

Autoregressive Knowledge Distillation through Imitation Learning

Alexander Lin^{1*} Jeremy Wohlwend^{2*} Howard Chen^{3*} Tao Lei⁴
Harvard University MIT Princeton University ASAPP, Inc.
Cambridge, MA, USA Cambridge, MA, USA Princeton, NJ, USA New York, NY, USA

¹alexanderlin01@g.harvard.edu

²jwohlwend@csail.mit.edu

³howardchen@cs.princeton.edu

⁴tao@asapp.com

Abstract

The performance of autoregressive models on natural language generation tasks has dramatically improved due to the adoption of deep, self-attentive architectures. However, these gains have come at the cost of hindering inference speed, making state-of-the-art models cumbersome to deploy in real-world, time-sensitive settings. We develop a compression technique for autoregressive models that is driven by an imitation learning perspective on knowledge distillation. The algorithm is designed to address the exposure bias problem. On prototypical language generation tasks such as translation and summarization, our method consistently outperforms other distillation algorithms, such as sequence-level knowledge distillation. Student models trained with our method attain 1.4 to 4.8 BLEU/ROUGE points higher than those trained from scratch, while increasing inference speed by up to 14 times in comparison to the teacher model.

1 Introduction

Autoregressive models are ubiquitous in natural language processing. Due to the sequential nature of text generation, they are often the tool of choice for tackling sequence-to-sequence problems such as translation (Sutskever et al., 2014), summarization (Rush et al., 2015), and dialogue (Eric and Manning, 2017). Furthermore, they form the backbone of several successful generative pre-training architectures (Howard and Ruder, 2018; Peters et al., 2018; Radford et al., 2019; Dai et al., 2019).

Two recent trends have made autoregressive models cumbersome to deploy in real-world, natural language generation (NLG) applications. First, state-of-the-art models have grown larger and larger, amounting to hundreds of millions and even billions of parameters (Dong et al., 2019; Liu and

Lapata, 2019; Raffel et al., 2019). The increase in size and depth dramatically slows down inference speed. Second, the architecture of choice for autoregressive models seems to have shifted from the recurrent neural network (RNN) (Bahdanau et al., 2014; Luong et al., 2015) to the Transformer (Vaswani et al., 2017). Though the Transformer’s self-attention mechanism improves performance, it also increases the computational complexity of the step-by-step generation algorithms that are used at test time. Thus, both of these trends have contributed to significantly increasing inference time costs, especially on CPUs and low-resource devices, hindering their use in production systems.

Knowledge distillation (KD) (Bucilu et al., 2006; Hinton et al., 2015) is one popular method for compressing neural models. It transfers the information learned by a large, pretrained *teacher* model to a smaller, untrained *student* variant. In comparison to other methods such as weight pruning and quantization, KD allows the compressed model’s architecture to significantly differ from that of the original teacher. This enables models trained with KD to achieve high performance while meeting inference cost requirements (e.g. memory usage, inference speed, etc.).

Sequence-level knowledge distillation (SeqKD), proposed by Kim and Rush (2016), is the dominant technique for autoregressive KD in the current NLG literature, especially for machine translation (Gu et al., 2017; Ren et al., 2019; Zhou et al., 2019). This method trains a student model using a modified dataset generated by the teacher model and the standard negative log-likelihood objective. While SeqKD is simple and efficient, we argue that it does not take advantage of the teacher’s full potential.

Training the student model with a *static* dataset leads to the exposure bias problem. During training, the student model learns to predict the next token given previous tokens provided by the data.

* Work done while employed by ASAPP, Inc.

However, at inference time, the student generates the entire sequence from scratch by repeatedly using its own outputs as context for subsequent steps. This training-inference inconsistency causes a decrease in generation quality. Alternatively, we propose that the student can leverage the teacher in a *dynamic* fashion during the learning process.

We devise a new compression algorithm for autoregressive models called *imitation-based knowledge distillation* (ImitKD). It is inspired by an imitation learning (IL) perspective on the autoregressive distillation problem. Our algorithm trains a student model within an IL framework by treating the teacher as an oracle, and allows the student to explore its own generation during training. The teacher corrects the student’s generation at every time step, thereby guiding the student in learning how to generate.

Experimental results in translation and summarization show that ImitKD is especially suitable for compressing deep Transformer models that achieve high performance into shallow RNNs that generate up to 14 times faster at inference time. Our method consistently outperforms other distillation algorithms (such as word-level KD and sequence-level KD), and yields student models that beat models trained without a teacher by 1.4 to 4.8 points on generation metrics such as BLEU and ROUGE.

2 Background

2.1 Autoregressive Distillation

We introduce notation and formalize the task of autoregressive distillation. An *autoregressive model* π specifies a distribution over a T -dimensional *target sequence* $\mathbf{y} = \{y_1, \dots, y_T\} \in \mathcal{Y}$ by decomposing its joint distribution into a product of univariate conditionals:

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

where $\mathbf{y}_{<t}$ denotes $\{y_1, \dots, y_{t-1}\}$ for $t > 1$ and \emptyset for $t = 1$. The joint distribution over \mathbf{y} may itself be conditional on some related *source feature* $\mathbf{x} \in \mathcal{X}$ (e.g. translation, summarization) or not (e.g. language modeling). Since the former case can generalize the latter by letting $\mathcal{X} = \emptyset$, we will specify the presence of \mathbf{x} in the rest of the paper.

