

Non-Parametric Adaptation for Neural Machine Translation

Ankur Bapna

Google AI

ankurbpn@google.com

Orhan Firat

Google AI

orhanf@google.com

Abstract

Neural Networks trained with gradient descent are known to be susceptible to catastrophic forgetting caused by parameter shift during the training process. In the context of Neural Machine Translation (NMT) this results in poor performance on heterogeneous datasets and on sub-tasks like rare phrase translation. On the other hand, non-parametric approaches are immune to forgetting, perfectly complementing the generalization ability of NMT. However, attempts to combine non-parametric or retrieval based approaches with NMT have only been successful on narrow domains, possibly due to over-reliance on sentence level retrieval. We propose a novel n-gram level retrieval approach that relies on local phrase level similarities, allowing us to retrieve neighbors that are useful for translation even when overall sentence similarity is low. We complement this with an expressive neural network, allowing our model to extract information from the noisy retrieved context. We evaluate our semi-parametric NMT approach on a heterogeneous dataset composed of WMT, IWSLT, JRC-Acquis and OpenSubtitles, and demonstrate gains on all 4 evaluation sets. The semi-parametric nature of our approach opens the door for non-parametric domain adaptation, demonstrating strong inference-time adaptation performance on new domains without the need for any parameter updates.

1 Introduction

Over the last few years, neural sequence to sequence models (Sutskever et al., 2014; Bahdanau et al., 2015; Cho et al., 2014) have revolutionized the field of machine translation by significantly improving translation quality over their phrase based counterparts (Sennrich et al., 2015; Wu et al., 2016; Zhou et al., 2016). With more gains arising from continued research on new neural network architectures and accompanying train-

ing techniques (Vaswani et al., 2017; Gehring et al., 2017; Chen et al., 2018), NMT researchers, both in industry and academia, have doubled down on their ability to train high capacity models on large corpora with gradient based optimization.

However, despite huge improvements in overall translation quality NMT has shown some glaring weaknesses, including idiom processing, and rare word or phrase translation (Koehn and Knowles, 2017; Isabelle et al., 2017; Lee et al., 2018) - tasks that should be easy if the model could retain learned information from individual training examples. NMT has also been shown to perform poorly when dealing with multi-domain data (Farajian et al., 2017a). This ‘catastrophic forgetting’ problem has been well-studied in traditional neural network literature, caused by parameter shift during the training process (McCloskey and Cohen, 1989; Santoro et al., 2016). Non-parametric methods, on the other hand, are resistant to forgetting but are prone to over-fitting due to their reliance on individual training examples. We focus on a non-parametric extension to NMT, hoping to combine the generalization ability of neural networks with the eidetic memory of non-parametric methods. Given a translation query, we rely on an external retrieval mechanism to find similar source-target instances in the training corpus, which are then utilized by the model.

There has been some work on semi-parametric NMT (Gu et al., 2017; Zhang et al., 2018b; Cao and Xiong, 2018), but its effectiveness has been confined to narrow domain datasets. Existing approaches have relied on sentence level similarity metrics for retrieval, which works well for domains with high train-test overlap, but fails to retrieve useful candidates for broad domains. Even if we could find training instances with overlapping phrases it’s likely that the information in most retrieved source-target pairs is noise for the pur-

pose of translating the current query.

To retrieve useful candidates when sentence similarity is low, we use n-gram retrieval instead of sentence retrieval. This results in neighbors which have high local overlap with the source sentence, even if they are significantly different in terms of overall sentence similarity. This is intuitively similar to utilizing information from a phrase table (Koehn et al., 2003) within NMT (Dahlmann et al., 2017), without losing the global context lost when constructing the phrase table. We also propose another simple extension using dense vectors for n-gram retrieval which allows us to exploit similarities beyond lexical overlap.

To effectively extract the signal from the noisy retrieved neighbors, we develop an extension of the approach proposed in (Cao and Xiong, 2018). While (Cao and Xiong, 2018) encode the retrieved targets without any context, we incorporate information from the current and retrieved sources while encoding the retrieved target, in order to distinguish useful information from noise.

We evaluate our semi-parametric NMT approach on two tasks.

