

A Stable and Effective Learning Strategy for Trainable Greedy Decoding

Yun Chen[†], Victor O.K. Li[†], Kyunghyun Cho^{‡,*}, Samuel R. Bowman[‡]

[†]The University of Hong Kong, [‡]New York University, *CIFAR Global Scholar

yun.chencreek@gmail.com, vli@eee.hku.hk,

kyunghyun.cho@nyu.edu, bowman@nyu.edu

Abstract

Beam search is a widely used approximate search strategy for neural network decoders, and it generally outperforms simple greedy decoding on tasks like machine translation. However, this improvement comes at substantial computational cost. In this paper, we propose a flexible new method that allows us to reap nearly the full benefits of beam search with nearly no additional computational cost. The method revolves around a small neural network *actor* that is trained to observe and manipulate the hidden state of a previously-trained decoder. To train this actor network, we introduce the use of a pseudo-parallel corpus built using the output of beam search on a base model, ranked by a target quality metric like BLEU. Our method is inspired by earlier work on this problem, but requires no reinforcement learning, and can be trained reliably on a range of models. Experiments on three parallel corpora and three architectures show that the method yields substantial improvements in translation quality and speed over each base system.

1 Introduction

Neural network sequence decoders yield state-of-the-art results for many text generation tasks, including machine translation (Bahdanau et al., 2015; Luong et al., 2015; Gehring et al., 2017; Vaswani et al., 2017; Dehghani et al., 2018), text summarization (Rush et al., 2015; Ranzato et al., 2015; See et al., 2017; Paulus et al., 2017) and image captioning (Vinyals et al., 2015; Xu et al., 2015). These decoders generate tokens from left to right, at each step giving a distribution over possible next tokens, conditioned on both the input and all the tokens generated so far. However, since the space of all possible output sequences is infinite and grows exponentially with sequence length, heuristic search methods such as greedy decod-

ing or beam search (Graves, 2012; Boulanger-Lewandowski et al., 2013) must be used at decoding time to select high-probability output sequences. Unlike greedy decoding, which selects the token of the highest probability at each step, beam search expands all possible next tokens at each step, and maintains the k most likely prefixes, where k is the beam size. Greedy decoding is very fast—requiring only a single run of the underlying decoder—while beam search requires an equivalent of k such runs, as well as substantial additional overhead for data management. However, beam search often leads to substantial improvement over greedy decoding. For example, Ranzato et al. (2015) report that beam search (with $k = 10$) gives a 2.2 BLEU improvement in translation and a 3.5 ROUGE-2 improvement in summarization over greedy decoding.

Various approaches have been explored recently to improve beam search by improving the method by which candidate sequences are scored (Li et al., 2016; Shu and Nakayama, 2017), the termination criterion (Huang et al., 2017), or the search function itself (Li et al., 2017). In contrast, Gu et al. (2017) have tried to directly improve greedy decoding to decode for an arbitrary decoding objective. They add a small *actor* network to the decoder and train it with a version of policy gradient to optimize sequence objectives like BLEU. However, they report that they are seriously limited by the instability of this approach to training.

In this paper, we propose a procedure to modify a trained decoder to allow it to generate text greedily with the level of quality (according to metrics like BLEU) that would otherwise require the relatively expensive use of beam search. To do so, we follow Cho (2016) and Gu et al. (2017) in our use of an *actor* network which manipulates the decoder’s hidden state, but introduce a stable and effective procedure to train this actor. In our

training procedure, the actor is trained with ordinary backpropagation on a model-specific artificial parallel corpus. This corpus is generated by running the un-augmented model on the training set with large-beam beam search, and selecting outputs from the resulting k -best list which score highly on our target metric.

Our method can be trained quickly and reliably, is effective, and can be straightforwardly employed with a variety of decoders. We demonstrate this for neural machine translation on three state-of-the-art architectures: RNN-based (Luong et al., 2015), ConvS2S (Gehring et al., 2017) and Transformer (Vaswani et al., 2017), and three corpora: IWSLT16 German-English,¹ WMT15 Finnish-English² and WMT14 German-English.³

2 Background

2.1 Neural Machine Translation

In sequence-to-sequence learning, we are given a set of source–target sentence pairs and tasked with learning to generate each target sentence (as a sequence of words or word-parts) from its source sentence. We first use an encoding model such as a recurrent neural network to transform a source sequence into an encoded representation, then generates the target sequence using a neural decoder.

Given a source sentence $\mathbf{x} = \{x_1, \dots, x_{T_s}\}$, a neural machine translation system models the distribution over possible output sentences $\mathbf{y} = \{y_1, \dots, y_T\}$ as:

$$P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \prod_{t=1}^T P(y_t|y_{<t}, \mathbf{x}; \boldsymbol{\theta}), \quad (1)$$

where $\boldsymbol{\theta}$ is the set of model parameters.