In autoregressive distillation, the goal is to learn a *student model* π that performs well at sequence generation by minimizing its loss with respect to a

pre-trained *teacher model* π^* . In many cases, the training objective can be formalized as:

$$\mathcal{L}(\pi) = \mathbb{E}_{\mathbf{y}|\mathbf{x} \sim \mathcal{D}} \left[\sum_{t=1}^T \ell^{\pi^*}(\mathbf{y}_{<t}, \mathbf{x}; \pi) \right], \quad (2)$$

where $\ell^{\pi^*}(\cdot; \pi)$ is the next-token loss function measuring the discrepancy between the teacher and student models given some prior context $\{\mathbf{y}_{<t}, \mathbf{x}\}$, and \mathcal{D} denotes a distribution (or dataset) of source-target pairs $\mathbf{x} \rightarrow \mathbf{y}$.

Due to the combinatorial nature of sequence generation, an autoregressive distillation method must maximize its learning efficiency by carefully selecting the distribution \mathcal{D} , i.e. how it explores the exponentially-sized space. We motivate this choice with the field of *imitation learning*, an active research area of reinforcement learning.

2.2 Distillation as Imitation Learning

Autoregressive text generation can be interpreted as a T -step Markov decision process (MDP). In particular, the autoregressive model π we wish to learn can be treated as a policy learner that maps a state to a distribution over actions. In our case, a state is a partial sequence $\mathbf{y}_{<t}$ for $t < T$, an action is the next token y_t , and the action space is the vocabulary. Given a state (partial sequence) and a chosen action (next token), the transition function is deterministic and simply concatenates them to form a new state (partial sequence).

The policy learner must be trained using some form of supervision. One option is to use reward-based reinforcement learning, which requires defining the numerical quality of a state/generation. However, for the autoregressive distillation problem, an arguably better choice is *imitation learning* (IL), which optimizes the policy by learning from demonstrations. In IL settings, an *oracle* policy π^* that is known to achieve high performance is provided during training. As a result, we can recast the overall goal as minimizing the divergence of the policy π from the oracle π^* . For example, it may be difficult to objectively define what it means for an aspiring translator to perform well (especially at the local token-by-token level). Yet, if we were given access to an expert translator, we could simply say the learner is performing well if they translate in the same way as the expert.

The IL framework is well suited for the setting of autoregressive distillation, since the student and teacher models naturally fill the respective roles of

the policy learner π and the oracle π^* . Thus, we can easily apply theoretical results and practical methods from the IL literature to the autoregressive distillation problem.

2.3 SeqKD as Behavioral Cloning

One distinguishing feature between different imitation learning methods pertains to how to define the state distribution \mathcal{D} in the training objective (Equation 2). Indeed, this is also one of the key design questions of autoregressive distillation. For instance, one simple and effective IL method is *behavioral cloning* (Ross and Bagnell, 2010), which obtains \mathcal{D} by running the oracle π^* on the MDP.

The popular sequence-level knowledge distillation (SeqKD) algorithm of Kim and Rush (2016) can be interpreted as behavioral cloning. For each source feature \mathbf{x} in the original training data, the teacher/oracle generates its (approximate) mode $\mathbf{y}^* = \arg \max_{\mathbf{y}'} \pi^*(\mathbf{y}' | \mathbf{x})$, typically using beam search. This new set of $\mathbf{x} \rightarrow \mathbf{y}^*$ pairs forms a teacher-generated dataset \mathcal{D}^* that serves as the state distribution for training the student. In addition, the negative log-likelihood of the teacher’s tokens $\mathbf{y}^* = \{y_1^*, \dots, y_T^*\}$ is used as the loss $\ell^{\pi^*}(\cdot; \pi)$. The overall training objective is:

$$\begin{aligned} \mathcal{L}_{\text{SeqKD}}(\pi) &= \mathbb{E}_{\mathbf{y}^* | \mathbf{x} \sim \mathcal{D}^*} \left[\sum_{t=1}^T -\log \pi(y_t^* | \mathbf{y}_{<t}^*, \mathbf{x}) \right]. \quad (3) \end{aligned}$$

The key advantage of SeqKD (as well as behavioral cloning) lies in its simplicity – we only need some samples from the teacher/oracle to work with. In comparison to vanilla supervised learning (which minimizes the negative log-likelihood of human-generated sequences), SeqKD has no additional training overhead other than the creation of the new dataset \mathcal{D}^* .

However, the simplicity of the algorithm also limits its potential. Ross and Bagnell (2010) argued that training a policy π via behavioral cloning incurs regret with respect to the oracle π^* that is a quadratic function of the time horizon T . Intuitively, behavioral cloning suffers from the *exposure bias problem*. During training, the student model learns to perform good actions for the teacher/oracle’s state distribution \mathcal{D}^* , but is never exposed to its own states. Thus, during testing (when the student must walk an MDP of self-generated states), the step-by-step errors compound over time, resulting in suboptimal generations.

We argue that in autoregressive distillation, the teacher/oracle can do more than produce a static dataset. It is a dynamic entity capable of interacting with the student throughout training. By querying the teacher with its own states, the student has the opportunity to ameliorate exposure bias and learn how to generate.

3 Imitation-Based Distillation Algorithm

In this section, we present our IL-based algorithm for autoregressive distillation. We begin by describing the key design principles and why we expect them to work well. Then, we elaborate on the algorithm’s implementation in detail.

3.1 Design Principles and Rationale

One key principle of our algorithm is that the student model must be trained on its own state distribution so that it will perform better at generation. In practice, we achieve this by sampling training examples from $\tilde{\mathcal{D}}$, a mixture of an initial distribution \mathcal{D} (e.g. a static training set) and the distribution \mathcal{D}_π of generations from the student π . We use \mathcal{D} to alleviate the cold-start problem, in which an untrained π generates poorly at the start of training.

