

Integrated Training for Sequence-to-Sequence Models Using Non-Autoregressive Transformer

Evgeniia Tokarchuk^{1,2}, Jan Rosendahl², Weiyue Wang², Pavel Petrushkov¹, Tomer Lancewicki¹, Shahram Khadivi¹, and Hermann Ney²

¹eBay Inc., Aachen, Germany

²RWTH Aachen University, Aachen, Germany

e.tokarchuk@uva.nl

{ppetrushkov, tlancewicki, skhadivi}@ebay.com

{rosendahl, wwang, ney}@cs.rwth-aachen.de

Abstract

Complex natural language applications such as speech translation or pivot translation traditionally rely on cascaded models. However, cascaded models are known to be prone to error propagation and model discrepancy problems. Furthermore, there is no possibility of using end-to-end training data in conventional cascaded systems, meaning that the training data most suited for the task cannot be used. Previous studies suggested several approaches for integrated end-to-end training to overcome those problems, however they mostly rely on (synthetic or natural) three-way data. We propose a cascaded model based on the non-autoregressive Transformer that enables end-to-end training without the need for an explicit intermediate representation. This new architecture (i) avoids unnecessary early decisions that can cause errors which are then propagated throughout the cascaded models and (ii) utilizes the end-to-end training data directly. We conduct an evaluation on two pivot-based machine translation tasks, namely French→German and German→Czech. Our experimental results show that the proposed architecture yields an improvement of more than 2 BLEU for French→German over the cascaded baseline.

1 Introduction

Many complex natural language applications such as speech translation (Sperber and Paulik, 2020) or pivot translation (Utiyama and Isahara, 2007; De Gispert and Marino, 2006) traditionally rely on cascaded models. The technique of model cascading is commonly used to solve problems that can be divided into a sequence of sub-problems where the solution to the first problem is used as an input to the second and so on. Typically cascaded systems include several consecutive and independently trained models, each of which aims to solve a particular sub-task. For example in a cascaded speech

translation system an automatic speech recognition model receives the audio signal as an input and generates a transcription as an output of the first sub-task. This output could be passed to a system that adds punctuation and capitalization to the sequence, before, as a final step, a machine translation system is applied.

Cascaded models are appealing if there is more training data for each of the sub-tasks than for the full task. Examples for such scenarios include automatic speech translation (AST), image captioning in non-English languages, and non-English machine translation. However, cascaded models are prone to error propagation, meaning that decision errors in the first model are forwarded to and possibly amplified by the second model. Usually, there is also a loss of information when passing information between models since the interface between models traditionally requires each model to output a discrete decision. This means that the deeper knowledge that the model may encode in its representation of the output is reduced to a ‘surface form’ of a particular prediction, which is passed on to the following model. Lastly, in conventional cascaded system there is no possibility to make use of end-to-end training data, meaning that the training data most suited for the task cannot be used.

To tackle these problems, several approaches for integrated end-to-end training of cascaded models have been proposed and applied to different NLP tasks (Bahar et al., 2021; Sperber et al., 2019; Sung et al., 2019). Integrated end-to-end training is usually achieved by merging the consecutive models and fine-tuning the resulting system on the end-to-end training data. Although the idea of this approach is simple, it remains an open challenge how to choose the interface between the models in such a way that they can be trained, e.g. by gradient propagation. Furthermore, most of these approaches rely on synthetic or natural multi-way

training data, i.e. data that does not only provide an (input, output) pair but also the correct label for all sub-tasks involved. For a detailed discussion of the literature, we refer to Section 2. In this work we focus on the task of pivot-based machine translation, i.e. the translation from a source (src) language via a pivot (piv) language to the desired target (trg) language, as an example for a two-stage task that is traditionally solved by model cascading.

We propose a cascaded model based on the non-autoregressive Transformer (NAT) that enables end-to-end training without the need for an explicit intermediate representation, that is inevitable in autoregressive models. This new architecture (i) avoids unnecessary early decisions that can cause errors which are then propagated throughout the cascaded models (ii) utilizes the src→trg, src→piv and piv→trg training data and (iii) communicates the full information from the src→piv model downstream by providing a natural interface between the src→piv and piv→trg models.

2 Related Work

Several approaches were proposed in recent years to address the weaknesses of the traditional cascaded models. Early works investigated the applications of the N-best list decoding both in speech translation and pivot-based translation (Woszczyna et al., 1993; Lavie et al., 1996; Och and Ney, 2004; Utiyama and Isahara, 2007). The N-best list decoding allows to pass multiple intermediate hypotheses and avoid unnecessary early decisions. An efficient alternative to the n -best list is lattices, which replaced the n -best list for the speech translation models (Zhang et al., 2005; Schultz et al., 2004; Matusov et al., 2008). However, the usage of the discrete decisions does not allow to train cascaded model jointly on src→trg data.