Given a parallel corpus $D_{x,y}$ of source–target sentence pairs, the neural machine translation model can be trained by maximizing the log-likelihood:

$$\hat{\boldsymbol{\theta}} = \operatorname{argmax}_{\boldsymbol{\theta}} \left\{ \sum_{\langle \mathbf{x}, \mathbf{y} \rangle \in D_{x,y}} \log P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) \right\}. \quad (2)$$

2.2 Decoding

Given estimated model parameters $\hat{\boldsymbol{\theta}}$, the decision rule for finding the translation with the highest

probability for a source sentence \mathbf{x} is given by

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} \left\{ P(\mathbf{y}|\mathbf{x}; \hat{\boldsymbol{\theta}}) \right\}. \quad (3)$$

However, since such exact inference requires the intractable enumeration of large and potentially infinite set of candidate sequences, we resort to approximate decoding algorithms such as greedy decoding, beam search, noisy parallel decoding (NPAD; Cho, 2016), or trainable greedy decoding (Gu et al., 2017).

Greedy Decoding In this algorithm, we generate a single sequence from left to right, by choosing the token that is most likely at each step. The output $\hat{\mathbf{y}} = \{\hat{y}_1, \dots, \hat{y}_T\}$ can be represented as

$$\hat{y}_t = \operatorname{argmax}_y \left\{ P(y|\hat{y}_{<t}, \mathbf{x}; \hat{\boldsymbol{\theta}}) \right\}. \quad (4)$$

Despite its low computational complexity of $O(|V| \times T)$, the translations selected by this method may be far from optimal under the overall distribution given by the model.

Beam Search Beam search decodes from left to right, and maintains $k > 1$ hypotheses at each step. At each step t , beam search considers all possible next tokens conditioned on the current hypotheses, and picks the k with the overall highest scores $\prod_{t'=1}^t P(y_{t'}|y_{<t'}, \mathbf{x}; \hat{\boldsymbol{\theta}})$. When all the hypotheses are complete (they end in an end-of-the-sentence symbol or reach a predetermined length limit), it returns the hypothesis with the highest likelihood. Tuning to find a roughly optimal beam size k can yield improvements in performance with sizes as high as 30 (Koehn and Knowles, 2017; Britz et al., 2017). However, the complexity of beam search grows linearly in beam size, with high constant terms, making it undesirable in some applications where latency is important, such as in on-device real-time translation.

NPAD Noisy parallel approximate decoding (NPAD; Cho, 2016) is a parallel decoding algorithm that can be used to improve greedy decoding or beam search. The main idea is that a better translation with a higher probability may be found by injecting unstructured random noise into the hidden state of the decoder network. Positive results with NPAD suggest that small manipulations to the decoder hidden state can correspond to substantial but still reasonable changes to the output sequence.

¹<https://wit3.fbk.eu/>

²<http://www.statmt.org/wmt15/translation-task.html>

³<http://www.statmt.org/wmt14/translation-task>

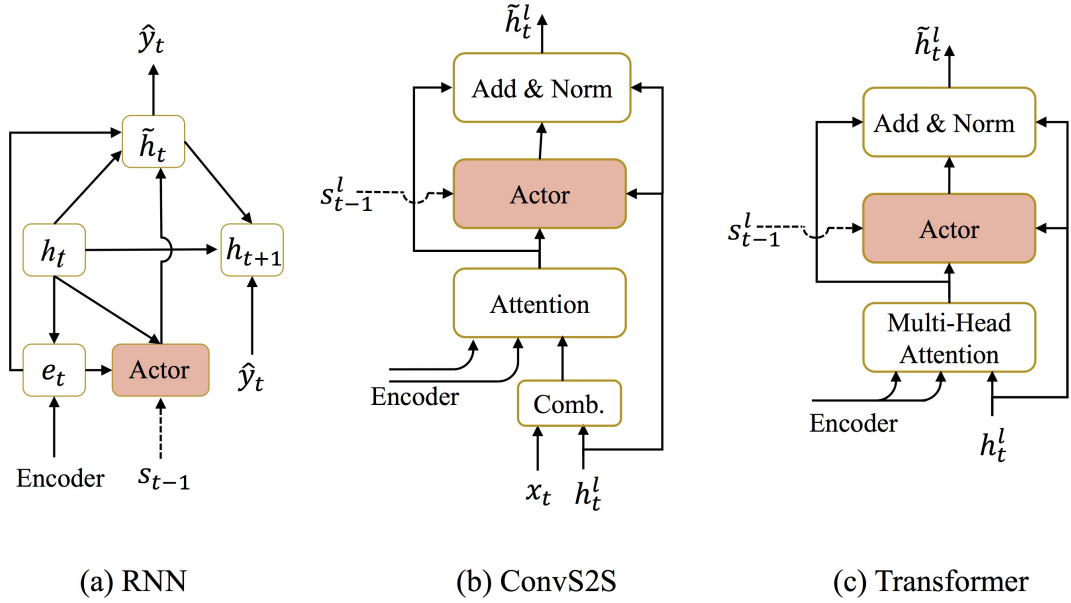


Figure 1: A single step of a generic actor interacting with a decoder of each of three types. The dashed arrows denote an optional recurrent connection in the actor network.