This idea builds upon the empirical and theoretical foundation of dataset aggregation (DAgger), one of the most popular imitation learning methods that improve upon behavioral cloning. DAgger (Ross et al., 2011) successively populates its training set by adding new data generated from the oracle-learner mixture. It then re-trains the policy learner on the aggregated dataset at each iteration. Under some assumptions (such as the loss function being strongly convex in π), Ross et al. (2011) proved that DAgger yields a policy π that has linear regret in T with respect to π^* . This is a significant improvement over the behavior cloning result and can be attributed to fixing exposure bias. We expect a similar strategy of mixing oracle and learner distributions to work well for non-convex neural networks, as shown in other applications (Zhang and Cho, 2016; Sun et al., 2017).

Another key principle of our algorithm is that the teacher model should play the role of the oracle and correct the student’s generations at each time step. In order for this training strategy to be successful, the teacher must be able to provide better actions than the student for the student’s own states.

To resolve this question, we perform a test in which a deep Transformer-based translation model

Decoding Method	Bleu \uparrow
Transformer only	33.8
RNN only	28.6
RNN first, Transformer completes	31.4

Table 1: Greedy decoding BLEU scores on the IWSLT validation set for the preliminary test.

completes the partial translations of a shallow RNN-based model. As shown in Table 1, the Transformer completions achieve much higher BLEU score than the RNN’s full generations. This validates our assumption that a strong teacher model can indeed play the role of the oracle and guide the student to better states.

3.2 The ImitKD Algorithm

Our imitation-based knowledge distillation algorithm (ImitKD) is given in Algorithm 1. The central training objective is:

$$\mathcal{L}_{\text{ImitKD}}(\pi) = \mathbb{E}_{\mathbf{y}|\mathbf{x}\sim\tilde{\mathcal{D}}} \left[\sum_{t=1}^T \ell^{\pi^*}(\mathbf{y}_{<t}, \mathbf{x}; \pi) \right] \quad (4)$$

where $\tilde{\mathcal{D}}$ is the data mixture defined by sampling from the initial dataset \mathcal{D} and generating with the student (lines 8-11). The probability $\beta_i \in [0, 1]$ (line 8) controls how often an example comes from \mathcal{D} . The loss function ℓ^{π^*} can be realized as the negative log-likelihood of the oracle’s optimal next token/action,

$$\ell_{\text{opt}}^{\pi^*}(\mathbf{y}_{<t}, \mathbf{x}; \pi) = -\log \pi(v^* | \mathbf{y}_{<t}, \mathbf{x}),$$

where $v^* = \arg \max_{v \in \mathcal{V}} \pi^*(v | \mathbf{y}_{<t}, \mathbf{x})$. Alternatively, ℓ^{π^*} can be the cross-entropy loss between the full distributions,

$$\begin{aligned} \ell_{\text{full}}^{\pi^*}(\mathbf{y}_{<t}, \mathbf{x}; \pi) \\ = - \sum_{v \in \mathcal{V}} \pi^*(v | \mathbf{y}_{<t}, \mathbf{x}) \cdot \log \pi(v | \mathbf{y}_{<t}, \mathbf{x}). \end{aligned}$$

Next, we describe some practical implementations in order to make Algorithm 1 suitable for compressing deep learning systems. One limitation of dataset aggregation is that the amount of training examples keeps growing, making each iteration successively more expensive. As an alternative to aggregation, we perform *data replacement* within each training batch.

As shown in Algorithm 1, we treat each mini-batch $\tilde{\mathcal{D}}_i$ as a new iteration of the dataset and perform a single step of stochastic gradient descent

Algorithm 1 Imitation-Based Distillation

- 1: Let \mathcal{D} be initial dataset.
- 2: Initialize π_1 at random.
- 3: **for** $i = 1, \dots, I$ **do**
- 4: Initialize new dataset $\tilde{\mathcal{D}}_i = \emptyset$.
- 5: **repeat** B **times**
- 6: Sample an example $e = \mathbf{y} | \mathbf{x} \sim \mathcal{D}$.
- 7: Sample uniformly $u \sim [0, 1]$.
- 8: **if** $u > \beta_i$ **then**
- 9: Generate $\hat{\mathbf{y}}$ from π_i given \mathbf{x} .
- 10: Replace example with $e = \hat{\mathbf{y}} | \mathbf{x}$.
- 11: **end if**
- 12: Append example e to $\tilde{\mathcal{D}}_i$.
- 13: Compute $\mathcal{L}_{\text{imitKD}}(\pi_i)$ on $\tilde{\mathcal{D}}_i$ with π^* .
- 14: Let $\pi_{i+1} = \pi_i - \alpha_i \cdot \partial \mathcal{L}_{\text{imitKD}} / \partial \pi_i$.
- 15: **end for**
- 16: **return** Best policy π on validation set.

on $\mathcal{L}_{\text{imitKD}}$ (Equation 4) with respect to the parameters of the previous model π_i to yield π_{i+1} . Thus, the number of iterations I becomes the number of mini-batches used to train the student model.

Our practical algorithmic changes are inspired by theory. The dataset aggregation algorithm (Ross et al., 2011) achieves its regret bounds because it reduces to the Follow-the-Leader algorithm for online learning (Kakade et al., 2009). Our training paradigm can be similarly interpreted as an online gradient descent algorithm, which has comparable guarantees for strongly convex losses (Hazan et al., 2007) and even certain non-strongly convex losses (Garber, 2019). Variants of this paradigm have also been employed in other deep learning work (Bengio et al., 2015; Sun et al., 2017).