- We evaluate our approach on a multi-domain English-French corpus constructed from narrow domain datasets like JRC-Acquis (Steinberger et al., 2006; Tiedemann) and OpenSubtitles (Tiedemann, 2009)¹, and the standard IWSLT and WMT bilingual corpora, as described in Sections 3 and 4. Our results, for the first time, indicate that semi-parametric NMT can be beneficial beyond narrow domain tasks, demonstrating gains of around 0.5 BLEU on WMT, and huge gains ranging from 2-10 BLEU points on IWSLT, JRC-Acquis and OpenSubtitles, when compared to a strong sequence to sequence baseline.
- The semi-parametric nature of our model enables non-parametric inference-time adaptation to new datasets, without the need for any parameter updates. When trained on WMT and evaluated on the other datasets, our model out-performs fine-tuning based adaptation (Luong and Manning, 2015) on JRC-Acquis and OpenSubtitles, and significantly improves performance over the non-adapted model on IWSLT.

¹<http://www.opensubtitles.org/>

2 Semi-parametric NMT

Standard approaches for Neural Machine Translation rely on seq2seq architectures (Sutskever et al., 2014; Bahdanau et al., 2015), where given a source sequence $X = \{x_1, x_2, \dots, x_{T_x}\}$ and a target sequence $Y = \{y_1, y_2, \dots, y_{T_y}\}$, the goal is to model the probability distribution, $p(y_t | X, y_1, \dots, y_{t-1})$.

Semi-parametric NMT (Dahlmann et al., 2017; Gu et al., 2017) approaches this learning problem with a different formulation, by modeling $p(y_t | X, y_1, \dots, y_{t-1}, \Phi_X)$ instead, where $\Phi_X = \{(X^1, Y^1) \dots (X^N, Y^N)\}$ is the set of sentence pairs where the source sentence is a neighbor of X , retrieved from the training corpus using some similarity metric. This relies on a two step approach - the retrieval stage finds training instances, (X^i, Y^i) , similar to the source sentence X , and the translation stage generates the target sequence Y given X and Φ_X . We follow this setup, proposing improvements to both stages in order to enhance the applicability of semi-parametric NMT to more general translation tasks.

2.1 Retrieval Approaches

Existing approaches have proposed using off the shelf search engines for the retrieval stage. However, our objective differs from traditional information retrieval, since the goal of retrieval in semi-parametric NMT is to find neighbors which might improve translation performance, which might not correlate with maximizing sentence similarity.

Our baseline strategy relies on a sentence level similarity score, similar to those used for standard information retrieval tasks (Robertson, 2004). We compare this against finer-grained n-gram retrieval using the same similarity metric. We also propose a dense vector based n-gram retrieval strategy, using representations extracted from a pre-trained NMT model.

2.1.1 IDF Based Sentence Retrieval

Our baseline approach relies on a simple inverse document frequency (IDF) based similarity score. We define the IDF score of any token, t , as $f_t = \log(\frac{\|C\|}{n_t})$, where $\|C\|$ is the number of sentence pairs in training corpus and n_t is the number of sentences t occurs in. Let any two sentence pairs in the corpus be (X^i, Y^i) and (X^j, Y^j) . Then we define the similarity between (X^i, Y^i) and

(X^j, Y^j) by,

$$\text{sim}(X^i, X^j) = 2 \times \sum_{t \in (X^i \cap X^j)} f_t - \sum_{t \in (X^i \cup X^j)} f_t \quad (1)$$

For every sentence in the training, dev and test corpora, we find the N most similar training sentence pairs and provide them as context to NMT.

2.1.2 IDF Based N-Gram Retrieval

Motivated by phrase based SMT, we retrieve neighbors which have high local, sub-sentence level overlap with the source sentence. We adapt our approach to retrieve n-grams instead of sentences. We note that the similarity metric defined above for sentences is equally applicable for n-gram retrieval.

Let $X = (t^1, \dots, t^T)$ be a sentence. Then the set of all possible n-grams of X , for a given n , can be defined as $S_X^n = \{(t_i, \dots, t_{i+n}) \mid 1 \leq i \leq T\}$ (also including padding at the end). To reduce the number of n-grams used to represent every sentence, we define the reduced set of n-grams for X to be $\hat{S}_X^n = \{(t_i, \dots, t_{i+n}) \mid 1 \leq i \leq T, i \bmod \frac{n}{2} = 1\}$.

We represent every sentence by their reduced n-gram set. For every n-gram in \hat{S}_X^n , we find the closest n-gram in the training set using the IDF similarity defined above. For each retrieved n-gram we find the corresponding sentence (In case an n-gram is present in multiple sentences, we choose one randomly). The set of neighbors of X is then the set of all sentences in the training corpus that contain an n-gram that maximizes the n-gram similarity with any n-gram in \hat{S}_X^n .

To capture phrases of different lengths we use multiple n-gram widths, n . In case a sentence has already been added to the retrieved set, we find the next most similar sentence to avoid having duplicates. The number of neighbors retrieved for each source sentence is proportional to its length.