Trainable Greedy Decoding Approximate decoding algorithms generally approximate the maximum-a-posteriori inference described in Equation 3. This is not necessarily the optimal basis on which to generate text, since (i) the conditional log-probability assigned by a trained NMT model does not necessarily correspond well to translation quality (Tu et al., 2017), and (ii) different application scenarios may demand different decoding objectives (Gu et al., 2017). To solve this, Gu et al. (2017) extend NPAD by replacing the unstructured noise with a small feedforward actor neural network. This network is trained using a variant of policy gradient reinforcement learning to optimize for a target quality metric like BLEU under greedy decoding, and is then used to guide greedy decoding at test time by modifying the decoder’s hidden states. Despite showing gains over the equivalent actorless model, their attempt to directly optimize the quality metric makes training unstable, and makes the model nearly impossible to optimize fully. This paper offers a stable and effective alternative approach to training such an actor, and further develops the architecture of the actor network.

3 Methods

We propose a method for training a small actor neural network, following the trainable greedy decoding approach of Gu et al. (2017). This actor

takes as input the current decoder state \mathbf{h}_t , an attentional context vector \mathbf{e}_t for the source sentence, and optionally the previous hidden state \mathbf{s}_{t-1} of the actor, and produces a vector-valued action \mathbf{a}_t which is used to update the decoder hidden state. The actor function can take on a variety of forms, and we explore four: a feedforward network with one hidden layer (*ff*), feedforward network with two hidden layers (*ff2*), a GRU recurrent network (*rnn*; Cho et al., 2014), and gated feedforward network (*gate*).

The feedforward *ff* actor function is computed as

$$\begin{aligned} \mathbf{z}_t &= \sigma([\mathbf{h}_t, \mathbf{e}_t]W^i + \mathbf{b}^i), \\ \mathbf{a}_t &= \tanh(\mathbf{z}_tW^o + \mathbf{b}^o), \end{aligned} \quad (5)$$

the *ff2* actor is computed as

$$\begin{aligned} \mathbf{z}_t^1 &= \sigma([\mathbf{h}_t, \mathbf{e}_t]W^i + \mathbf{b}^i), \\ \mathbf{z}_t^2 &= \sigma(\mathbf{z}_t^1W^z + \mathbf{b}^z), \\ \mathbf{a}_t &= \tanh(\mathbf{z}_t^2W^o + \mathbf{b}^o), \end{aligned} \quad (6)$$

the *rnn* actor is computed as

$$\begin{aligned} \mathbf{z}_t &= \sigma([\mathbf{h}_t, \mathbf{e}_t]U^z + \mathbf{s}_{t-1}W^z), \\ \mathbf{r}_t &= \sigma([\mathbf{h}_t, \mathbf{e}_t]U^r + \mathbf{s}_{t-1}W^r), \\ \tilde{\mathbf{s}}_t &= \tanh([\mathbf{h}_t, \mathbf{e}_t]U^h + (\mathbf{s}_{t-1} \circ \mathbf{r}_t)W^h), \\ \mathbf{s}_t &= (1 - \mathbf{z}_t) \circ \tilde{\mathbf{s}}_t + \mathbf{z}_t \circ \mathbf{s}_{t-1}, \\ \mathbf{a}_t &= \mathbf{s}_tU, \end{aligned} \quad (7)$$

and the *gate* actor is computed as

$$\begin{aligned} \mathbf{z}_t &= \sigma([\mathbf{h}_t, \mathbf{e}_t]U^z), \\ \mathbf{a}_t &= \mathbf{z}_t \circ \tanh([\mathbf{h}_t, \mathbf{e}_t]U). \end{aligned} \quad (8)$$

Once the action \mathbf{a}_t has been computed, the hidden state \mathbf{h}_t is simply replaced with the updated state $\tilde{\mathbf{h}}_t$:

$$\tilde{\mathbf{h}}_t = f(\mathbf{h}_t, \mathbf{e}_t) + \mathbf{a}_t. \quad (9)$$

Figure 1 shows a single step of the actor interacting with the underlying neural decoder of each of the three NMT architectures we use: the RNN-based model of Luong et al. (2015), ConvS2S (Gehring et al., 2017), and Transformer (Vaswani et al., 2017). We add the actor at the decoder layer immediately after the computation of the attentional context vector. For the RNN-based NMT, we add the actor network only to the last decoder layer, the only place attention is used. Here, it takes as input the hidden state of the last decoder layer \mathbf{h}_t^L and the source context vector \mathbf{e}_t , and outputs the action \mathbf{a}_t , which is added back to the attention vector $\tilde{\mathbf{h}}_t^L$. For ConvS2S and Transformer, we add an actor network to each decoder layer. This actor is added to the sublayer which performs multi-head or multi-step attention over the output of the encoder stack. It takes as input the decoder state \mathbf{h}_t^l and the source context vector \mathbf{e}_t^l , and outputs an action \mathbf{a}_t^l which is added back to get $\tilde{\mathbf{h}}_t^l$.

Training To overcome the severe instability reported by Gu et al. (2017), we introduce the use of a *pseudo-parallel corpus* generated from the underlying NMT model (Gao and He, 2013; Auli and Gao, 2014; Kim and Rush, 2016; Chen et al., 2017; Freitag et al., 2017; Zhang et al., 2017) for actor training. This corpus includes pairs that both (i) have a high model likelihood, so that we can coerce the model to generate them without much additional training or many new parameters and, (ii) represent high-quality translations, measured according to a target metric like BLEU. We do this by generating sentences from the original unaugmented model with large-beam beam search and selecting the best sentence from the resulting k -best list according to the decoding objective.