3.3 Data Mixture Selection and Annealing

Dataset replacement requires an initial dataset that can be potentially replaced at each step. A natural candidate for this initial dataset is the original supervised training data (denoted as \mathcal{D}'), which can be interpreted as a collection of samples from a human oracle. Alternatively, we can use the SeqKD dataset \mathcal{D}^* , which has generations from the teacher.

If we take samples from \mathcal{D}' or \mathcal{D}^* and replace some of them with student-generated samples, we effectively create a teacher-student dataset mixture. Unlike DAGger, this mixture occurs at the sequence level instead of the token/state level. An advantage of sequence-level mixtures is that they do not require generating with the teacher during each

training iteration, which can be quite expensive if the teacher is a large neural network. Instead, the teacher only needs to compute the batched loss, which is comparatively much cheaper.

The exact mixing schedule β_1, \dots, β_I is a customizable feature of Algorithm 1. Empirically, we have found an exponential decay to work well:

$$\beta_i = r^{i/I}, \quad (5)$$

where $r \in (0, 1]$ is the final mixing rate.

3.4 Speeding Up Training

Generating sequences \hat{y} on the fly at every iteration (line 9) can be a major computation bottleneck during training. We speed up this step by generating a pool of $B \cdot M$ examples in parallel only once every M iterations, where B is the batch size and M is a hyperparameter. One caveat of this modification is that at iteration i , the loss function may no longer be computed on examples generated by the most recent set of model parameters, but rather parameters from up to M iterations prior. Nonetheless, we have found that setting M to a small integer (e.g. 2-8) can speed up training time without impacting final model performance.

We use a greedy decoding or top- K sampling with small K to produce samples \hat{y} (line 9) in our algorithm. These two strategies are efficient to run, operate similarly to the generation employed at inference time, and have empirically worked well in our experiments. Of course, the generation strategy can be customized for different tasks.

4 Related Work

The distillation problem for autoregressive models was first tackled by Kim and Rush (2016), who introduced sequence-level knowledge distillation for neural machine translation. Subsequent works have used seqKD for non-autoregressive translation models (Gu et al., 2017; Zhou et al., 2019), low-resource settings (Chen et al., 2017), and ensemble distillation with multiple teachers (Kuncoro et al., 2016; Tan et al., 2019). Wei et al. (2019) proposed a behavioral cloning method for distilling autoregressive translation models into non-autoregressive translation models. In contrast, our method aims to address the learning challenges in autoregressive distillation, such as exposure bias.

Various methods other than standard supervised learning have been explored for training generative models of language. MIXER (Ranzato et al., 2015)

and Beam Search Optimization (Wiseman and Rush, 2016) also perform generation during training, but use sequence-level metrics (e.g. BLEU score) as training supervision. Similarly, SEARNN (Leblond et al., 2017) trains RNNs to iteratively generate sequences with beam search to compute the local loss of a single action during the decoding process. Scheduled sampling (Bengio et al., 2015) and its extensions (Goyal et al., 2017; Zhang et al., 2019) alleviate exposure bias by replacing some words in the true context with the model’s prediction. However, without a dynamic queryable oracle, these methods face the challenge of properly defining the training signal when the generated sequence no longer exists in the static training data. For example, directly reusing the tokens in the static dataset as the target next token leads to an inconsistent training procedure (Huszár, 2015). In contrast to these methods, distillation can fully leverage the teacher oracle, allowing us to design a simple and efficient imitation learning algorithm.

5 Experimental Setup

We test our autoregressive distillation method and all baselines on three language generation tasks – IWSLT 2014 German \rightarrow English translation, WMT 2016 English \rightarrow German translation, and CNN/DailyMail abstractive news summarization.

Datasets The IWSLT 2014 De \rightarrow En dataset consists of approximately 170K sequence pairs. Following standard practice (Bahdanau et al., 2016; Deng et al., 2018; Wang et al., 2019), we randomly sample 4% of this dataset as the validation set and let the remaining be the training set. The test set is the concatenation of the dev2010, tst2010, tst2011, and tst2012 files. We use a shared vocabulary of 14K lowercased BPE tokens (Sennrich et al., 2015).

The WMT 2016 En \rightarrow De dataset has 4.5 million training pairs. We use the same preprocessing of the prior work (Ott et al., 2018), newstest2013 as the validation set and newstest2014 as the test set. The vocabulary consists of 32K cased BPE tokens.

The CNN/DailyMail summarization dataset has 287K, 13K and 12K pairs in the training, validation and test sets, respectively. Following prior work (See et al., 2017), we truncate documents to 400 tokens and summaries to 100 tokens in the training set. During evaluation, we generate up to 128 tokens. We use a pre-trained BERT (Devlin et al., 2018) tokenizer with a vocabulary of 30K lowercased tokens (Liu and Lapata, 2019).