Most recent works are focusing instead on the joint or integrated training for sequence-to-sequence cascaded models. Thus, (Cheng et al., 2017) suggested a joint training approach for the pivot-based neural machine translation. In their work, two attention-based RNN models (Bahdanau et al., 2015) are trained jointly with different connection terms in the objective function and the src→trg as a bridging corpus. Another approach is to apply the transfer-learning technique for pivot-based NMT (Kim et al., 2019), meaning that the direct src→trg model is initialized with the respective weights from the pre-trained models, and

fine-tuned on src→trg corpus through the trainable adapter. Pivot-based NMT is typically used in a low-resource src→trg setup, and multilingual NMT systems proved to be successful in this scenario (Johnson et al., 2017; Aharoni et al., 2019; Zhang et al., 2020). To tackle a low-resource NMT problem, (Kim et al., 2019) also explore different ways to extend the back-translation idea (Sennrich et al., 2016a) for src→piv→trg scenarios. However, since this work aims to provide the general framework for the integrated training of cascaded sequence-to-sequence models, we do not aim for comprehensive comparisons with multilingual NMT systems and various data augmentation strategies. We refer to (Kim et al., 2019) for in-depth comparison studies.

In speech translation, the tight model integration for the cascaded models also attracted attention from the community. (Anastasopoulos and Chiang, 2018; Wang et al., 2019; Sperber et al., 2019) discussed either use of attention or hidden state vectors as a connection interface for the tight model integration in cascaded systems. Recently, (Bahar et al., 2021) proposed to use posterior distribution as an input to the encoder of the second model.

3 Background

3.1 Sequence-to-Sequence modeling

The modeling of the sequence-to-sequence problems, namely converting the source sequence f_1^J in one domain to the target sequence e_1^I in another domain, is nowadays usually done using encoder-decoder deep neural networks (Sutskever et al., 2014; Bahdanau et al., 2015; Vaswani et al., 2017). The purpose of the encoder is to map the input sequence f_1^J to a continuous, hidden vector representation h , from which the decoder decodes the target sequence.

In applications such as machine translation, the Transformer (Vaswani et al., 2017), an attention-based sequence-to-sequence model, is considered state of the art (Barrault et al., 2020).

Commonly the probability distribution over the target sequences in sequence-to-sequence models is expressed by a left-to-right factorization:

$$p(e_1^I | f_1^J) = \prod_{i=1}^I p(e_i | e_1^{i-1}, f_1^J). \quad (1)$$

These models are also called *autoregressive*, meaning that each consecutive token in the target se-

quence depends on the left context of the same sequence.

3.2 Non-Autoregressive NMT

In contrast to the autoregressive modelling approach, the non-autoregressive Transformer (Gu et al., 2018) assumes that all tokens in the target sequence are generated independently of each other. This means in particular that there is no need for a search procedure at inference time since target tokens can be generated and optimized in parallel. However, current approaches also need an explicit length model as additional input to the decoder. Gu et al. (2018) utilize the standard Transformer architecture and provide several modifications in order to obtain a non-autoregressive machine translation system.

Recent works proposed to relax the independence constraints during training and use *iterative decoding* for the NAT, meaning that instead of only one decoding pass, the model relies on the multiple passes, and conditional dependence might be used on the consecutive passes to achieve better performance (Ghazvininejad et al., 2019; Gu et al., 2019; Lee et al., 2018; Stern et al., 2019). Such decoding procedure allows shrinking the gap between the performance of the autoregressive and non-autoregressive models.

3.3 Pivot-based Machine Translation

A cascading system p_{s2t} for pivot-based machine translation consists of a $src \rightarrow piv$ model p_{s2p} and a $piv \rightarrow trg$ model p_{p2t} , which typically have a disjoint parameter set. While both models are trained independently, they work in cooperation when producing the translation, i.e., the most likely target sequence \hat{e}_1^I for the given source sequence f_1^J . The pivot sequence z_1^K can be viewed as a latent variable, and the target sequence probability can be expressed by summing over all pivot sequences:

$$\begin{aligned} p_{s2t}(e_1^I | f_1^J) &= \sum_{z_1^K} p_{p2t}(e_1^I | z_1^K, f_1^J) p_{s2p}(z_1^K | f_1^J) \\ &= \sum_{z_1^K} p_{p2t}(e_1^I | z_1^K) p_{s2p}(z_1^K | f_1^J). \end{aligned}$$

Since the sum over all possible pivot hypothesis z_1^K is intractable in practice, instead *two-pass decoding* is used as an approximation to obtain the target

hypotheses:

$$\hat{z}_1^K = \operatorname{argmax}_{K, z_1^K} \prod_{k=1}^K p_{s2p}(z_k | z_1^{k-1}, f_1^J) \quad (2)$$

$$\hat{e}_1^I = \operatorname{argmax}_{I, e_1^I} \prod_{i=1}^I p_{p2t}(e_i | e_1^{i-1}, \hat{z}_1^K). \quad (3)$$

We investigate the stability and potential for improvement of this interface in the Section 6.1.