2.1.3 Dense Vector Based N-Gram Retrieval

We also extend our n-gram retrieval strategy with dense vector based n-gram representations. The objective behind using a dense vector based approach is to incorporate information relevant to the translation task in the retrieval stage. We use a pre-trained Transformer Base (Vaswani et al., 2017) encoder trained on WMT to generate sub-word level dense representations for the sentence. The representation for each n-gram is now defined to be the mean of the representations of all its constituent sub-words. We use the L_2 distance of

n-gram representations as the retrieval criterion. Note that we use a sub-word level decomposition of sentences for dense retrieval, as compared to word-level for IDF based retrieval (i.e., n-grams are composed of sub-words instead of words).

Following the approach described for IDF based n-gram retrieval, we use multiple values of n , and remove duplicate neighbors while creating the retrieved set.

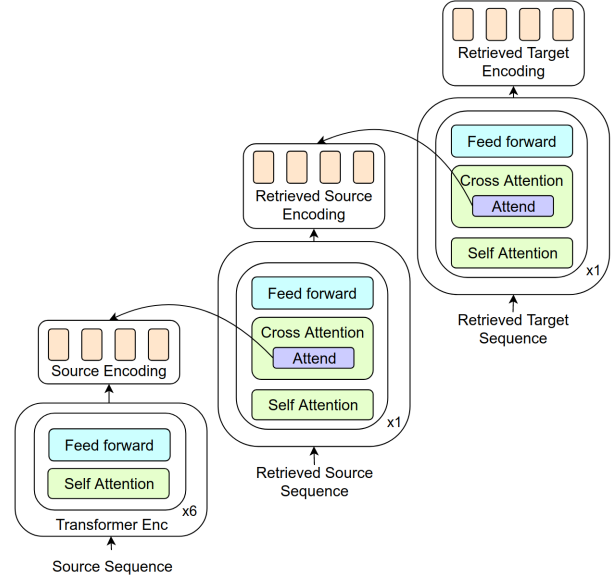


Figure 1: Architecture of the Conditional Source Target Memory. The retrieved targets, Y^i , are encoded in a transformer encoder, incorporating the attention context from the retrieved sources, X^i . In turn, the retrieved sources, X^i , are encoded while incorporating context from the current translation source, X .

2.2 NMT with Context Retrieval

To incorporate the retrieved neighbors, Φ_X , within the NMT model, we first encode them using Transformer layers, as described in subsection 2.2.1. This encoded memory is then used within the decoder via an attention mechanism, as described in subsection 2.2.2.

2.2.1 Conditional Source Target Memory

We now describe how each retrieved translation pair, (X^i, Y^i) , is encoded. This architecture is illustrated in Figure 1.

- We first encode the retrieved source, X^i , in a Transformer layer. Apart from self-attention, we incorporate information from the encoder representation of the current source, X , using decoder style cross-attention.

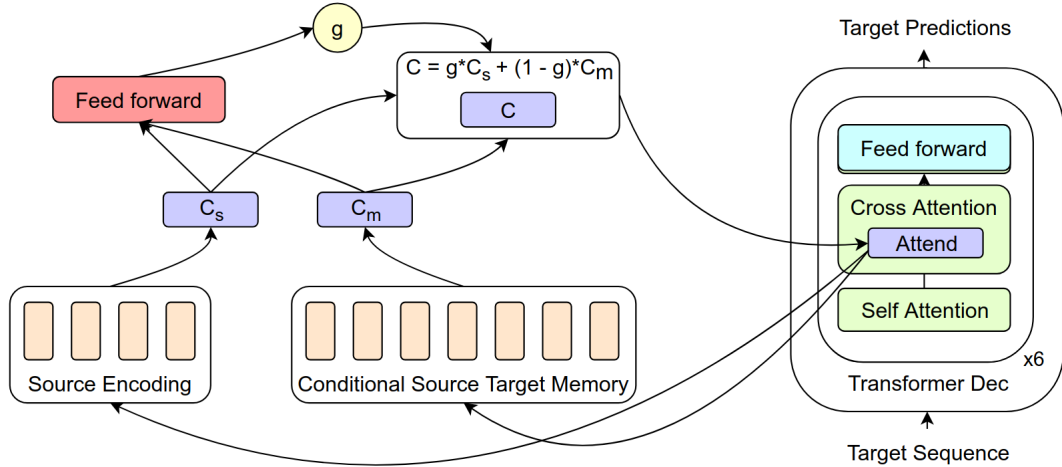


Figure 2: Architecture of the gated attention mechanism used in the multi-source transformer decoder.

