseline and EBM methods with both data filtering and retrieval, we report the results of reinforcement learning baselines and ablation experiments. Different types of ablations were examined for each algorithm. For the REINFORCE and PPO baselines, we tried training without any additional monolingual data, using the labeled data with different shuffling as the unlabeled data (-Mono). (+Mono) refers to the original setup of using the whole *pre-processed data pool* as the unlabeled data. For QE-DYNAMIC and QE-STATIC, we ran three types of ablation experiments, training without additional

Language Pair		EN-DE			
Data Size		50K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	24.24	77.23	79.01	92.07	
QE-STATIC	23.70	77.13	78.71	92.05	
QE-DYNAMIC	25.29	77.97	79.66	92.20	
Language Pair		EN-ZH			
Data Size		50K			
Metric	BLEU	KIWI	COMET	XCOMET	
Supervised	20.02	75.94	77.71	76.51	
QE-STATIC	20.01	76.29	77.73	76.14	
QE-DYNAMIC	20.10	76.35	77.88	76.54	

Table 2: Main results for high resource language pairs (**German, Chinese**). All numbers represent the average of three experiments with different seeds. EBM methods were run with both labeled data filtering and adjacent unlabeled batch retrieval. The best result for each metric in each dataset is highlighted in bold.

monolingual data (-Mono), training with only unlabeled batch retrieval (NN) or labeled batch filtering (FILTER), and training with neither NN nor FILTER.

5.1 Reinforcement Learning vs Energy-based Training

EBM methods outperform reinforcement learning methods in all metrics. We hypothesize that this superiority arises from the difference in the granularity of information the QE score holds in each algorithm. RL methods treat the score given by the quality estimation model as a simple scalar value, while EBM methods treat it as a loss that can back-propagate to update the translation model’s parameters. For REINFORCE or PPO, a reward is given per sequence, meaning that it assigns the same reward for all the tokens in the same sequence. Since it is stripped of gradients when used as part of the loss term, it holds no information other than the quality of the translation pair relative to other translation pairs. The score value acts as only a “scaling factor” for the gradients of the log probability term (similar to cross-entropy) in the loss. Meanwhile, in the case of energy-based methods, the score holds fine-grained information about each token’s contribution to the sequence-level quality estimation. The gradients with respect to the score hold information about how each parameter of the NMT model should be updated to produce a translation that is most compatible with the source sentence.

			EN-BN				EN-MR			
	Mono	FILTER/NN	BLEU	KIWI	COMET	XCOMET	BLEU	KIWI	COMET	XCOMET
Supervised	-	-	5.76	65.46	65.72	39.64	6.43	54.61	57.64	31.98
	-	FILTER	4.59	62.48	62.96	39.33	3.10	47.38	52.03	28.90
REINFORCE	+Mono	-	5.22	64.02	64.03	39.92	5.86	52.81	55.91	31.55
	-Mono	-	5.02	62.53	63.07	38.22	5.48	51.53	55.40	30.84
PPO	+Mono	-	5.53	64.43	65.09	40.28	4.77	49.15	53.98	29.93
	-Mono	-	4.03	62.13	62.89	38.86	4.77	49.45	54.39	29.88
QE-STATIC	+Mono	-	5.56	64.49	64.94	39.17	5.47	52.28	56.68	30.47
		NN	5.79	65.20	65.60	40.29	5.92	54.19	57.01	31.77
		FILTER	5.56	66.07	66.65	41.24	6.31	57.10	58.46	32.72
	-Mono	FILTER & NN	6.22	<u>66.66</u>	65.86	42.33	6.22	<u>57.56</u>	58.37	33.16
		-	5.62	64.94	66.13	39.89	5.84	53.33	56.62	31.59
QE-DYNAMIC	+Mono	-	6.04	63.73	66.21	38.91	6.00	54.14	57.25	31.45
		NN	6.00	66.17	<u>66.91</u>	40.78	6.59	54.48	57.75	31.92
		FILTER	5.72	66.43	<u>66.43</u>	<u>41.59</u>	6.44	57.46	57.69	<u>33.24</u>
	-Mono	FILTER & NN	6.04	67.14	67.06	41.43	6.45	58.44	<u>58.45</u>	33.26
		-	<u>6.05</u>	64.59	65.65	40.22	<u>6.56</u>	54.63	54.72	31.82

Table 4: Detailed results for **Bengali, Marathi**. All numbers represent the average of three experiments with different seeds. EBM (+Mono) methods were run with all possible combinations of data-centric techniques. RL baselines were run without any data-centric techniques. The best result for each metric is highlighted in bold, and the second best is underlined.