Variant	Context (States)	Next Token	Next Token Loss Function
Vanilla	Data	Data	$\ell_{\text{opt}}^{\pi^*}$ with v^* from data
SeqKD	Teacher	Teacher	$\ell_{\text{opt}}^{\pi^*}$ with v^* from beam search
ImitKD	Student/Data Mix	Teacher	$\ell_{\text{opt}}^{\pi^*}$ with $v^* = \text{argmax}_v \pi^*(v \mathbf{y}_{<t}, \mathbf{x})$
ImitKD*	Student/Teacher Mix	Teacher	$\ell_{\text{opt}}^{\pi^*}$ with $v^* = \text{argmax}_v \pi^*(v \mathbf{y}_{<t}, \mathbf{x})$
Vanilla + Full	Data	Teacher	$\ell_{\text{full}}^{\pi^*}$
SeqKD + Full	Teacher	Teacher	$\ell_{\text{full}}^{\pi^*}$
ImitKD + Full	Student/Data Mix	Teacher	$\ell_{\text{full}}^{\pi^*}$
ImitKD* + Full	Student/Teacher Mix	Teacher	$\ell_{\text{full}}^{\pi^*}$

Table 2: Summary of training variants. All “+ Full” variants are trained with the full teacher-student cross entropy.

Models Transformers often attain state-of-the-art performance on common language generation tasks. On the other hand, RNNs (without self-attention) generate much faster at inference time. Thus, from a practitioner’s standpoint, it may be most desirable to compress a high-performing Transformer into a lightweight RNN. For all tasks, we use the state-of-the-art Transformer architecture (Vaswani et al., 2017) as the teacher model. The teacher models are trained using vanilla supervised learning. For WMT, we directly use the pre-trained Transformer model provided by the Fairseq library (Ott et al., 2018, 2019).

In all tasks, we use a recurrent neural network, specifically SRU (Lei et al., 2017), as the student model. For completeness, we also train Transformer, GRU (Cho et al., 2014), and LSTM (Hochreiter and Schmidhuber, 1997) based student models on the IWSLT translation task, illustrating the effectiveness of our distillation method for various neural architectures. All RNN-based models follow the seq2seq, encoder-decoder architecture (Sutskever et al., 2014) and employ a single scaled dot-product attention between the encoder and decoder (Bahdanau et al., 2014; Luong et al., 2015).

All models are trained using the Adam optimizer (Kingma and Ba, 2014) with an inverse-square-root learning rate scheduler and learning rate warmup (Vaswani et al., 2017). Our experiments were conducted using Flamb, a PyTorch-based model training and evaluation library (Wohlgend et al., 2019). More implementation details such as hyperparameter settings are provided in Appendix A.

Variants For the student models, we compare a wide range of training variants, including baselines such as vanilla supervised learning (which directly uses the original training set) and sequence-level knowledge distillation (SeqKD). All SeqKD vari-

ants form the teacher-generated dataset using beam search with beam size $K = 5$. For our imitation-based method, we experiment with annealing from the original training set (ImitKD) or the teacher-generated SeqKD dataset (ImitKD*). We also experiment with different token-level losses; base variants are trained with the optimal next token while “+ Full” variants are trained with the full cross entropy. Table 2 summarizes all variants and highlights their differences. Note that the Vanilla + Full baseline – referred to as “WordKD” by Kim and Rush (2016) – has appeared in other distillation works (e.g. Tan et al., 2019; Sanh et al., 2019).

Evaluation We use BLEU score (Papineni et al., 2002) for translation and report ROUGE-1, ROUGE-2 and ROUGE-L scores (Lin, 2004) for summarization. For all models, the training checkpoint with the highest BLEU/ROUGE-1 score on the validation set is used for test set evaluation. We also report the perplexity metric for all tasks.

6 Results

IWSLT De→En Translation Table 3 compares all distillation methods on the IWSLT dataset. The teacher model is an 8-layer Transformer. We use a 3-layer SRU, a 2-layer SRU and a 2-layer Transformer as student models. For all three student models, our ImitKD method outperforms all baselines in terms of BLEU score with beam size 1 (Bleu_1), BLEU score with beam size 5 (Bleu_5) and perplexity (PPL). The improvement on Bleu score ranges from 1.4 to 4.8 points compared to the Vanilla training method. The 3-layer SRU model trained with ImitKD + Full even slightly exceeds the performance of the teacher model. Furthermore, our method consistently outperforms SeqKD by up to 1.4 BLEU, highlighting the benefit of training the student model with its own state distribution.

Variant	PPL ↓	Bleu ₁ ↑	Bleu ₅ ↑
Transf. (8-layer)			
Teacher	5.6	34.4	35.2
SRU (3-layer)			
Vanilla	7.4	30.0	31.2
SeqKD	153.0	33.0	33.1
ImitKD	14.7	34.1	34.4
Vanilla + Full	5.4	34.2	34.8
SeqKD + Full	6.1	34.3	34.8
ImitKD + Full	5.3	34.8	35.4
SRU (2-layer)			
Vanilla	7.4	29.5	30.6
SeqKD	102.1	32.0	32.4
ImitKD	12.7	33.3	33.7
Vanilla + Full	6.0	33.0	33.8
SeqKD + Full	6.8	33.1	33.7
ImitKD + Full	5.7	33.7	34.5
Transf. (2-layer)			
Vanilla	6.4	32.8	33.4
SeqKD	23.4	34.0	34.2
ImitKD	7.5	34.3	34.6
Vanilla + Full	5.9	33.8	34.2
SeqKD + Full	7.1	34.0	34.4
ImitKD + Full	5.9	34.4	34.8

Table 3: Results on IWSLT test dataset.

Variant	SRU	GRU	LSTM	Transf.
Vanilla	28.6	28.6	27.7	32.4
SeqKD	31.4	31.2	30.5	33.3
ImitKD	32.7	32.7	32.4	33.7

Table 4: BLEU scores of different student architectures on the IWSLT validation set. We use a beam size of 1.