More specifically, let $\langle \mathbf{x}, \mathbf{y} \rangle$ be a sentence pair in the training data and $\mathcal{Z} = \{\mathbf{z}^1, \dots, \mathbf{z}^k\}$ be the k -best list from beam search on the pretrained NMT model, where k is the beam size. We define the objective score of the translation \mathbf{z} w.r.t. the gold-standard translation \mathbf{y} according to a target metric such as BLEU (Papineni et al., 2002), NIST

(Doddington, 2002), negative TER (Snover et al., 2006), or METEOR (Lavie and Denkowski, 2009) as $O(\mathbf{z}, \mathbf{y})$. Then we choose the sentence $\tilde{\mathbf{z}}$ that has the highest score to become our new target sentence:

$$\tilde{\mathbf{z}} = \operatorname{argmax}_{\mathbf{z}=\mathbf{z}^1, \dots, \mathbf{z}^k} O(\mathbf{z}, \mathbf{y}). \quad (10)$$

Once we obtain the pseudo-corpus $D_{x,z} = \{\langle \mathbf{x}_i, \tilde{\mathbf{z}}_i \rangle_{i=1}^n\}$, we keep the underlying model fixed and train the actor by maximizing the log-likelihood of the actor parameters with these pairs:

$$\hat{\theta}_a = \operatorname{argmax}_{\theta_a} \left\{ \sum_{\langle \mathbf{x}, \mathbf{z} \rangle \in D_{x,z}} \log P(\mathbf{z} | \mathbf{x}; \hat{\theta}, \theta_a) \right\} \quad (11)$$

In this way, the actor network is trained to manipulate the neural decoder’s hidden state at decoding time to induce it to produce better-scoring outputs under greedy or small-beam decoding.

4 Experiments

4.1 Setting

We evaluate our approach on IWSLT16 German-English, WMT15 Finnish-English, and WMT14 De-En translation in both directions with three strong translation model architectures.

For IWSLT16, we use *tst2013* and *tst2014* for validation and testing, respectively. For WMT15, we use *newstest2013* and *newstest2015* for validation and testing, respectively. For WMT14, we use *newstest2013* and *newstest2014* for validation and testing, respectively. All the data are tokenized and segmented into subword symbols using byte-pair encoding (BPE; Sennrich et al., 2016) to restrict the size of the vocabulary. Our primary evaluations use tokenized and cased BLEU. For METEOR and TER evaluations, we use *multeval*⁴ with tokenized and case-insensitive scoring. All the underlying models are trained from scratch, except for ConvS2S WMT14 English-German translation, for which we use the trained model (as well as training data) provided by Gehring et al. (2017).⁵

⁴<https://github.com/jhclark/multeval>

⁵<https://s3.amazonaws.com/fairseq-py/models/wmt14.v2.en-de.fconv-py.tar.bz2>

	BLEU↑						tok/s↑					
	greedy	beam4	tg	greedy	beam4	tg	greedy	beam4	tg	greedy	beam4	tg
IWSLT16	De → En						En → De					
RNN	23.57	<u>24.90</u>	23.59	62.8	45.0	60.4	20.05	<u>21.11</u>	19.88	48.1	32.5	45.7
ConvS2S	27.44	<u>28.80</u>	28.74	191.1	87.2	167.5	22.88	<u>24.02</u>	24.42	136.5	64.0	124.0
Transformer	27.15	<u>28.74</u>	28.36	63.9	31.0	59.8	23.87	25.03	<u>25.46</u>	57.9	26.5	51.2
WMT15	Fi → En						En → Fi					
RNN	12.45	<u>13.22</u>	13.02	51.5	33.1	43.4	9.77	<u>10.81</u>	10.57	44.0	31.2	43.8
ConvS2S	15.43	<u>16.86</u>	17.17	24.8	11.4	16.2	12.65	<u>13.97</u>	14.33	25.0	11.7	16.9
Transformer	13.76	<u>14.61</u>	14.49	31.4	13.4	29.8	12.38	<u>13.55</u>	12.95	29.8	12.8	27.9
WMT14	De → En						En → De					
RNN	23.08	<u>24.62</u>	24.54	38.4	26.6	36.4	18.87	<u>20.59</u>	19.89	33.2	22.4	32.5
ConvS2S	27.52	<u>28.79</u>	28.56	22.5	9.9	14.6	24.86	<u>25.71</u>	26.04	19.9	9.1	13.6
Transformer	26.44	<u>27.31</u>	26.96	32.9	14.3	30.9	22.01	<u>22.74</u>	22.31	28.5	12.2	26.1

Table 1: Generation quality (BLEU) and speed (tokens/sec). Speed is measured for sentence-by-sentence generation without mini-batching on the test set on CPU. We show the result by the underlying model with greedy decoding (greedy), beam search with $k = 4$ (beam4) and our trainable greedy decoder (tg).