- The retrieved target, Y^i , is encoded in a similar manner, attending the encoded representation of X^i generated in the previous step.

The encoded representations for all targets, $\{Y^i, 1 \leq i \leq N\}$, are then concatenated along the time axis to form the Conditional Source Target Memory (CSTM).

2.2.2 Gated Multi-Source Attention

We use gated multi-source attention to combine the context from the source encoder representations and the CSTM. This is similar to the gated attention employed by (Cao and Xiong, 2018). We use a Transformer based decoder that attends to both, the encoder outputs and the CSTM, in every cross-attention layer. The rest of the decoder architecture remains unchanged.

Let the context vectors obtained by applying multi-head attention to the source and memory, with query q_t be c_t^s and c_t^m respectively. Then the gated context vector, c_t^t , is given by,

$$g_t = \sigma(W_{gs}c_t^s + W_{gm}c_t^m) \quad (2)$$

$$c_t = g_t * c_t^s + (1 - g_t) * c_t^m \quad (3)$$

where g_t is the scalar gating variable at time-step t , and W_{gs} and W_{gm} are learned parameters. These steps are illustrated in Figure 2.

3 Experiments

3.1 Data and Evaluation

We compare the performance of a standard Transformer Base model and our semi-parametric NMT approach on an English-French translation task.

We create a new heterogeneous dataset, constructed from a combination of the WMT training set (36M pairs), the IWSLT bilingual corpus (237k pairs), JRC-Acquis (797k pairs)² and OpenSubtitles (33M pairs)³. For WMT, we use newstest 13 for validation and newstest 14 for test. For IWSLT, we use a combination of the test corpora from 2012-14 for validation and test 2015 for eval. For OpenSubtitles and JRC-Acquis, we create our own splits for validation and test, since no benchmark split is publicly available. After deduping, the JRC-Acquis test and validation set contain 6574 and 5121 sentence pairs respectively. The OpenSubtitles test and validation sets contain 3975 and 3488 pairs. For multi-domain training, the validation set is a concatenation of the four individual validation sets.

All datasets are tokenized with the Moses tokenizer (Koehn et al., 2007) and mixed without any sampling. We use a shared vocabulary SentencePiece Model (Kudo and Richardson, 2018) for sub-word tokenization, with a vocabulary size of 32000 tokens. We train each model for 1M steps, and choose the best checkpoint from the last 5 checkpoints based on validation performance. BLEU scores are computed with tokenized true-cased output and references with *multi-bleu.perl* from Moses.

For IDF based sentence retrieval, for each sentence in the training, dev and test corpus, we use $N = 10$ neighbors per example during both, training and evaluation. For the N-Gram level retrieval strategies, we used $N = 10$ neighbors dur-

²From <http://opus.nlpl.eu/JRC-Acquis.php>

³From <http://opus.nlpl.eu/OpenSubtitles.php>

Model	Data	newstest 14	IWSLT 2015	OpenSub	JRC-Acquis
TransformerBase	Multi Domain (MD)	41.92	43.17	26.67	56.19
+ CSTM	MD + IDF Sentence	40.89	42.35	28.25	65.38
+ CSTM	MD + IDF N-Gram	41.92	45.09	28.74	66.39
+ CSTM	MD + Dense N-Gram	42.41	45.02	29.06	66.92

Table 1: Comparison of test translation quality (BLEU) with different retrieval strategies. Multi-domain is a concatenation of all 4 datasets. IDF Sentence, IDF-N-Gram and Dense N-Gram correspond to multi-domain datasets constructed with the different retrieval strategies.

source	‘The top copy of the passenger waybill <i>shall be kept on the</i> bus or coach throughout the journey to which it refers .’
neighbor source	‘The top copy of the journey form <i>shall be kept on the</i> vehicle during the whole of the journey to which it refers .’
baseline translation	‘La copie supérieure de la lettre de transport de voyageurs doit être conservée dans l’ autobus ou l’ autocar tout au long du voyage auquel elle se rapporte .’
neighbor target	‘L’ original de la feuille de route doit se trouver à bord du véhicule <i>pendant toute la durée du voyage pour lequel elle a été établie .’</i>
translation	‘L’ original de la feuille de route doit se trouver à bord de l’ autobus ou de l’ autocar <i>pendant toute la durée du voyage pour lequel elle a été établie .’</i>
reference	‘L’ original de la feuille de route doit se trouver à bord de l’ autobus ou de l’ autocar <i>pendant toute la durée du voyage pour lequel elle a été établie .’</i>