4 Model Integration

Starting from the conventional cascaded model, as described in Section 3.3, we propose to connect the two consecutive encoder-decoder models through an end-to-end trainable interface. The $src \rightarrow piv$ model consists of both Encoder_{s2p} and Decoder_{s2p} , similarly the $piv \rightarrow trg$ model consists of Encoder_{p2t} and Decoder_{p2t} . We introduce an interface which connects Decoder_{s2p} to the Encoder_{p2t} . The main requirement for this connection interface is to be differentiable to make the gradient propagation possible. In order to fulfill this requirement, we follow the previous work (see more in Section 2) and choose to focus on two possible interfaces:

- **Decoder States Interface:** Pass the final sequence of hidden states vectors of the last $src \rightarrow piv$ Decoder_{s2p} layer as an input to the Encoder_{p2t} . The input embedding layer and positional encoding layer are omitted in the Encoder_{p2t} , and the hidden states vector is then used directly as an input to the next self-attention block (see Figure 1a).
- **Decoder Posteriors Interface:** Pass the probability distribution $p_{s2p}(z_1^K | i, f_1^J)$ of the Decoder_{s2p} . The embedding matrix E from Encoder_{p2t} is used to calculate a ‘soft embedding’

$$\sum_{v \in V} E_v p_{s2p}(z_k = v | f_1^J).$$

Hence, the Decoder_{s2p} and Encoder_{p2t} are connected through the softmax layer, as shown shown in Figure 1b.

Note that the decoder posteriors interface requires the $src \rightarrow piv$ and $piv \rightarrow trg$ model to share a common vocabulary V .

Two autoregressive encoder-decoder models can be connected through these interfaces as shown

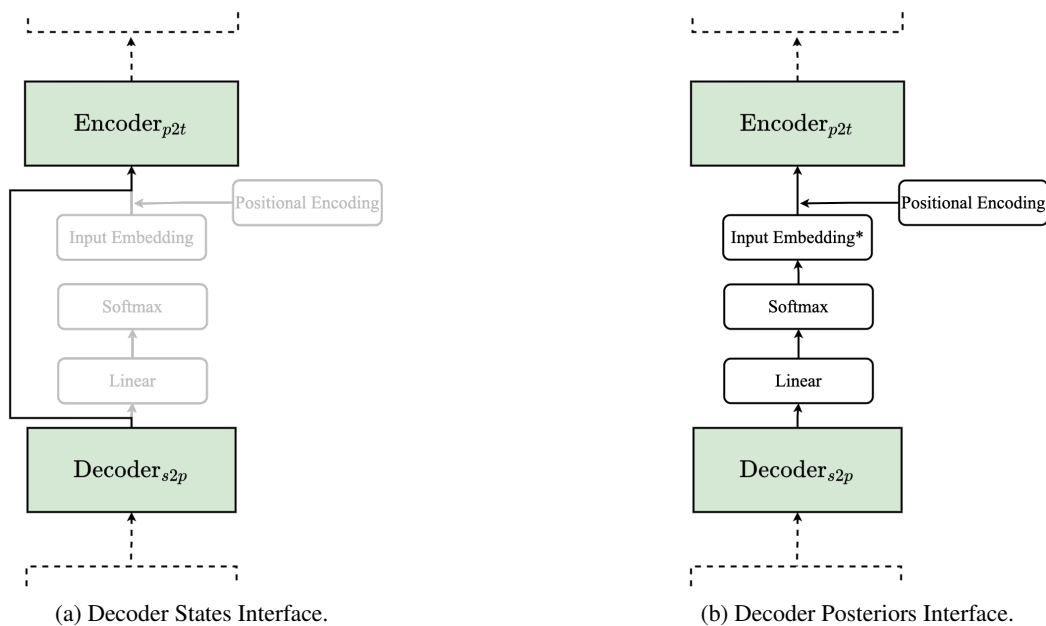


Figure 1: Two proposed connection interfaces between $\text{src} \rightarrow \text{piv}$ and $\text{piv} \rightarrow \text{trg}$ models for integrated training. The blocks in gray represents are omitted layers of the original cascaded Transformer architecture. For simplicity we do not show the Encoder_{s2p} and Decoder_{p2t} .

*Note that the input embedding is now a full fledged matrix multiplication, not a multiplication with a one-hot vector which is equivalent to a column selection.

in Figure 2a. However, at training time the Decoder_{s2p} requires a pivot sequence as an input. If there is no access to the three-way $\text{src} \rightarrow \text{piv} \rightarrow \text{trg}$ data, the pivot sequence has to be obtained by doing a search in training, which is computationally very prohibitive in a real world task, or via forward or backward translation beforehand (synthetic data). The disadvantage of using synthetic data is that the pivot sequences remain static throughout the training, this means that the cascaded $\text{src} \rightarrow \text{piv} \rightarrow \text{trg}$ model is trained on pivot sequences which become less relevant the more training updates the $\text{src} \rightarrow \text{piv}$ models receives. To avoid a sub-optimal, discrete intermediate representation while still benefit from the model integration, we propose to replace $\text{src} \rightarrow \text{piv}$ autoregressive Transformer with a non-autoregressive one as shown in Figure 2b. The usage of NAT allows to replace the pivot sequence with a sequence of unknowns during the training on $\text{src} \rightarrow \text{trg}$ data. Since the decoder states interface do not use the embeddings of the Encoder_{p2t} , similar to other works, the Encoder_{p2t} can be safely omitted in the integrated model (Figure 2c).