	BLEU↑			
	tg	tg+beam4	tg	tg+beam4
IWSLT16	De → En		En → De	
RNN	23.59	25.03	19.88	20.72
ConvS2S	28.74	29.50	24.42	24.74
Transformer	28.36	28.95	25.46	25.89
WMT15	Fi → En		En → Fi	
RNN	13.02	13.49	10.57	11.04
ConvS2S	17.17	17.51	14.33	14.87
Transformer	14.49	14.79	12.95	13.45
WMT14	De → En		En → De	
RNN	24.54	24.86	19.89	20.56
ConvS2S	28.56	28.46	26.04	26.08
Transformer	26.96	27.21	22.31	21.92

Table 2: Generation quality (BLEU↑) using the proposed trainable greedy decoder without and with beam search ($k = 4$). Results without beam search (tg) are also appeared in Table 1.

RNN We use OpenNMT-py (Klein et al., 2017)⁶ to implement our model. It is composed of an encoder with two-layer bidirectional RNN, and a decoder with another two-layer RNN. We refer to OpenNMT’s default setting ($rnn_size = 500$, $word_vec_size = 500$) and the setting in Artetxe et al. (2018) ($rnn_size = 600$, $word_vec_size = 300$), and choose similar hyper-parameters: $rnn_size = 500$, $word_vec_size = 300$ for IWSLT16 and $rnn_size = 600$, $word_vec_size = 500$ for WMT. We use the input-feeding decoder and global attention with the general alignment function (Luong et al., 2015).

⁶<https://github.com/OpenNMT/OpenNMT-py>

ConvS2S We implement our model based on fairseq-py.⁷ We follow the settings in *fconv_iwslt_de_en* and *fconv_wmt_en_de* for IWSLT16 and WMT, respectively.

Transformer We implement our model based on the code from Gu et al. (2018).⁸ We follow their hyperparameter settings for all experiments.

In the results below, we focus on the *gate* actor and pseudo-parallel corpora constructed by choosing the sentence with the best BLEU score from the k -best list produced by beam search with $k = 35$. Experiments motivating these choices are shown later in this section.

4.2 Results and Analysis

The results (Table 1) show that the use of the actor makes it practical to replace beam search with greedy decoding in most cases: We lose little or no performance, and doing so yields an increase in decoding efficiency, even accounting for the small overhead added by the actor. Among the three architectures, ConvS2S—the one with the most and largest layers—performs best. We conjecture that this gives the decoder more flexibility with which to guide decoding. In cases where model throughput is less important, our method can also be combined with beam search at test time to yield results somewhat better than either could achieve alone. Table 2 shows the result when combining our method with beam search.

⁷<https://github.com/facebookresearch/fairseq-py>

⁸<https://github.com/salesforce/nonauto-nmt>

src	Am Vormittag wollte auch die Arbeitsgruppe Migration und Integration ihre Beratungen fortsetzen .
ref	During the morning , the Migration and Integration working group also sought to continue its discussions .
greedy	The morning also wanted to continue its discussions on migration and integration .
beam4	In the morning , the working group on migration and integration also <u>wanted to continue</u> its discussions .
beam35*	In the morning , the migration and integration working group also <u>wanted to continue</u> its discussions .
tg	The morning , the Migration and Integration Working Group wanted to <u>continue</u> its discussions .
tg+beam4	In the morning , the Migration and Integration Working Group wanted to <u>continue</u> its discussions .
src	Die meisten Mails werden unterwegs mehrfach von Software-Robotern gelesen .
ref	The majority of e-mails are read several times by software robots en route to the recipient .
greedy	Most mails are <u>read</u> by software robots on the go .
beam4	Most mails are <u>read</u> by software robots on the go .
beam35*	Most e-mails are <u>read several times</u> by software robots on the road .
tg	Most mails are <u>read several times</u> by software robots on the road .
tg+beam4	Most mails are <u>read several times</u> by software robots on the road .
src	Ich suche schon seit einiger Zeit eine neue Wohnung fr meinen Mann und mich .
ref	I have been looking for a new home for my husband and myself for some time now .
greedy	I have been looking for a new apartment for some time for my husband and myself .
beam4	I have been looking for a new apartment for some time for my husband and myself .
beam35*	I have been looking for a new apartment <u>for my husband and myself for some time now</u> .
tg	I have been looking for a new apartment <u>for my husband and myself for some time now</u> .
tg+beam4	I have been looking for a new apartment <u>for my husband and myself for some time now</u> .

Table 3: Translation examples from the WMT14 De-En test set with Transformer. We show translations generated by the underlying transformer using greedy decoding, beam search with $k = 4$, and beam search with $k = 35$ and the oracle BLEU scorer (\star). We also show the translations using our trainable greedy decoder both without and with beam search. Phrases of interest are underlined.

	IWSLT16 De-En				WMT14 De-En			
	ref	greedy	k35*	tg	ref	greedy	k35*	tg
Base Model	20.4	65.3	61.5	64.2	23.5	65.2	63.8	65.1
+Trainable Greedy Decoder	19.1	70.4	65.3	75.1	18.9	76.0	72.6	82.8

Table 4: Word-level likelihood (%) averaged by sentence for the IWSLT16 and WMT14 De-En test sets with Transformer. Each row represents the model used to evaluate word-level likelihood, and each column represents a different source of translations, including the reference (ref), greedy decoding on the base model (greedy), beam search with $k = 35$ on the base model and the BLEU scorer (k35*), and trainable greedy decoder (tg).