Table 2: A comparison of model outputs on a sample from the JRC-Acquis dataset. This model was trained using IDF based sentence level retrieval with Conditional Source Target Memory. The different colors and text formatting (underlined, italic, bold) represent different overlapping phrases within the model output, the retrieved target and the reference translation.

ing training, and neighbors corresponding to all n-grams during decoding. This was meant to limit memory requirements and enable the model to fit on P100s during training. We used n-gram width, $n = \{6, 10, 18\}$, for both IDF and dense vector based n-gram retrieval approaches. For scalability reasons, we restricted the retrieval set to the in-domain training corpus, i.e. neighbors for all train, dev and test sentences in the JRC-Acquis corpus were retrieved from the JRC-Acquis training split, and similarly for the other datasets.

3.2 Hyper-parameters and Optimization

For our baseline model we use the standard Transformer Base model (Vaswani et al., 2017). For the semi-parametric model, all our hyper-parameters for attention (8 attention heads), model dimensions (512) and hidden dimensions (2048), including those used in the CSTM memory are equivalent to Transformer Base.

The Transformer baselines are trained on 16 GPUs, with the learning rate, warm-up schedule and batching scheme described in (Vaswani et al.,

2017). The semi-parametric models were trained on 32 GPUs with each replica split over 2 GPUs, one to train the translation model and the other for computing the CSTM. We used a conservative learning rate schedule (3, 40K) (Chen et al., 2018) to train the semi-parametric models.

We apply a dropout rate (Srivastava et al., 2014) of 0.1 to all inputs, residuals, attentions and ReLU connections in both models. We use Adam (Kingma and Ba, 2014) to train all models, and apply label smoothing with an uncertainty of 0.1 (Szegedy et al., 2015). In addition to the transformer layers, layer normalization (Ba et al., 2016) was applied to the output of the CSTM. All models are implemented in Tensorflow-Lingvo (Shen et al., 2019).

4 Results

We compare the test performance of a multi-domain Transformer Base and our semi-parametric model using dense vector based n-gram retrieval and CSTM in Table 1. Apart from significantly improving performance by

source	‘Consciousness also is what makes life worth living .’
neighbor source	‘So in the last 10 years and the hope for the future , we ’ve seen the beginnings of a science of positive psychology , a science of what makes life worth living .’
baseline translation	‘La conscience est aussi ce <i>qui rend la vie valable</i> .’
neighbor target	‘Donc , depuis 10 ans , et , espérons-le , à l’ avenir nous assistons à l’ émergence d’ une science de la psychologie positive : une science qui fait en sorte que la vie vaille la peine d’ être vécue .’
translation	‘La conscience est aussi ce qui fait que la vie vaut la peine d’ être vécue .’
reference	‘La conscience est aussi ce qui fait que la vie vaut la peine d’ être vécue .’

Table 3: A comparison of model outputs on a sample from IWSLT. This model was trained using IDF based n-gram retrieval with Conditional Source Target Memory. N-Gram level retrieval results in finding neighbors with high local overlap, even when the rest of the sentences are dissimilar.

source	‘I was expecting to see gnashing of teeth and a fight breaking out at the gate .’
neighbor source	‘One could almost hear the collective gnashing of teeth in the US , especially in the Congress .’
baseline translation	‘J’ espérais voir <i>des dents brûlantes</i> et une bataille éclater à la porte .’
neighbor target	‘On a presque entendre les dents grincer aux États-Unis , surtout au Congrès .’
translation	‘Je m’ attendais à voir des grincements de dents et une bagarre éclater à la porte .’
reference	‘Je m’ attendais à voir des grincements de dents et une bagarre éclater à la porte .’

Table 4: A comparison of model outputs on a sample from WMT. This model was trained using IDF based n-gram retrieval with Conditional Source Target Memory. N-Gram level retrieval results in finding neighbors with high local overlap, even when the rest of the sentences are dissimilar.

more than 10 BLEU points on JRC-Acquis, 2-3 BLEU on OpenSubtitles and IWSLT, we notice a moderate gain of 0.5 BLEU points on WMT 14.

4.1 Comparison of retrieval strategies

We compare the performance of all 3 retrieval strategies in Table 1. The semi-parametric model with sentence level retrieval out-performs the seq2seq model by a huge margin on JRC-Acquis and OpenSubtitles. A sample from the JRC-Acquis dataset where the semi-parametric approach improves significantly over the neural approach is included in Table 2. We notice that there is a lot of overlap between the source sentence and the retrieved source, resulting in the semi-parametric model copying large chunks from the retrieved target. However, its performance is noticeably worse on WMT and IWSLT. Based on a manual inspection of the retrieved candidates, we attribute these losses to retrieval failures. For

broad domain datasets like WMT and IWSLT sentence retrieval fails to find good candidates.