Training such a cascaded model can be done with the following steps:

- *Pre-training*:
 - Train $\text{src} \rightarrow \text{piv}$ model on $\text{src} \rightarrow \text{piv}$ cor-

- pora
- Train $\text{piv} \rightarrow \text{trg}$ model on $\text{piv} \rightarrow \text{trg}$ cor-pora

- *Concatenation*: Concatenate the models in the cascade through the interface and initialize respective components with the pre-trained weights.
- *Fine-tuning*: fine-tune the resulting integrated model on the $\text{src} \rightarrow \text{trg}$ data.

This yields a $\text{src} \rightarrow \text{trg}$ architecture in which all parameters are pre-trained and which makes use of all parameters from the pre-trained models, with the exception of one linear layer and an embedding matrix in the decoder states interface. Please note that although we are focusing on pivot-based NMT as our target task, we argue that the proposed integration method can be easily adapted to any Transformer-based cascaded model.

5 Experimental Results

To test and verify the proposed cascaded model, we conduct experiments on French \rightarrow German and German \rightarrow Czech data from the WMT 2019 news translation task¹.

¹<http://www.statmt.org/wmt19/>

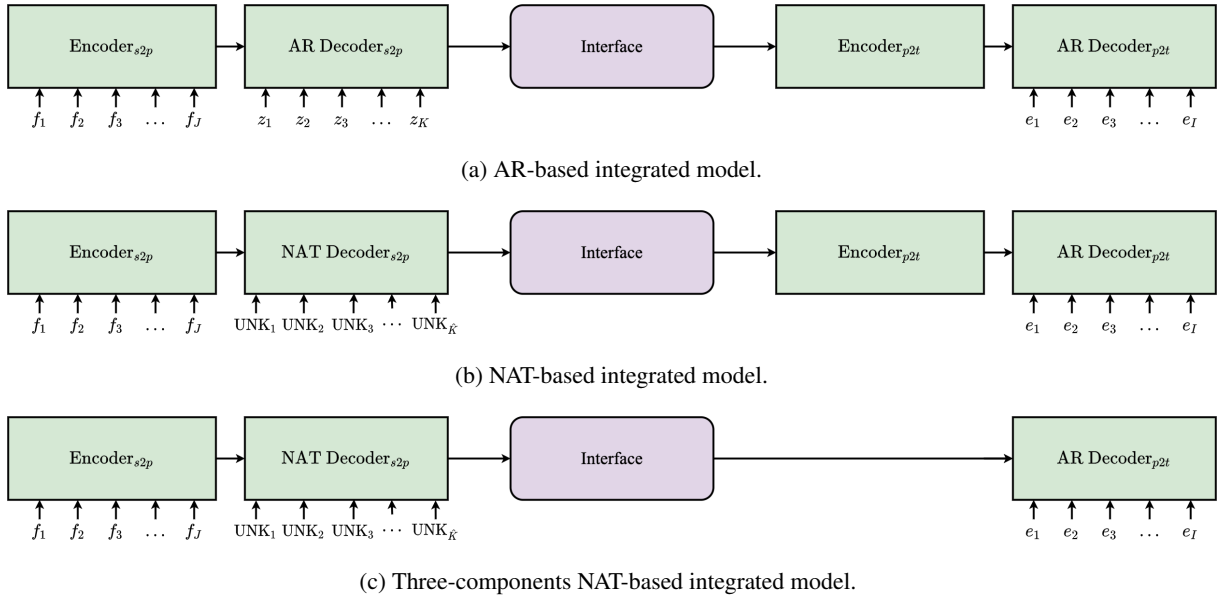


Figure 2: Different variants of the encoder-decoder model integration through the connection interface.

5.1 Data

Training data for French→German includes Europarl corpus version 7 (Koehn, 2005), CommonCrawl² corpus and the newstest2008-2010. The total number of parallel sentences is 2.3M.

The original German→Czech task was constrained to unsupervised translation, but we utilized the available parallel data to relax these constraints. The corpus consists of NewsCommentary version 14 (Tiedemann, 2012) and we extended it by including newssyscomb2009³ and the concatenation of previous years test sets newstest2008-2010 from the news translation task. The total amount of parallel sentences is 230K.

For both tasks we use newstest2011 as the development set and newstest2012 as the test sets. The data statistics, including pre-training data, are collected in Table 1.

		Sentences	Words (target)
direct data	French→German	2.3M	53M
pre-train	French→English	35M	905M
	English→German	9.7M	221M
direct data	German→Czech	230K	4.5M
pre-train	German→English	9.1M	180M
	English→Czech	49M	486M

Table 1: Training data overview.

²<https://commoncrawl.org/>
³<http://www.statmt.org/wmt09/system-combination-task.html>

5.2 Preprocessing

For each parallel corpus, we apply a standard preprocessing procedure: First, we tokenize each corpus using the Moses⁴ tokenizer. Then a true-casing model is trained on all training data and applied to both training and test data. In the final step, we train *byte-pair encoding* (BPE) (Sennrich et al., 2016b) with 32000 merge operations. In order to enable model integration, we train BPE jointly on all available data for the respective language.

