

Consistency by Agreement in Zero-shot Neural Machine Translation

Maruan Al-Shedivat*

Carnegie Mellon University
Pittsburgh, PA 15213

alshedivat@cs.cmu.edu

Ankur P. Parikh

Google AI Language
New York, NY 10011

aparikh@google.com

Abstract

Generalization and reliability of multilingual translation often highly depend on the amount of available parallel data for each language pair of interest. In this paper, we focus on zero-shot generalization—a challenging setup that tests models on translation directions they have not been optimized for at training time. To solve the problem, we (i) reformulate multilingual translation as probabilistic inference, (ii) define the notion of zero-shot consistency and show why standard training often results in models unsuitable for zero-shot tasks, and (iii) introduce a consistent agreement-based training method that encourages the model to produce equivalent translations of parallel sentences in auxiliary languages. We test our multilingual NMT models on multiple public zero-shot translation benchmarks (IWSLT17, UN corpus, Europarl) and show that agreement-based learning often results in 2-3 BLEU zero-shot improvement over strong baselines without any loss in performance on supervised translation directions.

1 Introduction

Machine translation (MT) has made remarkable advances with the advent of deep learning approaches (Bojar et al., 2016; Wu et al., 2016; Crego et al., 2016; Junczys-Dowmunt et al., 2016). The progress was largely driven by the encoder-decoder framework (Sutskever et al., 2014; Cho et al., 2014) and typically supplemented with an attention mechanism (Bahdanau et al., 2014; Luong et al., 2015b).

Compared to the traditional phrase-based systems (Koehn, 2009), neural machine translation (NMT) requires large amounts of data in order to reach high performance (Koehn and Knowles, 2017). Using NMT in a multilingual setting exacerbates the problem by the fact that given k languages




Figure 1: Agreement-based training of a multilingual NMT. At training time, given English-French ($En \leftrightarrow Fr$) and English-German ($En \leftrightarrow De$) parallel sentences, the model not only is trained to translate between the pair but also to agree on translations into a third language.

translating between all pairs would require $O(k^2)$ parallel training corpora (and $O(k^2)$ models).

In an effort to address the problem, different multilingual NMT approaches have been proposed recently. Luong et al. (2015a); Firat et al. (2016a) proposed to use $O(k)$ encoders/decoders that are then intermixed to translate between language pairs. Johnson et al. (2016) proposed to use a single model and prepend special symbols to the source text to indicate the target language, which has later been extended to other text preprocessing approaches (Ha et al., 2017) as well as language-conditional parameter generation for encoders and decoders of a single model (Platanios et al., 2018).

Johnson et al. (2016) also show that a single multilingual system could potentially enable *zero-shot* translation, *i.e.*, it can translate between language pairs not seen in training. For example, given 3 languages—German (De), English (En), and French (Fr)—and training parallel data only for (De, En) and (En, Fr), at test time, the system could additionally translate between (De, Fr).

Zero-shot translation is an important problem. Solving the problem could significantly improve data efficiency—a single multilingual model would

*Work done at Google.

be able to generalize and translate between any of the $O(k^2)$ language pairs after being trained only on $O(k)$ parallel corpora. However, performance on zero-shot tasks is often unstable and significantly lags behind the supervised directions. Moreover, attempts to improve zero-shot performance by fine-tuning (Firat et al., 2016b; Sestorain et al., 2018) may negatively impact other directions.

In this work, we take a different approach and aim to improve the training procedure of Johnson et al. (2016). First, we analyze multilingual translation problem from a probabilistic perspective and define the notion of *zero-shot consistency* that gives insights as to why the vanilla training method may not yield models with good zero-shot performance. Next, we propose a novel training objective and a modified learning algorithm that achieves consistency via agreement-based learning (Liang et al., 2006, 2008) and improves zero-shot translation. Our training procedure encourages the model to produce equivalent translations of parallel training sentences into an auxiliary language (Figure 1) and is provably zero-shot consistent. In addition, we make a simple change to the neural decoder to make the agreement losses fully differentiable.

We conduct experiments on IWSLT17 (Mauro et al., 2017), UN corpus (Ziems et al., 2016), and Europarl (Koehn, 2017), carefully removing complete pivots from the training corpora. Agreement-based learning results in up to +3 BLEU zero-shot improvement over the baseline, compares favorably (up to +2.4 BLEU) to other approaches in the literature (Cheng et al., 2017; Sestorain et al., 2018), is competitive with pivoting, and does not lose in performance on supervised directions.

2 Related work

A simple (and yet effective) baseline for zero-shot translation is pivoting that chain-translates, first to a pivot language, then to a target (Cohn and Lapata, 2007; Wu and Wang, 2007; Utiyama and Isahara, 2007). Despite being a pipeline, pivoting gets better as the supervised models improve, which makes it a strong baseline in the zero-shot setting. Cheng et al. (2017) proposed a joint pivoting learning strategy that leads to further improvements.

Lu et al. (2018) and Arivazhagan et al. (2018) proposed different techniques to obtain “neural interlingual” representations that are passed to the decoder. Sestorain et al. (2018) proposed another fine-tuning technique that uses dual learning (He

et al., 2016), where a language model is used to provide a signal for fine-tuning zero-shot directions.

Another family of approaches is based on distillation (Hinton et al., 2014; Kim and Rush, 2016). Along these lines, Firat et al. (2016b) proposed to fine tune a multilingual model to a specified zero-shot-direction with pseudo-parallel data and Chen et al. (2017) proposed a teacher-student framework. While this can yield solid performance improvements, it also adds multi-staging overhead and often does not preserve performance of a single model on the supervised directions. We note that our approach (and agreement-based learning in general) is somewhat similar to distillation at training time, which has been explored for large-scale single-task prediction problems (Anil et al., 2018).