Switching to n-gram level retrieval brings the WMT performance close to the seq2seq approach, and IWSLT performance to 2 BLEU points above the baseline model. Representative examples from IWSLT and WMT where n-gram retrieval improves over sentence level retrieval can be seen in Tables 3 and 4. Despite the majority of the retrieved neighbor having nothing in common with the source sentence, n-gram retrieval is able to find neighbors that contain local overlaps.

Using dense n-gram retrieval allows us to move beyond lexical overlap and retrieve semantically similar n-grams even when the actual tokens are different. As a result, dense n-gram retrieval improves performance over all our models on all 4 datasets. An illustrative example from WMT is included in Table 5.

source	‘The artist died last Sunday at the age of 71 .’
neighbor source	‘A former minister George Thomson passed away last week at the age of 87 .’
baseline translation	‘L’ artiste est mort dimanche dernier à l’ âge de 71 ans .’
neighbor target	‘George Thomson , ancien ministre , est décédé la semaine dernière à l’ âge de 87 ans .’
translation	‘L’ artiste est décédé dimanche dernier à l’ âge de 71 ans .’
reference	‘L’ artiste est décédé dimanche dernier , à l’ âge de 71 ans .’

Table 5: A comparison of model outputs on a sample from WMT. This model was trained using dense vector based n-gram retrieval with Conditional Source Target Memory. Dense vector based n-gram retrieval allows us to find semantically similar phrases, even when the lexical context is dissimilar.

Model	newstest 14	IWSLT 2015	OpenSub	JRC-Acquis
No Memory	41.92	43.17	26.67	56.19
TM	41.64	44.32	27.38	64.25
CTM	41.87	44.76	27.74	65.18
CSTM	42.41	45.02	29.06	66.92

Table 6: Comparison of test translation quality (BLEU) with different memory architectures. All models are trained on the Dense N-Gram Multi-Domain dataset. CSTM corresponds to the proposed Conditional Source Target Memory. CTM corresponds to Conditional Target Memory, where we ignore the retrieved sources while encoding the retrieved targets, and directly attend the encoding of the current source, X . TM corresponds to encoding the retrieved targets without any context.

4.2 Memory Ablation Experiments

We report the performance of the various memory ablations in Table 6. We first remove the retrieved sources, X^i , from the CSTM, resulting in an architecture where the encoding of a retrieved target, Y^i , only incorporates information from the source X , represented by the row CTM in the table. This results in a clear drop in performance on all datasets. We ablate further by removing the attention to the original source X , resulting in a slightly smaller drop in performance (represented by TM). These experiments indicate that incorporating context from the sources significantly contributes to performance, by allowing the model to distinguish between relevant context and noise.

5 Non-Parametric Adaptation

Using a semi-parametric formulation for MT opens up the possibility of non-parametric adaptation. The biggest advantage of this approach is the possibility of training a single massively customizable model which can be adapted to any new dataset or document at inference time, by just updating the retrieval dataset.

We evaluate our model’s performance on non-parametric adaptation and compare it against a fully fine-tuned model. In this setting, we train a

baseline model and a dense n-gram based semi-parametric model on the WMT training corpus. We only retrieve and train on examples from the WMT corpus during training. We use the same hyper-parameters and training approaches used for the multi-domain experiments, as in Section 3.

The baseline model is then fine-tuned independently on JRC-Acquis, OpenSubtitles and IWSLT. The semi-parametric model is adapted non-parametrically to these three datasets, without any parameter updates. Adaptation is achieved via the retrieval mechanism - while evaluating, we retrieve similar examples from their respective training datasets. To quantify headroom, we also fine-tune our semi-parametric model on each of these datasets.

The results for non-parametric adaptation experiments are documented in Table 7. We notice that the non-parametric adaptation strategy significantly out-performs the base model on all 4 datasets. More importantly, we find that our approach is capable of adapting to both, JRC-Acquis and OpenSubtitles, via just the retrieval apparatus, and out-performs the fully fine-tuned model indicating that non-parametric adaptation might be a reasonable approach when adapting to a lot of narrow domains or documents.