			EN-DE				EN-ZH			
	Mono	FILTER/NN	BLEU	KIWI	COMET	XCOMET	BLEU	KIWI	COMET	XCOMET
Supervised	-	-	24.24	77.23	79.01	92.07	20.02	75.94	77.71	76.51
	-	FILTER	23.50	77.11	78.73	92.01	19.42	75.17	77.00	75.81
REINFORCE	+Mono	-	24.45	77.21	79.07	91.92	20.00	75.91	77.61	76.58
	-Mono	-	24.57	77.36	79.17	92.16	20.08	75.82	77.54	76.08
PPO	+Mono	-	25.41	77.91	79.73	92.16	20.05	75.73	77.60	76.35
	-Mono	-	25.36	<u>78.16</u>	<u>79.85</u>	92.27	20.16	75.79	77.67	76.22
QE-STATIC	+Mono	-	24.31	77.30	79.10	92.11	19.90	75.74	77.51	76.04
		NN	22.85	76.71	78.61	86.87	19.79	75.76	77.46	76.10
		FILTER	25.16	77.64	79.32	92.25	<u>20.17</u>	76.38	77.88	<u>76.60</u>
	-Mono	FILTER & NN	23.70	77.13	78.71	92.05	20.01	76.29	77.73	76.14
		-	25.47	77.94	79.79	<u>92.52</u>	19.88	75.89	77.59	76.37
QE-DYNAMIC	+Mono	-	<u>25.51</u>	78.10	79.84	92.41	19.87	75.80	77.61	76.32
		NN	25.62	78.35	80.14	92.73	20.05	75.98	77.72	76.64
		FILTER	25.48	77.93	79.65	92.28	20.27	76.25	<u>77.75</u>	76.20
	-Mono	FILTER & NN	25.29	77.97	79.66	92.20	20.10	<u>76.35</u>	77.88	76.54
		-	25.40	78.08	79.80	92.34	19.95	75.71	77.54	76.54

Table 5: Detailed results for **German, Chinese**. All numbers represent the average of three experiments with different seeds. EBM (+Mono) methods were run with all possible combinations of data-centric techniques. RL baselines were run without any data-centric techniques. The best result for each metric is highlighted in bold, and the second best is underlined.

	Mono	FILTER/NN	BLEU	KIWI	COMET	XCOMET
Supervised	-	-	14.11	68.31	70.02	60.05
	-	FILTER	12.65	65.54	67.68	59.01
REINFORCE	+Mono	-	13.88	67.49	69.16	59.99
	-Mono	-	13.79	66.81	68.79	59.33
PPO	+Mono	-	13.94	66.81	69.10	59.68
	-Mono	-	13.58	66.38	68.70	59.31
QE-Static	+Mono	-	13.81	67.45	69.56	59.45
		NN	13.58	67.96	69.67	58.76
		FILTER	14.30	69.30	70.57	60.70
	-Mono	FILTER & NN	14.04	69.41	70.17	60.92
		-	14.20	68.02	70.03	60.09
QE-Dynamic	+Mono	-	14.36	67.94	70.23	59.77
		NN	14.56	68.74	<u>70.63</u>	60.52
		FILTER	14.48	<u>69.52</u>	70.38	60.83
	-Mono	FILTER & NN	14.47	69.97	70.76	<u>60.86</u>
		-	<u>14.49</u>	68.25	69.43	60.23

Table 6: Detailed results averaged for all four languages (**German, Chinese, Bengali, Marathi**). The best result for each metric is highlighted in bold, and the second best is underlined.

MR translation. For low resource target languages, applying each technique separately as well as combining them enhances translation quality for both QE-DYNAMIC and QE-STATIC.

Table 5 contains comparisons for high resource languages. For both German and Chinese, applying unlabeled batch retrieval consistently improves performance for QE-DYNAMIC, but not for QE-STATIC. We hypothesize that unlabeled batch retrieval is more beneficial for QE-DYNAMIC because the energy model is fine-tuned with the labeled batch, and is likely to make the most accurate predictions on those batches. Using an unlabeled batch that is similar to the labeled batch during NMT model fine-tuning best exploits the tuned energy model.