5.3 Model and Training

We implement the models described in Section 4 using the *fairseq* (Ott et al., 2019) sequence-to-sequence extendable framework. As non-autoregressive src→piv model, we choose the Conditional Masked Language Model (CMLM) (Ghazvininejad et al., 2019) with 6 layers for both encoder and decoder, and a standard 6 layer ‘base’ Transformer for the piv→trg system (Vaswani et al., 2017). For each interface, the length of the pivot sequence is set to the length of the source sequence by default. More on the length modeling is discussed in the Section 6.4. For the decoder states interface, the last decoder is used for all the experiments.

For model fine-tuning, the Adam optimizer (Kingma and Ba, 2015) with $\beta = \{(0.9, 0.98)\}$ and the learning rate 0.5×10^{-5} is used for all the models. The

⁴<http://www.statmt.org/moses/>

learning rate is reduced during training based on the inverse square root of the update. Additionally, 10,000 and 4,000 warm-up updates have been used for French→German and German→Czech accordingly. The dropout is set to 0.1 for French→German and 0.3 for German→Czech. We set the effective batch size to 65,536 following the fairseq recommendations for the non-autoregressive models. Although CMLM provides the Mask-Predict decoding algorithm (Ghazvininejad et al., 2019), in our work we only use one iteration and obtain probability distribution and hidden states from the fully masked sequence, which means that each token is only conditioned on the source tokens. Results are reported using the *sacreBLEU*⁵ implementation of BLEU (Papineni et al., 2002).

We compare our models against three baselines:

- *direct baseline*: The direct baseline is the Transformer base model, which is trained only on src→trg (direct) parallel data.
- *AR pivot baseline*: A baseline system composed of cascading a src→piv and a piv→trg autoregressive (AR) models. These two models are autoregressive Transformer ‘base’ models with six layers of encoder and decoder, respectively. The individual models are trained on either src→piv or piv→trg data. There is no fine-tuning on the src→trg data, and results are reported based on the inference only.
- *NA pivot baseline*: Similarly to the AR baseline, we provide the results for the non-autoregressive (NA) pivot baseline. The main difference is that the non-autoregressive CMLM model is selected as the src→piv model. We follow standard training procedure for the CMLM as described in (Ghazvininejad et al., 2019), and as for hyperparameters, we rely on the fairseq guidelines⁶. While pre-training, a random mask is applied to the decoder input, meaning that the number of observed and masked tokens varies for each batch. During decoding, we employ five decoding iterations to achieve better performance on the src→piv model. The Transformer base piv→trg model is trained in the

same way as for the AR pivot baseline.

Additionally, we compare our NA integrated model with the AR integrated model (2a) based on the synthetic data generation (Hilmes, 2020). Synthetic data is generated by the forward pass of the src→piv model offline before fine-tuning on the src→trg data, meaning that the pivot hypotheses stay the same during fine-tuning.

We report the best results for the proposed cascaded model with the different interfaces in Table 2. The best checkpoint is selected based on BLEU score of the development set. The results show up to 2.1% BLEU improvements for the decoder states and decoder posteriors interfaces on French→German compare to the pivot baseline. On the other hand, there is a 2.0% BLEU degradation of the performance while using decoder posteriors interface on German→Czech compare to the pivot baseline and up to 2.3% BLEU degradation using decoder states interface. We suppose that such degradation can be based on the training data size since the German→Czech is ten times smaller than French→German. To check on our assumption, we perform additional analysis with the different training data partitions in Section 6.2. Moreover, according to the decoder states interface results, the usage of the additional encoder showed its usefulness compared to the three-components architecture.

6 Analysis

6.1 Error Propagation

Error propagation is a well-known problem of cascaded models. In the following we investigate how significantly errors in one model influence the following models. To this end, we monitor both the individual model performance and the end-to-end cascaded performance by running experiments on a three-way test set that consists of (source, pivot, target) triples. For that purpose, we extract 3000 overlapping sentences from NewsCommentary v14 for WMT French→English and WMT English→German to create a new test set that is disjoint with the training data. We train a 6-layer ‘base’ Transformer for French→English (src→piv) and another for English→German (piv→trg). In order to analyse the impact of disturbances and simulate errors in the French→English system, we generate a weaker hypothesis by:

⁵<https://github.com/mjpost/sacrebleu>

⁶https://github.com/pytorch/fairseq/blob/master/examples/nonautoregressive_translation/scripts.md

		French→German		German→Czech	
		BLEU[%]		BLEU[%]	
		dev	test	dev	test
AR	direct baseline	20.0	20.4	13.5	14.0
	pivot baseline	19.5	20.7	18.8	18.1
NA Int.	Pivot hypothesis (NA pivot baseline)	17.1	18.1	17.3	16.6
	Decoder States w/o Encoder _{p2t}	20.9	21.8	15.5	15.5
		w Encoder _{p2t}	21.5	22.8	16.5
	Decoder Posteriors	21.6	22.7	16.8	17.0
AR Int.	Decoder States [†]	20.6	21.2	16.6	16.8
	Decoder Posteriors [†]	20.5	21.1	17.9	17.1

Table 2: Results for integrated training with different non-autoregressive (NA) interfaces on src→trg data in comparison to autoregressive (AR) baseline model. All pivot/cascaded models are pre-trained on the respective data. We use `newstest{2011, 2012}` as dev and test respectively. Results marked with [†] are taken from (Hilmes, 2020).