Adaptation Strategy	newstest 14	IWSLT 2015	OpenSub	JRC-Acquis
Base	41.16	39.75	22.92	53.1
Fine-tuning	-	42.87	26.55	62.99
Non-Parametric (NP)	41.57	40.95	27.09	64.93
NP + Fine-tuning	-	43.82	29.12	66.72

Table 7: Comparison of test translation quality (BLEU) with different adaptation strategies. The base model (Transformer Base) is trained on the WMT dataset. Fine-tuning corresponds to fine-tuning based adaptation, where we initialize the domain-specific model from the WMT pre-trained Base model, and fine-tune it on the in-domain dataset for a few epochs. Non-parametric corresponds to our semi-parametric NMT model, adapted to in-domain data during inference by retrieving neighbors from the in-domain training corpus.

In-domain fine-tuning on top of non-parametric adaptation further improves by 2 BLEU points on all datasets, increasing the gap even further with the seq2seq adapted models.

6 Related Work

Tools incorporating information from individual translation pairs, or translation memories (Lagoudaki; Reinke, 2013), have been widely utilized by human translators in the industry. There have been a few efforts attempting to combine non-parametric methods with NMT (Gu et al., 2017; Zhang et al., 2018b; Cao and Xiong, 2018), but the key difference of our approach is the introduction of local, sub-sentence level similarity in the retrieval process, via n-gram level retrieval. Combined with our architectural improvements, motivated by the target encoder and gated attention from (Cao and Xiong, 2018) and the extended transformer model from (Zhang et al., 2018a), our semi-parametric NMT model is able to outperform purely neural models in broad multi-domain settings.

Some works have proposed using phrase tables or the outputs of Phrase based MT within NMT (Dahmann et al., 2017; Zhang et al., 2017; Zhou et al., 2017). While this reduces the noise present within the retrieved translation pairs, it requires training and maintaining a separate SMT system which might introduce errors of its own.

Another class of methods requires fine-tuning the entire NMT model to every instance at inference time, using retrieved examples (Farajian et al., 2017b; Wuebker et al., 2015), but these approaches require running expensive gradient descent steps before every translation.

Beyond NMT, there have been a few other attempts to incorporate non-parametric approaches into neural generative models (Guu et al., 2018;

Hayati et al., 2018; Weston et al., 2018). This strong trend towards combining neural generative models with non-parametric methods is an attempt to counter the weaknesses of neural networks, especially their failure to remember information from individual training instances and the diversity problem of seq2seq models (Vijayakumar et al., 2016; Jiang and de Rijke, 2018).

While our approach relies purely on retrieval from the training corpus, there has been quite a lot of work, especially on Question Answering, that attempts to find additional signals to perform the supervised task in the presence of external knowledge sources (Chen et al., 2017; Wang et al., 2018). Retrieving information from unsupervised corpora by utilizing multilingual representations (Guo et al., 2018) might be another interesting extension of this work.

7 Conclusions and Future Work

We make two major technical contributions in this work which enable us to improve the quality of semi-parametric NMT on broad domain datasets. First, we propose using n-gram retrieval, with standard Inverse Document Frequency similarity and with dense vector representations, that takes into account local sentence similarities that are critical to translation. As a result we are able to retrieve useful candidates even for broad domain tasks with little train-test overlap. Second, we propose a novel architecture to encode retrieved source-target pairs, allowing the model to distinguish useful information from noise by encoding the retrieved targets in context of the current translation task.

We demonstrate, for the first time, that semi-parametric methods can beat neural models by significant margins on multi-domain Machine Translation. By successfully training semi-parametric

neural models on a broad domain dataset (WMT), we also open the door for non-parametric adaptation, showing huge improvements on new domains without any parameter updates.

While we constrain this work to retrieved context, our architecture can be utilized to incorporate information from other sources of context, including documents, bilingual dictionaries etc. Using dense representations for retrieval also allows extending semi-parametric neural methods to other input modalities, including images and speech.

With this work, we hope to motivate further investigation into semi-parametric neural models for and beyond Neural Machine Translation.