Examples Table 3 shows a few selected translations from the WMT14 German-English test set. In manual inspection of these examples and others, we find that the actor encourages models to recover missing tokens, optimize word order, and correct prepositions.

Likelihood We also compare word-level likelihood for different decoding results assigned by the base model and the actor-augmented model. For a sentence pair $\langle \mathbf{x}, \mathbf{y} \rangle$, word-level likelihood is defined as

$$P_w = \frac{1}{T} \sum_{t=1}^T P(y_t | y_{<t}, \mathbf{x}; \theta). \quad (12)$$

Table 4 shows the word-level likelihood averaged over the test set for IWSLT16 and WMT14 German to English translation with Transformer.

Our trainable greedy decoder learns a much more peaked distribution and assigns a much higher probability mass to its greedy decoding result than the base model. When evaluated under the base model, the translations from trainable greedy decoding have smaller likelihood than the translations from greedy decoding using the base model for both datasets. This indicates that the trainable greedy decoder is able to find a sequence that is not highly scored by the underlying model, but that corresponds to a high value of the target metric.

Magnitude of Action Vector We also record the L_2 norm of the action, decoder hidden state, and attentional source context vectors on the validation set. Figure 2 shows these values over the course of training on the IWSLT16 De-En validation set with Transformer. The norm of the action starts

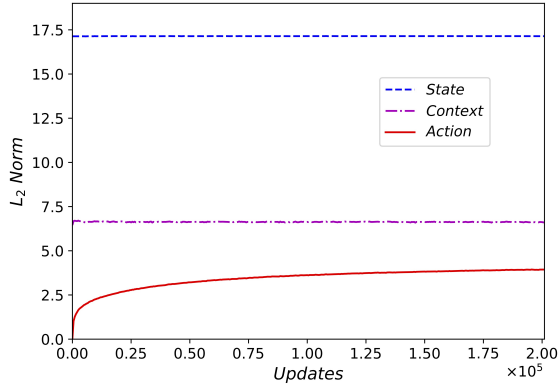


Figure 2: The norms of the three activation vectors on the IWSLT16 De-En validation set with Transformer. *Action*, *Context* and *State* represent the norm of the action, attentional source context vector and decoder hidden state, respectively.

small, increases rapidly early in training, and converges to a value well below that of the decoder hidden state. This suggests that the action adjusts the decoders hidden state only slightly, rather than overwriting it.

4.3 Effects of Model Settings

Actor Architecture Figure 3 shows the trainable greedy decoding result on IWSLT16 De-En validation set with different actor architectures. We observe that our approach is stable across different actor architectures and is relatively insensitive to the hyperparameters of the actor. For the same type of actor, the performance increases gradually with the hidden layer size. The use of a recurrent connection within the actor does not meaningfully improve performance, possibly since all actors can use the recurrent connections of the underlying decoder. Since the gate actor contains no additional hyperparameters and was observed to learn quickly and reliably, we use it in all other experiments.

Here, we also explore a simple alternative to the use of the actor: creating a pseudo-parallel corpus with each model, and then training each model, unmodified and entirely, directly on this new corpus. This experiment (*cont.* in Figure 3) yields results that are comparable to, but not better than, the results seen with the actors. However, this comes with substantially greater computational complexity at training time, and, if the same trained model is to be optimized for multiple target metrics, greater storage costs as well.

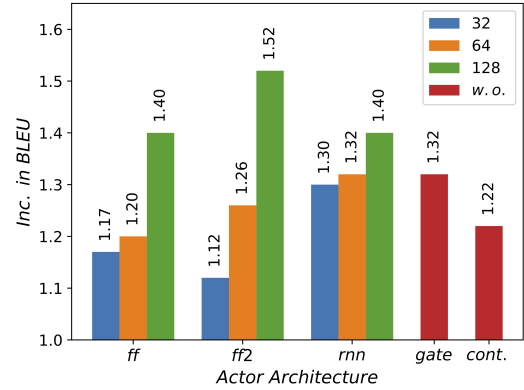


Figure 3: The effect of the actor architecture and hidden state size on trainable greedy decoding results over the IWSLT16 De-En validation set with Transformer (BLEU \uparrow), shown with a baseline (*cont.*) in which the underlying model, rather than the actor, is trained on the pseudo-parallel corpus. The Y-axis starts from 1.0. *w.o.* indicates an actor with no hidden layer. 0.0 corresponds to 33.04 BLEU.

Beam Size Figure 4a shows the effect of the beam size used to generate the pseudo-parallel corpus on the IWSLT16 De-En validation set with Transformer. Trainable greedy decoding improves over greedy decoding even when we set $k = 1$, namely, running *greedy* decoding on the augmented model to construct the new training corpus. With increased beam size k , the BLEU score consistently increases, but we observe diminishing returns beyond roughly $k = 35$, and we use that value elsewhere.