To further demonstrate the effectiveness of ImitKD across different model types, we report validation set Bleu₁ for various 2-layer neural architectures in Table 4. Our ImitKD method outperforms the baselines in all cases, with the gains being especially large for recurrent architectures.

WMT En→De Translation We present the results on the WMT dataset in Table 5. The teacher is a 6-layer Transformer and the student is a 4-layer SRU. Here, we see that ImitKD performs closer to SeqKD. These results reveal that direct behavioral cloning (SeqKD) can be quite effective when the amount of oracle demonstrations is sufficiently high, e.g. several millions of examples. Nonetheless, ImitKD and ImitKD* can improve on SeqKD by training the student not only with the teacher’s

Variant	PPL ↓	Bleu ₁ ↑	Bleu ₅ ↑
Teacher	3.2	28.7	29.2
Vanilla	5.5	22.0	23.1
SeqKD	9.0	24.9	25.5
ImitKD	7.4	24.6	25.5
ImitKD*	8.4	25.3	25.8
Vanilla + Full	5.2	23.8	24.5
SeqKD + Full	5.6	24.7	25.3
ImitKD + Full	5.6	25.3	25.9
ImitKD* + Full	5.6	25.0	25.8
△Vanilla	18.7	13.8	15.1
△SeqKD	42.3	17.1	17.9
△ImitKD	15.0	17.8	19.0
△ImitKD*	17.8	18.6	19.5

Table 5: Results on WMT dataset. ImitKD* is trained on a student/teacher dataset mixture. △ indicates that the model is trained with 25× less data.

Variant	PPL ↓	R1 ↑	R2 ↑	RL ↑
Teacher	12.5	39.0	17.6	35.7
Vanilla	14.7	36.1	15.6	32.8
SeqKD	52.9	36.4	16.1	33.1
ImitKD	17.2	37.3	16.4	34.1
ImitKD*	37.1	37.7	16.7	34.5
Vanilla + Full	13.6	36.2	16.0	32.9
SeqKD + Full	20.2	37.4	16.5	34.0
ImitKD + Full	14.0	38.4	17.1	34.9
ImitKD* + Full	17.9	38.1	17.1	34.6

Table 6: Results on CNN/DailyMail dataset. All models generate using beam search $K = 5$ decoding.

states but also with its own states. Among all variants, ImitKD + Full performs the best while avoiding the overhead of creating a teacher-modified dataset. Furthermore, we see that ImitKD is especially effective in low-data regimes. As shown in the bottom block of Table 5, ImitKD methods achieve much stronger results over baselines when we reduce the WMT training data to the same size as IWSLT.

CNN/DailyMail Summarization In Table 6, we present the CNN/DailyMail results using a 6-layer Transformer as the teacher and a 2-layer SRU as the student. Once again, the best student is ImitKD + Full, which achieves ROUGE scores that are within 1 point of the teacher’s. We see that ImitKD variants outperform the baselines on all ROUGE metrics, showcasing the utility of our method on a different NLG task.

Task	Model	Architecture	Size	% Compress	CPU Time	× Faster
IWSLT	Teacher	Transf. (8-layer)	20.1 M	—	269 ms / 518 ms	—
	Student	SRU (3-layer)	14.0 M	68.7%	55 ms / 97 ms	4.9 / 5.3
	Student	SRU (2-layer)	8.6 M	42.7%	37 ms / 56 ms	7.2 / 9.3
	Student	Transf. (2-layer)	8.5 M	42.3%	78 ms / 144 ms	3.4 / 3.6
WMT	Teacher	Transf. (6-layer)	209.9 M	—	816 ms / 1466 ms	—
	Student	SRU (4-layer)	34.2 M	16.3%	174 ms / 306 ms	4.7 / 4.8
CNN/DM	Teacher	Transf. (6-layer)	59.8 M	—	1900 ms / 12138 ms	—
	Student	SRU (2-layer)	14.4 M	24.1%	258 ms / 826 ms	7.4 / 14.7

Table 7: Model types, along with statistics on size (in number of parameters) and CPU inference time per decoding (in milliseconds) for beam search with beam size $K \in \{1, 5\}$.




Figure 1: Bleu score versus sequence decoding length across different models and training variants. Each point on the graph represents the Bleu score of all sequences whose length is within a bin of width 20.

Size and Speed Analysis In Table 7, we analyze how our distillation technique can reduce computational costs, using the IWSLT (Table 3), WMT (Table 5), and CNN/DailyMail (Table 6) teacher/student pairs as case studies. By training small student models with ImitKD, we can substantially decrease model size and increase inference speed, while minimizing performance loss. Shallow, recurrent architectures are especially attractive, because they can generate 4-14 times faster than deep Transformer teachers, and 2-3 times faster than Transformer students of similar size.

Performance Analysis at Different Lengths

Figure 1 breaks down BLEU score vs. decoding length for IWSLT models trained with different algorithms (Vanilla, SeqKD, ImitKD). We show results for the three different types of RNNs (i.e. SRU, GRU, and LSTM) and Transformer in Table 4. All models have two layers.

As expected, we observe that the generation quality (in terms of BLEU score) degrades as the decoding length increases. This phenomenon can be explained by the global error compounding with each additional decision step (Ross et al., 2011) and has

been reported in previous works (Bahdanau et al., 2014; Zhang et al., 2019). As shown in Figure 1, models trained with the vanilla objective, especially RNN-based models, suffer the most from this problem. SeqKD improves the performance across all sequence lengths, but still experiences some BLEU score degradation for longer sequences. ImitKD further improves the BLEU score across all bins, and more importantly, the improvement is most significant for longer sequences. This analysis suggests that ImitKD explicitly addresses the exposure bias problem for training student models.