For high resource languages, filtering mostly improves the performance of the NMT model, with the exception of QE-DYNAMIC for German.

6 Conclusion

We propose a new method of improving neural machine translation by using quality estimators as trainable energy loss networks. Our method performs better than supervised fine-tuning and reinforcement learning baselines in both high and low-resource language directions, especially showing greater improvements for low-resource languages. Based on our experiments, the best performance is expected when applying QE-DYNAMIC with labeled data filtering and retrieval of nearest unlabeled batches.

beled batches.

Limitations

Our proposed training scheme has several limitations. First, although it reduces latency during inference compared to other methods such as energy-based reranking, the joint update requires a large amount of computation and memory during training, since we need to calculate and store gradients for both models. Second, it is difficult to use the quality estimation model as an energy model as is when there is a vocabulary mismatch between the energy model and the NMT model. This problem may be solved by adapting the quality estimator to operate in latent space, which is a potential future research direction.

Acknowledgement

We sincerely thank Jonghyun Song, Seungwoo Choi, and anonymous reviewers for their helpful discussion and feedback. This work was supported in part by the National Research Foundation of Korea (NRF) grant (RS-2023-00280883, RS-2023-00222663) and by the National Super computing Center with super computing resources including technical support (KSC-2023-CRE-0176), with the aid of computing resources from Artificial Intelligence Industry Center Agency and Google cloud platform research credits, and partially supported by New Faculty Startup Fund from Seoul National University.

References

- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. [Estimating or propagating gradients through stochastic neurons for conditional computation](#). *Preprint*, arXiv:1308.3432.
- Sumanta Bhattacharyya, Amirmohammad Rooshenas, Subhajt Naskar, Simeng Sun, Mohit Iyyer, and Andrew McCallum. 2021. [Energy-based reranking: Improving neural machine translation using energy-based models](#). 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 4528–4537, Online. Association for Computational Linguistics.
- Mauro Cettolo, Marcello Federico, Luisa Bentivogli, Jan Niehues, Sebastian Stüker, Katsuhito Sudoh, Koichiro Yoshino, and Christian Federmann. 2017. [Overview of the IWSLT 2017 evaluation campaign](#). In *Proceedings of the 14th International Conference on Spoken Language Translation*, pages 2–14, Tokyo, Japan. International Workshop on Spoken Language Translation.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. [InfoXML: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Yuntian Deng, Anton Bakhtin, Myle Ott, Arthur Szlam, and Marc’Aurelio Ranzato. 2020. [Residual energy-based models for text generation](#). In *International Conference on Learning Representations*.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. [Raft: Reward ranked finetuning for generative foundation model alignment](#). *Preprint*, arXiv:2304.06767.
- Markus Freitag, Nitika Mathur, Chi-kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Tom Kocmi, Frederic Blain, Daniel Deutsch, Craig Stewart, Chrysoula Zerva, Sheila Castilho, Alon Lavie, and George Foster. 2023. [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.
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2023. [xcomet: Transparent machine translation evaluation through fine-grained error detection](#). *Preprint*, arXiv:2310.10482.
- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, Wolfgang Macherey, Arnaud Doucet, Orhan Firat, and Nando de Freitas. 2023. [Reinforced self-training \(rest\) for language modeling](#). *Preprint*, arXiv:2308.08998.
- Tianxing He, Bryan McCann, Caiming Xiong, and Ehsan Hosseini-Asl. 2021. [Joint energy-based model training for better calibrated natural language understanding models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1754–1761, Online. Association for Computational Linguistics.
- 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](#). *Preprint*, arXiv:2401.12873.
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A method for stochastic optimization](#). *Preprint*, arXiv:1412.6980.
- Yann Lecun, Sumit Chopra, Raia Hadsell, Marc Aurelio Ranzato, and Fu Jie Huang. 2006. *A tutorial on energy-based learning*. MIT Press.
- Jay Yoon Lee, Dhruv Patel, Purujit Goyal, Wenlong Zhao, Zhiyang Xu, and Andrew McCallum. 2022. [Structured energy network as a loss](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 20862–20875. Curran Associates, Inc.
- Zhuang Ma and Michael Collins. 2018. [Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3698–3707, Brussels, Belgium. Association for Computational Linguistics.
- Richard Yuanzhe Pang, Vishakh Padmakumar, Thibault Sellam, Ankur Parikh, and He He. 2023. [Reward gaming in conditional text generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4746–4763, Toronto, Canada. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. [Adapterhub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020): Systems Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on*

- Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, and Yejin Choi. 2023. [Is reinforcement learning \(not\) for natural language processing: Benchmarks, baselines, and building blocks for natural language policy optimization](#). *Preprint*, arXiv:2210.01241.
- Miguel Moura Ramos, Patrick Fernandes, António Farinhas, and André F. T. Martins. 2023. [Aligning neural machine translation models: Human feedback in training and inference](#). *Preprint*, arXiv:2311.09132.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C. Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte M. Alves, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. [Cometkiwi: Ist-unbabel 2022 submission for the quality estimation shared task](#). *Preprint*, arXiv:2209.06243.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *Preprint*, arXiv:1707.06347.
- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. [Policy gradient methods for reinforcement learning with function approximation](#). In *Advances in Neural Information Processing Systems*, volume 12. MIT Press.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#). *Preprint*, arXiv:2008.00401.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2021. [Multilingual translation from denoising pre-training](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3450–3466, Online. Association for Computational Linguistics.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semaire Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian 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, Madian 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](#). *Preprint*, arXiv:2307.09288.
- Lifu Tu, Richard Yuanzhe Pang, Sam Wiseman, and Kevin Gimpel. 2020. [ENGINE: Energy-based inference networks for non-autoregressive machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2826, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, and Shengyi Huang. 2020. [Trl: Transformer reinforcement learning](#). <https://github.com/huggingface/trl>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Chrysoula Zerva, Frédéric Blain, Ricardo Rei, Piyawat Lertvittayakumjorn, José G. C. de Souza, Diptesh Kanojia, Duarte Alves, Constantin Orăsan, Marina Fomicheva, André F. T. Martins, and Lucia Specia. 2022. [Findings of the WMT 2022 shared task on quality estimation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 69–99, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

A Hyperparameters

Symbol	Function	Value
α	weight for unlabeled batch	$\min(0, steps/1000) * 0.001$
β	weight for labeled batch	$\max(0, 1 - 10 * \alpha)$
K	number of hypotheses for energy loss in NMT model	5
N	number of hypotheses for contrastive learning in energy model	5

Table 7: Hyperparameters related to training.

B Experimental Details

- When retrieving the nearest unlabeled batch for each labeled batch, the embeddings used to calculate cosine similarity are taken from the “all-mpnet-base-v2” checkpoint in the SBERT library (Reimers and Gurevych, 2019).
- At the start of training, MBART weights are initialized to the “facebook/mbart-large-50” checkpoint from Huggingface (Wolf et al., 2020).
- At the start of training, the QE model’s weights are initialized to the “Unbabel/wmt22-cometkiwi-da” checkpoint from Huggingface (Wolf et al., 2020).
- Adapter implementations are taken from AdapterHub (Pfeiffer et al., 2020).

C Computation Time

Operation	Data	Time(s)
filter by quality (batch size = 1000)	train data after length filtering (2769)	24.5
compute corpus embeddings (batch size = 32)	train data after quality filtering (2283 = 0.8245 * 2769)	0.77
compute nearest neighbors	filtered train data (2283 = 0.8245 * 2769)	0.86

Table 8: Computation time of precomputed operations for data filtering and for nearest neighbor retrieval. The time was computed using 2769 data samples from the English-Bengali dataset after applying length-based filtering. When we filter out poorly aligned samples from this pool, 2283 samples are left. Before starting the model updates, we compute embeddings for each source sentence and also compute the nearest neighbor for each sentence based on the embeddings. For approximately 2K samples, these operations can each be completed in less than a second.

Algorithm	Time(s) per batch
Supervised	0.2
REINFORCE	8.4
PPO	9.5
QE-STATIC	8.9
QE-DYNAMIC	17.2

Table 9: Latency of a single training step for each algorithm. The time was computed using the English-Bengali dataset, with a batch size of 4 and gradient accumulation step of 4. QE-STATIC runs at a similar speed to REINFORCE, PPO and QE-dynamic at half the speed.

D Translation Samples