Training Corpus Construction There are a variety of ways one might use the output of beam search to construct a pseudo-parallel corpus: We could use the single highest-scoring output (by BLEU, or our target metric) for each input (*top1*), use all 35 beam search outputs (*full*), use all those outputs that score higher than the threshold, namely the base model’s greedy decoding output (*thd*), or combine the *top1* results with the gold-standard translations (*comb.*). We show the effect of training corpus construction in Figure 4b. *para* denotes the baseline approach of training the actor with the original parallel corpus used to train the underlying NMT model. Among the four novel approaches, *full* obtains the worst performance, since the beam search outputs contain translations that are far from the gold-standard translation. We choose the best performing *top1* strategy.

Obj.		BLEU \uparrow	MTR \uparrow	TER \downarrow
Greedy	-	27.15	29.0	54.4
Beam4	-	28.74	29.9	51.9
	sBLEU	<u>28.36</u>	<u>29.7</u>	52.0
	sMTR	<u>28.36</u>	29.6	51.8
	-sTER	28.05	29.6	<u>51.4</u>

Table 5: Results when trained with different decoding objectives on IWSLT16 De-En translation using Transformer. MTR denotes METEOR. We report greedy decoding and beam search ($k = 4$) results using the original model, and results with trainable greedy decoding (lower half).

Decoding Objectives As our approach is capable of using an arbitrary decoding objective, we investigate the effect of different objectives on BLEU, METEOR (MTR) and TER scores with Transformer for IWSLT16 De-En translation. Table 5 shows the final result on the test set. When trained with one objective, our model yields relatively good performance on that objective. For example, negative sentence-level TER (i.e., -sTER) leads to -3.0 TER improvement over greedy decoding and -0.5 TER improvement over beam search. However, since these objectives are all well correlated with each other, training with different objectives do not differ dramatically.

5 Related Work

Data Distillation Our work is directly inspired by work on knowledge distillation, which uses a similar pseudo-parallel corpus strategy, but aims at training a compact model to approximate the function learned by a larger model or an ensemble of models (Hinton et al., 2015). Kim and Rush (2016) introduce knowledge distillation in the context of NMT, and show that a smaller student network can be trained to achieve similar performance to a teacher model by learning from pseudo-corpus generated by the teacher model. Zhang et al. (2017) propose a new strategy to generate a pseudo-corpus, namely, fast sequence-interpolation based on the greedy output of the teacher model and the parallel corpus. Freitag et al. (2017) extend knowledge distillation on an ensemble and oracle BLEU teacher model. However, all these approaches require the expensive procedure of retraining the full student network.

Pseudo-Parallel Corpora in Statistical MT Pseudo-parallel corpora generated from beam search have been previously used in statistical

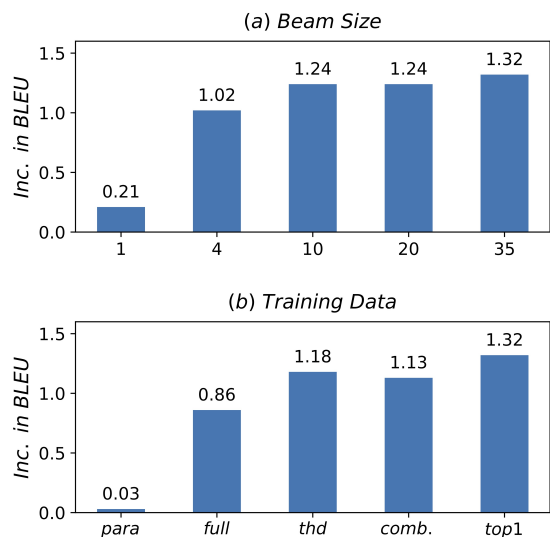


Figure 4: (a) The effect of beam size on the IWSLT16 De-En validation with Transformer and (b) the effect of the training corpus composition in the same setting. *para*: parallel corpus; *full*: all 35 beam search outputs; *thd*: beam search outputs that score higher than the base model’s greedy decoding output; *top1*: beam search output with the highest bleu score; *comb.*: *top1+para*. 0.0 corresponds to 33.04 BLEU.

machine translation (SMT) (Chiang, 2012; Gao and He, 2013; Auli and Gao, 2014; Dakwale and Monz, 2016). Gao and He (2013) integrate a recurrent neural network language model as an additional feature into a trained phrase-based SMT system and train it by maximizing the expected BLEU on k -best list from the underlying model. Our work revisits a similar idea in the context trainable greedy decoding for neural MT.

Decoding for Multiple Objectives Several works have proposed to incorporate different decoding objectives into training. Ranzato et al. (2015) and Bahdanau et al. (2016) use reinforcement learning to achieve this goal. Shen et al. (2016) and Norouzi et al. (2016) train the model by defining an objective-dependent loss function. Wiseman and Rush (2016) propose a learning algorithm tailored for beam search. Unlike these works that optimize the entire model, Li et al. (2017) introduce an additional network that predicts an arbitrary decoding objective given a source sentence and a prefix of translation. This prediction is used as an auxiliary score in beam search. All of these methods focus primarily on improving beam search results, rather than those with greedy decoding.