7 Conclusion

In this work, we developed a new knowledge distillation technique inspired by imitation learning for compressing large and cumbersome autoregressive models. We demonstrated the empirical success of our method over popular baselines on several natural language generation tasks.

Although our experiments focused on sequence-to-sequence settings, one future possibility is to explore using ImitKD for compressing large language models aimed at transfer learning.

References

- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.
- Cristian Bucilu, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541.
- Yun Chen, Yang Liu, Yong Cheng, and Victor OK Li. 2017. A teacher-student framework for zero-resource neural machine translation. *arXiv preprint arXiv:1705.00753*.
- 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. *arXiv preprint arXiv:1406.1078*.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*.
- Yuntian Deng, Yoon Kim, Justin Chiu, Demi Guo, and Alexander Rush. 2018. Latent alignment and variational attention. In *Advances in Neural Information Processing Systems*, pages 9712–9724.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *arXiv preprint arXiv:1905.03197*.
- Mihail Eric and Christopher D Manning. 2017. A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue. *arXiv preprint arXiv:1701.04024*.
- Dan Garber. 2019. Logarithmic regret for online gradient descent beyond strong convexity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 295–303.
- Kartik Goyal, Chris Dyer, and Taylor Berg-Kirkpatrick. 2017. Differentiable scheduled sampling for credit assignment. *arXiv preprint arXiv:1704.06970*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Elad Hazan, Amit Agarwal, and Satyen Kale. 2007. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- Ferenc Huszár. 2015. How (not) to train your generative model: Scheduled sampling, likelihood, adversary? *arXiv preprint arXiv:1511.05101*.
- Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. 2009. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, pages 793–800.
- Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Adhiguna Kuncoro, Miguel Ballesteros, Lingpeng Kong, Chris Dyer, and Noah A Smith. 2016. Distilling an ensemble of greedy dependency parsers into one mst parser. *arXiv preprint arXiv:1609.07561*.
- Rémi Leblond, Jean-Baptiste Alayrac, Anton Osokin, and Simon Lacoste-Julien. 2017. Searnn: Training rnns with global-local losses. *arXiv preprint arXiv:1706.04499*.
- Tao Lei, Yu Zhang, Sida I Wang, Hui Dai, and Yoav Artzi. 2017. Simple recurrent units for highly parallelizable recurrence. *arXiv preprint arXiv:1709.02755*.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. *arXiv preprint arXiv:1904.01038*.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. *arXiv preprint arXiv:1806.00187*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Ofir Press and Lior Wolf. 2016. Using the output embedding to improve language models. *arXiv preprint arXiv:1608.05859*.
- Alec Radford, Jeffrey Wu, Dario Amodei, Daniela Amodei, Jack Clark, Miles Brundage, and Ilya Sutskever. 2019. Better language models and their implications. *OpenAI Blog <https://openai.com/blog/better-language-models>*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2019. FastSpeech: Fast, robust and controllable text to speech. In *Advances in Neural Information Processing Systems*, pages 3165–3174.
- Stéphane Ross and Drew Bagnell. 2010. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. *arXiv preprint arXiv:1509.00685*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Wen Sun, Arun Venkatraman, Geoffrey J Gordon, Byron Boots, and J Andrew Bagnell. 2017. Deeply aggravated: Differentiable imitation learning for sequential prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3309–3318. JMLR. org.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*.
- 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*, pages 5998–6008.
- Dilin Wang, Chengyue Gong, and Qiang Liu. 2019. Improving neural language modeling via adversarial training. *arXiv preprint arXiv:1906.03805*.
- Bingzhen Wei, Mingxuan Wang, Hao Zhou, Junyang Lin, and Xu Sun. 2019. Imitation learning for non-autoregressive neural machine translation. *arXiv preprint arXiv:1906.02041*.
- Sam Wiseman and Alexander M Rush. 2016. Sequence-to-sequence learning as beam-search optimization. *arXiv preprint arXiv:1606.02960*.
- Jeremy Wohlwend, Nicholas Matthews, and Ivan Itzcovich. 2019. Flambé: A customizable framework for machine learning experiments. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 181–188.

Jiakai Zhang and Kyunghyun Cho. 2016. Query-efficient imitation learning for end-to-end autonomous driving. *arXiv preprint arXiv:1605.06450*.

Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. Bridging the gap between training and inference for neural machine translation. *arXiv preprint arXiv:1906.02448*.

Chunting Zhou, Graham Neubig, and Jiatao Gu. 2019. Understanding knowledge distillation in non-autoregressive machine translation. *arXiv preprint arXiv:1911.02727*.

A Appendices

A.1 DAgger Algorithm

The dataset aggregation (DAgger) algorithm (Ross et al., 2011) minimizes the following objective:

$$\mathcal{L}_{\text{Imit}}(\pi) = \mathbb{E}_{s_1, \dots, s_T \sim \mathcal{D}} \left[\sum_{t=1}^T \ell^{\pi^*}(s_t; \pi) \right], \quad (6)$$

where \mathcal{D} is a distribution (or dataset) of T -step state trajectories and $\ell^{\pi^*}(s, \pi)$ is the action-discrepancy loss between the oracle π^* and the policy learner π in state s . The full DAgger algorithm is given in Algorithm 2.