- Applying artificial character-level noise: With a probability of p_{noise} each character in the decoded pivot hypothesis is replaced with a random character from the character set of the sentence
- Using a weaker checkpoint than the baseline
- Reducing the beam size to 1 (greedy search)

By applying these procedures, we control the performance of the src→piv model while maintaining a stable performance for the piv→trg model. As is shown in Figure 3, the errors in the src→piv model are actually deflated by the piv→trg system, since a loss of 1.0 BLEU in the src→piv system results in only a drop of around 0.5 BLEU for the cascaded src→trg system.

Similarly, we conduct experiments in the other direction. By improving the quality of the prediction from the src→piv model, we study the potential gain for the src→trg task. For that purpose, we translate each source sentence to a 10-best list of pivot sentences. Using the pivot reference from the three-way test set we can select the single best hypothesis based on the sentence-level BLEU

The sentence with the best BLEU score among ten candidates is then passed to the piv→trg model. This cheating experiment results in an improvement of 6.2% absolute BLEU on the src→piv model, which in turn however only results in 1.4% absolute BLEU improvement on the cascaded src→trg model. We conclude that (i) the piv→trg models weakens both improvements and errors of

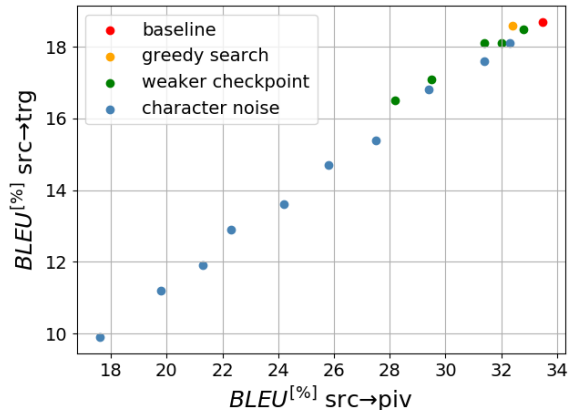


Figure 3: Impact of errors in the src→piv model on the performance of the cascaded src→trg system.

the src→piv model and (ii) the ambiguities in an src→piv 10-best list hold room for an improvement of over 1.0 BLEU.

6.2 Effect of Training Data Size

To investigate how much the NAT-based integrated model quality depends on the training data size, we train our model on randomly sampled 50%, 30%, and 10% selections of the original French→German training corpus. To prevent overfitting on a small corpus, we increase the dropout rate to 0.3 compared to 0.1 on full French→German corpus. The Table 3 shows that when training on 10% of the original data, the discrepancy between the best model performance is around 2.4% BLEU. This setup simulates the

data conditions of German→Czech since the total amount of training sentences in German→Czech corpus is around 10% of the French→German corpus. Based on our experimental results, we suppose that the integrated model needs some minimum amount of parallel src→trg data to achieve the acceptable performance.

data percentage	BLEU ^[%]
100%	21.5
50%	21.0
30%	20.6
10%	19.1

Table 3: French→German dev set results using different training data partitions. The data percentage refers to the relative size of the training corpus comparing to the full French→German training set. All experiments use the decoder states interface for NAT-based integrated training.

6.3 Effect of Model Pre-training

In our experiments for the NAT-based integrated model, we solely rely on the models’ pre-training, which means that instead of random initialization for the NAT-based integrated model components, we utilize the weights from the respective pre-trained models. In this section, we study the importance of model pre-training and its impact on the final model performance. For that purpose, we train the NAT-based integrated model with various initialization options.

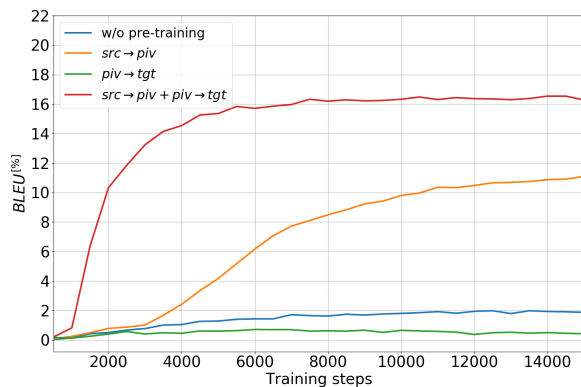


Figure 4: German→Czech dev set results for different parameter pre-training schemes. src→piv indicates that both Encoder_{s2p} and Decoder_{s2p} are pre-trained and all other parameters are randomly initialized. We use a similar notation for the other pre-training schemes. All experiments use the decoder states interface for NAT-based integrated training.