6 Conclusion

This paper introduces a novel method, based on an automatically-generated pseudo-parallel corpus, for training an actor-augmented decoder to optimize for greedy decoding. Experiments on three models and three datasets show that the training strategy makes it possible to substantially improve the performance of an arbitrary neural sequence decoder on any reasonable translation metric in either greedy or beam-search decoding, all with only a few trained parameters and minimal additional training time.

As our model is agnostic to both the model architecture and the target metric, we see the exploration of more diverse and ambitious model–target metric pairs as a clear avenue for future work.

Acknowledgments

This work was partly supported by Samsung Advanced Institute of Technology (Next Generation Deep Learning: from pattern recognition to AI), Samsung Electronics (Improving Deep Learning using Latent Structure) and the Facebook Low Resource Neural Machine Translation Award. KC thanks support by eBay, TenCent, NVIDIA and CIFAR. This project has also benefited from financial support to SB by Google and Tencent Holdings.

References

- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. Unsupervised neural machine translation. In *International Conference on Learning Representations*.
- Michael Auli and Jianfeng Gao. 2014. Decoder integration and expected bleu training for recurrent neural network language models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 136–142, Baltimore, Maryland. Association for Computational Linguistics.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Joseph Lowe, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. *preprint arXiv:1607.07086*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. 2013. Audio chord recognition with recurrent neural networks. In *Proceedings of the 14th International Society for Music Information Retrieval Conference*.
- Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. 2017. Massive exploration of neural machine translation architectures. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1442–1451, Copenhagen, Denmark. Association for Computational Linguistics.
- Yun Chen, Yang Liu, Yong Cheng, and Victor O.K. Li. 2017. A teacher-student framework for zero-resource neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1925–1935. Association for Computational Linguistics.
- David Chiang. 2012. Hope and fear for discriminative training of statistical translation models. *Journal of Machine Learning Research*, 13(Apr):1159–1187.
- Kyunghyun Cho. 2016. Noisy parallel approximate decoding for conditional recurrent language model. *preprint arXiv:1605.03835*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Praveen Dakwale and Christof Monz. 2016. Improving statistical machine translation performance by oracle-bleu model re-estimation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 38–44.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2018. Universal transformers. *preprint arXiv:1807.03819*.
- George Doddington. 2002. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proceedings of the second international conference on Human Language Technology Research*, pages 138–145. Morgan Kaufmann Publishers Inc.
- Markus Freitag, Yaser Al-Onaizan, and Baskaran Sankaran. 2017. Ensemble distillation for neural machine translation. *preprint arXiv:702.01802*.
- Jianfeng Gao and Xiaodong He. 2013. Training mrf-based phrase translation models using gradient ascent. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies*, pages 450–459, Atlanta, Georgia. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann Dauphin. 2017. Convolutional sequence to sequence learning. In *International Conference on Machine Learning*.
- Alex Graves. 2012. Sequence transduction with recurrent neural networks. *preprint arXiv:211.3711*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O. K. Li, and Richard Socher. 2018. Non-autoregressive neural machine translation. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Jiatao Gu, Kyunghyun Cho, and Victor O.K. Li. 2017. Trainable greedy decoding for neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1978, Copenhagen, Denmark. Association for Computational Linguistics.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *preprint arXiv:1503.02531*.
- Liang Huang, Kai Zhao, and Mingbo Ma. 2017. When to finish? optimal beam search for neural text generation (modulo beam size). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2134–2139.
- Yoon Kim and Alexander M. Rush. 2016. Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. Opennmt: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Alon Lavie and Michael J Denkowski. 2009. The meteor metric for automatic evaluation of machine translation. *Machine translation*, 23(2-3):105–115.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. A simple, fast diverse decoding algorithm for neural generation. *preprint arXiv:1611.08562*.
- Jiwei Li, Will Monroe, and Daniel Jurafsky. 2017. Learning to decode for future success. *preprint arXiv:1701.06549*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Mohammad Norouzi, Samy Bengio, zhifeng Chen, Navdeep Jaitly, Mike Schuster, Yonghui Wu, and Dale Schuurmans. 2016. Reward augmented maximum likelihood for neural structured prediction. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1723–1731. Curran Associates, Inc.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *preprint arXiv:1705.04304*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *preprint arXiv:1511.06732*.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. 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.

- Raphael Shu and Hideki Nakayama. 2017. Later-stage minimum bayes-risk decoding for neural machine translation. *preprint arXiv:1704.03169*.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of association for machine translation in the Americas*, volume 200.
- Zhaopeng Tu, Yang Liu, Lifeng Shang, Xiaohua Liu, and Hang Li. 2017. Neural machine translation with reconstruction.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30*.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pages 3156–3164. IEEE.
- Sam Wiseman and Alexander M. Rush. 2016. Sequence-to-sequence learning as beam-search optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, Austin, Texas. Association for Computational Linguistics.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, pages 2048–2057.
- Xiaowei Zhang, Wei Chen, Feng Wang, Shuang Xu, and Bo Xu. 2017. Towards compact and fast neural machine translation using a combined method. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1475–1481, Copenhagen, Denmark. Association for Computational Linguistics.