Acknowledgments

We would like to thank Naveen Arivazhagan, Macduff Hughes, Dmitry Lepikhin, Mia Chen, Yuan Cao, Ciprian Chelba, Zhifeng Chen, Melvin Johnson and other members of the Google Brain and Google Translate teams for their useful inputs and discussions. We would also like to thank the entire Lingvo development team for their foundational contributions to this project.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *International Conference on Learning Representations*.
- Qian Cao and Deyi Xiong. 2018. [Encoding gated translation memory into neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3042–3047. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. *arXiv preprint arXiv:1704.00051*.
- Mia Xu Chen, Orhan Firat, Ankur Bapna, et al. 2018. The best of both worlds: Combining recent advances in neural machine translation. *arXiv preprint arXiv:1804.09849*.
- 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*.
- Leonard Dahlmann, Evgeny Matusov, Pavel Petrushkov, and Shahram Khadivi. 2017. Neural machine translation leveraging phrase-based models in a hybrid search. *arXiv preprint arXiv:1708.03271*.
- M Amin Farajian, Marco Turchi, Matteo Negri, Nicola Bertoldi, and Marcello Federico. 2017a. Neural vs. phrase-based machine translation in a multi-domain scenario. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 280–284.
- M Amin Farajian, Marco Turchi, Matteo Negri, and Marcello Federico. 2017b. Multi-domain neural machine translation through unsupervised adaptation. In *Proceedings of the Second Conference on Machine Translation*, pages 127–137.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). *CoRR*, abs/1705.03122.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor OK Li. 2017. Search engine guided non-parametric neural machine translation. *arXiv preprint arXiv:1705.07267*.
- Mandy Guo, Qinlan Shen, Yinfei Yang, Heming Ge, Daniel Cer, Gustavo Hernandez Abrego, Keith Stevens, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Effective parallel corpus mining using bilingual sentence embeddings. *arXiv preprint arXiv:1807.11906*.
- Kelvin Guu, Tatsunori B Hashimoto, Yonatan Oren, and Percy Liang. 2018. Generating sentences by editing prototypes. *Transactions of the Association of Computational Linguistics*, 6:437–450.
- Shirley Anugrah Hayati, Raphael Olivier, Pravalika Avvaru, Pengcheng Yin, Anthony Tomasic, and Graham Neubig. 2018. Retrieval-based neural code generation. *arXiv preprint arXiv:1808.10025*.
- Pierre Isabelle, Colin Cherry, and George Foster. 2017. A challenge set approach to evaluating machine translation. *arXiv preprint arXiv:1704.07431*.
- Shaojie Jiang and Maarten de Rijke. 2018. Why are sequence-to-sequence models so dull? *EMNLP 2018*, page 81.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th annual meeting of the ACL on interactive poster and*

- demonstration sessions*, pages 177–180. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. *arXiv preprint arXiv:1706.03872*.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, pages 48–54. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Elina Lagoudaki. Translation memories survey 2006: Users perceptions around tm use.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fannjiang, and David Sussillo. 2018. Hallucinations in neural machine translation.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Uwe Reinke. 2013. State of the art in translation memory technology. *Translation: Computation, Corpora, Cognition*, 3(1).
- Stephen Robertson. 2004. Understanding inverse document frequency: on theoretical arguments for idf. *Journal of documentation*, 60(5):503–520.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. One-shot learning with memory-augmented neural networks. *arXiv preprint arXiv:1605.06065*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Jonathan Shen, Patrick Nguyen, Yonghui Wu, et al. 2019. Lingvo: a modular and scalable framework for sequence-to-sequence modeling. *arXiv preprint arXiv:1902.08295*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Ralf Steinberger, Bruno Pouliquen, Anna Widiger, Camelia Ignat, Tomaz Erjavec, Dan Tufis, and Dániel Varga. 2006. The jrc-acquis: A multilingual aligned parallel corpus with 20+ languages. *arXiv preprint cs/0609058*.
- 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.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. Re-thinking the inception architecture for computer vision. *CoRR*, abs/1512.00567.
- Jörg Tiedemann. Parallel data, tools and interfaces in opus.
- Jörg Tiedemann. 2009. News from opus-a collection of multilingual parallel corpora with tools and interfaces.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerald Tesauro, Bowen Zhou, and Jing Jiang. 2018. R3: Reinforced ranker-reader for open-domain question answering. In *Conference on Artificial Intelligence*.
- Jason Weston, Emily Dinan, and Alexander H Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. *arXiv preprint arXiv:1808.04776*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Joern Wuebker, Spence Green, and John DeNero. 2015. Hierarchical incremental adaptation for statistical machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1059–1065.
- Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 2018a. [Improving the transformer translation model with document-level context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 533–542. Association for Computational Linguistics.

- Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2017. Improving neural machine translation through phrase-based forced decoding. *arXiv preprint arXiv:1711.00309*.
- Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2018b. Guiding neural machine translation with retrieved translation pieces. *arXiv preprint arXiv:1804.02559*.
- Jie Zhou, Ying Cao, Xuguang Wang, Peng Li, and Wei Xu. 2016. [Deep recurrent models with fast-forward connections for neural machine translation](#). *CoRR*, abs/1606.04199.
- Long Zhou, Wenpeng Hu, Jiajun Zhang, and Chengqing Zong. 2017. Neural system combination for machine translation. *arXiv preprint arXiv:1704.06393*.