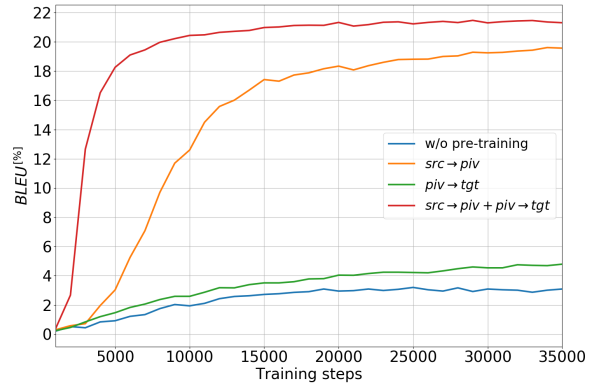


Figure 5: French→German dev set results for different parameter pre-training schemes. All experiments use the decoder states interface for NAT-based integrated training.

Figure 4 and Figure 5 show that initialization of src→piv encoder and decoder is crucial for the final model performance. Without initialization or with pre-training only piv→trg encoder and decoder, it is impossible to train the end-to-end system. We see a similar trend while using the decoder posteriors interface.

6.4 Length Modeling

Length modeling for the non-autoregressive decoder is one of the bottlenecks for our proposed NAT-based integrated model. The pivot sequence length has to be set in advance, and it can not be refined. In most of our experiments, we set the length of the intermediate sequence to be equal to the source sequence length both in training and test time. As a result, we do not fine-tune the length model using the src→trg data. Moreover, the assumption that source length should match the pivot length does not hold for every language pair. In Table 4 we experiment with using different length estimates and report how it affects the end-to-end translation quality.

The results show that better length modeling can lead to more than 2% BLEU improvements. However, for our experiments, we have not tried any sophisticated length prediction methods. We suppose that further exploration will be beneficial for the integrated model performance.

6.5 Decoder Iterations

The iterative refinement of the hypotheses by a non-autoregressive decoder plays an essential role in achieving better performance (Ghazvininejad et al., 2019; Gu et al., 2019). We observe that, the NA

	French→German	German→Czech
length source	BLEU ^[%]	BLEU ^[%]
random	19.2	14.6
source	21.6	16.8
target	18.9	16.5
predicted	21.3	17.2

Table 4: Results for the different pivot length estimates on the dev set. Length source `random` refers to the length choice based on uniform distribution in the interval $[2, 100)$. `predicted` refers to the usage of the CMLM length prediction component for length assignment. `source` and `target` indicate the length choice based on the source sequence or target sequence lengths. All experiments use the decoder posteriors interface for NAT-based integrated training.

baseline with one decoder iteration of the `src→piv` model results in 8.2 BLEU on the French→German development set, while five iterations of the same decoder yield 17.1 BLEU. However, simply increasing the number of iterations during decoding with the integrated model does not lead to similar improvements. Note that the output of the NA decoder is handed to an encoder, which a) more expressive than a softmax layer and b) is trained on the single-iteration output. This mismatch between training and decoding could be the reason why decoder iterations are not beneficial for the integrated model. Additionally, we experimented with decoder iterations during training of the integrated model, but it breaks the gradient propagation. Although our initial experiments with the iterations have been unsuccessful, we think that they can be applied for training using such approaches as Gumbel-Softmax (Jang et al., 2017).

6.6 Knowledge Distillation

Sequence-level knowledge distillation (KD) (Kim and Rush, 2016) proved to be useful for the training of non-autoregressive models (Zhou et al., 2020). Although it improves the `src→piv` model performance, our initial experiments show that KD results in a 0.1-0.3 BLEU degradation on the integrated model.

7 Conclusion

In this work, we propose a novel architecture for the integrated training of cascaded models based on a non-autoregressive Transformer. We train the model on `src→piv`, `piv→trg`, and `src→trg` data overcoming a drawback of conventional cascaded models. Moreover, it provides a natural inter-

face between two Transformer-based models and avoids unnecessary early decisions for intermediate representations. Our experimental results on the task of pivot-based machine translations show that the NAT-based integrated model outperforms the pivot baseline by up to 2.1% BLEU on WMT French→German.

We analyze the integrated model and conclude that the `src→piv` system is crucial for the final translation performance. Further work is required to apply established NAT improvements to this new architecture, such as iterative decoding in the cascaded training and further experiments on knowledge distillation in the `src→piv` pre-training, both of which show significant improvements in standalone systems (Ghazvininejad et al., 2019; Gu et al., 2018, 2019; Zhou et al., 2020). Additionally, more sophisticated techniques for length modeling, such as an external length model or multiple length candidates, can be applied in the future to improve the quality of the pivot hypotheses.

Even though we test our cascaded architecture on the task for pivot-based machine translation, we can use the architecture in any application, where a combination of sequential models is beneficial.