Algorithm 2 Dataset Aggregation

- 1: Let $\mathcal{D} = \emptyset$ be initial dataset.
 - 2: Initialize π_1 at random.
 - 3: **for** $i = 1, \dots, I$ **do**
 - 4: Let mixture policy $\tilde{\pi}_i = \beta_i \pi^* + (1 - \beta_i) \pi_i$.
 - 5: Initialize new dataset $\mathcal{D}^i = \emptyset$.
 - 6: **repeat** B **times**
 - 7: Run MDP on $\tilde{\pi}_i$, sample $\{s_1, \dots, s_T\}$.
 - 8: Append new states $\{s_1, \dots, s_T\}$ to \mathcal{D}^i .
 - 9: Aggregate $\mathcal{D} = \mathcal{D} \cup \mathcal{D}^i$.
 - 10: Train π_{i+1} on \mathcal{D} to min $\mathcal{L}_{\text{Imit}}$ with π^* .
 - 11: **end for**
 - 12: **return** Best policy π on validation set.
-

A.2 Implementation Details

In all experiments, all RNN-based models with hidden dimension N consist of a bidirectional encoder with hidden dimension $N/2$ and a left-to-right decoder with hidden dimension N .

For BLEU score evaluation, we use the NLTK library.¹ For ROUGE score evaluation, we use the py-rouge library.²

Preliminary Study For Table 1, we train both an 8-layer Transformer and a 2-layer RNN (specifically SRU) on the IWSLT dataset using standard supervised learning. The architectural and training details are the same as those outlined in the IWSLT experiments. At test time, both the Transformer and the RNN perform greedy decoding. On average, ground-truth translations in the IWSLT test set have 24.5 tokens. The ‘‘RNN first, Transformer completes’’ mixed decoding strategy generates 12

tokens (i.e. half on average) with the RNN and the rest with the Transformer. We measure generation quality using Bleu score.

IWSLT The IWSLT 2014 German \rightarrow English dataset is taken directly from the source website.³

We train an 8-layer Transformer teacher model with model dimension 256, feedforward dimension 1024, and four attention heads as the teacher model. The 2-layer student SRU model has a hidden dimension 512, and the 3-layer model has hidden dimension 1024 and projection dimension 256. The student Transformer model has model dimension 256, feedforward dimension 768 and 4 attention heads.

All models have word embedding dimension 256 and exhibit weight tying between the decoder embeddings and the output layer (Press and Wolf, 2016). We train models for 80K steps with batch size 128 using the Adam optimizer with base learning rate 0.1. We use an inverse-square-root learning rate scheduler (Vaswani et al., 2017) with 10K warmup steps for the teacher and 5K warmup steps for all students. Validation set metrics are recorded every 1K steps. For all ImitKD variants, we set the final mixing rate $r = 0.005$ (i.e. very close to 0), and use top- K sampling with $K = 5$ as the generation algorithm during training. We use $M = 4$ as the batch parallelization parameter.

In Table 4, the 2-layer SRU and the 2-layer Transformer follow the same architecture as those in Table 3. To standardize architecture across RNNs, the GRU and the LSTM have the same embedding dimension (i.e. 256) and hidden dimension (i.e. 512) as the SRU.

WMT The WMT 2016 dataset is taken from the Fairseq library.⁴

We use a pre-trained Transformer-large model from the Fairseq library (Ott et al., 2018, 2019) as our teacher model. It has embedding dimension 1024, model dimension 1024, and feedforward dimension 4096. The student is a 4-layer SRU with hidden size 1024, projection size 256, and embedding size 256. The student is trained for 15 epochs with batch size 512, base learning rate 0.1, and 4K warmup steps. We record validation metrics every 1/4 of the epoch. The encoder embeddings, decoder embeddings, and decoder output layer share

¹https://www.nltk.org/_modules/nltk/translate/bleu_score.html

²<https://github.com/Diego999/py-rouge>

³<https://sites.google.com/site/iwslt2014evaluation2014/data-provided>

⁴<https://github.com/pytorch/fairseq/tree/master/examples/translation>

the same weight parameters. We tune the final mixing rate $r \in \{0.5, 0.1, 0.005\}$ for our ImitKD variants.

CNN/Dailymail The CNN/DailyMail dataset is taken from Professor Kyunghyun Cho’s website, a commonly used source for this dataset.⁵ The teacher model is a 6-layer Transformer-base model with embedding dimension 512, model dimension 512, and feedforward dimension 2048. The student is a 2-layer SRU with embedding dimension 256, hidden size 1024, and projection size 256. We use a batch size of 128. For both models, the learning rate follows an inverse-square root schedule with warmup of 2K steps. Validation set metrics are recorded every 2K steps. The teacher has a base learning rate of 0.03, while the student has a base learning rate of 0.1. The teacher benefits from larger effective batch sizes by accumulating gradients every eight steps. On the other hand, the student does not seem to benefit from gradient accumulation and therefore takes a gradient step after processing each batch. All ImitKD variants use final mixing rate $r = 0.1$ and greedy decoding during training. We use $M = 4$ as the batch parallelization parameter.

Size and Speed Analysis CPU generation times for all models were measured on a 2019 MacBook Pro with a 2.6GHz 6-core Intel Core i7 processor. Time estimates reported in Table 7 were averaged over examples in the test set of the corresponding dataset.

Performance Analysis at Different Lengths

For each IWSLT variant, we ran greedy decoding (i.e. beam search decoding with beam size $K = 1$) on the test set. Then, we sorted the decoded sequences by length into the following bins: [0, 20], [21, 40], [41, 60], [61, 80], [81, 100], [101, 120]. Each point in Figure 1 is the Bleu score of all sequences within one of these bins for the corresponding IWSLT variant.

⁵<https://cs.nyu.edu/~kcho/DMQA/>