Acknowledgments

Authors affiliated with RWTH Aachen University have partially received funding from the European Research Council (ERC) (under the European Union’s Horizon 2020 research and innovation programme, grant agreement No 694537, project “SEQCLAS”) and eBay Inc. The work reflects only the authors’ views, and none of the funding agencies is responsible for any use that may be made of the information it contains.

We would like to thank Yunsu Kim and Benedikt Hilmes for providing the autoregressive integrated model results.

References

- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Antonios Anastasopoulos and David Chiang. 2018. [Tied multitask learning for neural speech translation](#). In *Proceedings of the 2018 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 82–91, New Orleans, Louisiana. Association for Computational Linguistics.
- Parnia Bahar, Tobias Bieschke, Ralf Schlueter, and Hermann Ney. 2021. Tight integrated end-to-end training for cascaded speech translation. In *IEEE Spoken Language Technology Workshop*, Shenzhen, China. To appear.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *ICRL 2015*.
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. Findings of the 2020 conference on machine translation (wmt20). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online. Association for Computational Linguistics.
- Yong Cheng, Qian Yang, Yang Liu, Maosong Sun, and Wei Xu. 2017. Joint training for pivot-based neural machine translation. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3974–3980, Melbourne, Australia.
- Adrià De Gispert and Jose B Marino. 2006. Catalan-english statistical machine translation without parallel corpus: bridging through spanish. In *Proc. of 5th International Conference on Language Resources and Evaluation (LREC)*, pages 65–68.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. Mask-predict: Parallel decoding of conditional masked language models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6112–6121, Hong Kong, China. Association for Computational Linguistics.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O. K. Li, and Richard Socher. 2018. Non-autoregressive neural machine translation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. Levenshtein transformer. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 11181–11191. Curran Associates, Inc.
- Benedikt Hilmes. 2020. Investigation on the model architecture for pivot-based neural machine translation. Bachelor’s thesis, Department of Computer Science, RWTH Aachen University, August.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- 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.
- Yunsu Kim, Petre Petrov, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2019. Pivot-based transfer learning for neural machine translation between non-english languages. In *Conference on Empirical Methods in Natural Language Processing*, Hong Kong, China.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.
- Alon Lavie, Donna Gates, Marsal Gavaldà, Laura Mayfield, Alex Waibel, and Lori Levin. 1996. Multilingual translation of spontaneously spoken language in a limited domain. In *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, Brussels, Belgium. Association for Computational Linguistics.
- Evgeny Matusov, Björn Hoffmeister, and Hermann Ney. 2008. Asr word lattice translation with exhaustive reordering is possible. In *Interspeech*, pages 2342–2345, Brisbane, Australia.
- Franz Josef Och and Hermann Ney. 2004. The alignment template approach to statistical machine translation. *Computational Linguistics*, 30(4):417–449.

- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Tanja Schultz, S. Jou, S. Vogel, and S. Saleem. 2004. Using word lattice information for a tighter coupling in speech translation systems. In *INTERSPEECH*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [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.
- Matthias Sperber, Graham Neubig, Jan Niehues, and Alex Waibel. 2019. [Attention-passing models for robust and data-efficient end-to-end speech translation](#). *Transactions of the Association for Computational Linguistics*, 7:313–325.
- Matthias Sperber and Matthias Paulik. 2020. [Speech translation and the end-to-end promise: Taking stock of where we are](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7409–7421, Online. Association for Computational Linguistics.
- Mitchell Stern, William Chan, Jamie Kiros, and Jakob Uszkoreit. 2019. [Insertion transformer: Flexible sequence generation via insertion operations](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5976–5985. PMLR.
- T. Sung, J. Liu, H. Lee, and L. Lee. 2019. [Towards end-to-end speech-to-text translation with two-pass decoding](#). In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7175–7179.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- Masao Utiyama and Hitoshi Isahara. 2007. [A comparison of pivot methods for phrase-based statistical machine translation](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 484–491, Rochester, New York. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Dingmin Wang, Meng Fang, Yan Song, and Juntao Li. 2019. [Bridging the gap: Improve part-of-speech tagging for Chinese social media texts with foreign words](#). In *Proceedings of the 5th Workshop on Semantic Deep Learning (SemDeep-5)*, pages 12–20, Macau, China. Association for Computational Linguistics.
- M. Woszczyna, N. Coccaro, A. Eisele, A. Lavie, A. McNair, T. Polzin, I. Rogina, C. P. Rose, T. Sloboda, M. Tomita, J. Tsutsumi, N. Aoki-Waibel, A. Waibel, and W. Ward. 1993. [Recent advances in janus: A speech translation system](#). In *Proceedings of the Workshop on Human Language Technology, HLT ’93*, page 211–216, USA. Association for Computational Linguistics.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich. 2020. [Improving massively multilingual neural machine translation and zero-shot translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1628–1639, Online. Association for Computational Linguistics.
- R. Zhang, Gen ichiro Kikui, H. Yamamoto, and W. Lo. 2005. A decoding algorithm for word lattice translation in speech translation. In *IWSLT*, pages 23–29.
- Chunting Zhou, Jiatao Gu, and Graham Neubig. 2020. [Understanding knowledge distillation in non-autoregressive machine translation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.