A setting harder than zero-shot is that of fully unsupervised translation (Ravi and Knight, 2011; Artetxe et al., 2017; Lample et al., 2017, 2018) in which no parallel data is available for training. The ideas proposed in these works (e.g., bilingual dictionaries (Conneau et al., 2017), backtranslation (Sennrich et al., 2015a) and language models (He et al., 2016)) are complementary to our approach, which encourages agreement among different translation directions in the zero-shot multilingual setting.

3 Background

We start by establishing more formal notation and briefly reviewing some background on encoder-decoder multilingual machine translation from a probabilistic perspective.

3.1 Notation

Languages. We assume that we are given a collection of k languages, $L_1, \dots, L_k$, that share a common vocabulary, V . A language, L_i , is defined by the marginal probability $\mathbb{P}(\mathbf{x}_i)$ it assigns to sentences (i.e., sequences of tokens from the vocabulary), denoted $\mathbf{x}_i := (x^1, \dots, x^l)$, where l is the length of the sequence. All languages together define a joint probability distribution, $\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_k)$, over k -tuples of *equivalent sentences*.

Corpora. While each sentence may have an equivalent representation in all languages, we assume that we have access to only partial sets of equivalent sentences, which form *corpora*. In this work, we consider *bilingual* corpora, denoted C_{ij} , that contain pairs of sentences sampled from $\mathbb{P}(\mathbf{x}_i, \mathbf{x}_j)$ and *monolingual* corpora, denoted C_i , that contain sentences sampled from $\mathbb{P}(\mathbf{x}_i)$.




Figure 2: Translation graph: Languages (nodes), parallel corpora (solid edges), and zero-shot directions (dotted edges).

Translation. Finally, we define a *translation task* from language L_i to L_j as learning to model the conditional distribution $\mathbb{P}(\mathbf{x}_j | \mathbf{x}_i)$. The set of k languages along with translation tasks can be represented as a directed graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with a set of k nodes, $\mathcal{V}$, that represent languages and edges, $\mathcal{E}$, that indicate translation directions. We further distinguish between two disjoint subsets of edges: (i) supervised edges, $\mathcal{E}_s$, for which we have parallel data, and (ii) zero-shot edges, $\mathcal{E}_0$, that correspond to zero-shot translation tasks. Figure 2 presents an example translation graph with supervised edges ($\text{En} \leftrightarrow \text{Es}$, $\text{En} \leftrightarrow \text{Fr}$, $\text{En} \leftrightarrow \text{Ru}$) and zero-shot edges ($\text{Es} \leftrightarrow \text{Fr}$, $\text{Es} \leftrightarrow \text{Ru}$, $\text{Fr} \leftrightarrow \text{Ru}$). We will use this graph as our running example.

3.2 Encoder-decoder framework

First, consider a purely bilingual setting, where we learn to translate from a source language, L_s , to a target language, L_t . We can train a translation model by optimizing the conditional log-likelihood of the bilingual data under the model:

$$\hat{\theta} := \arg \max_{\theta} \sum_{\mathcal{C}_{st}} \log \mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{x}_s) \quad (1)$$

where $\hat{\theta}$ are the estimated parameters of the model.

The encoder-decoder framework introduces a latent sequence, $\mathbf{u}$, and represents the model as:

$$\mathbb{P}_{\theta}(\mathbf{x}_t | \mathbf{x}_s) = \mathbb{P}_{\theta}^{\text{dec}}(\mathbf{x}_t | \mathbf{u} = f_{\theta}^{\text{enc}}(\mathbf{x}_s)) \quad (2)$$

where $f_{\theta}^{\text{enc}}(\mathbf{x}_s)$ is the encoder that maps a source sequence to a sequence of latent representations, $\mathbf{u}$, and the decoder defines $\mathbb{P}_{\theta}^{\text{dec}}(\mathbf{x}_t | \mathbf{u})$.¹ Note that $\mathbf{u}$ is usually deterministic with respect to $\mathbf{x}_s$ and accurate representation of the conditional distribution highly depends on the decoder. In neural machine translation, the exact forms of encoder and decoder are specified using RNNs (Sutskever et al., 2014),

¹Slightly abusing the notation, we use θ to denote all parameters of the model: embeddings, encoder, and decoder.

CNNs (Gehring et al., 2016), and attention (Bahdanau et al., 2014; Vaswani et al., 2017) as building blocks. The decoding distribution, $\mathbb{P}_{\theta}^{\text{dec}}(\mathbf{x}_t | \mathbf{u})$, is typically modeled autoregressively.

3.3 Multilingual neural machine translation

In the multilingual setting, we would like to learn to translate in *all directions* having access to only few parallel bilingual corpora. In other words, we would like to learn a collection of models, $\{\mathbb{P}_{\theta}(\mathbf{x}_j | \mathbf{x}_i)\}_{i,j \in \mathcal{E}}$. We can assume that models are independent and choose to learn them by maximizing the following objective:

$$\mathcal{L}^{\text{ind}}(\theta) = \sum_{i,j \in \mathcal{E}_s} \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{C}_{ij}} \log \mathbb{P}_{\theta}(\mathbf{x}_j | \mathbf{x}_i) \quad (3)$$

In the statistics literature, this estimation approach is called *maximum composite likelihood* (Besag, 1975; Lindsay, 1988) as it composes the objective out of (sometimes weighted) terms that represent conditional sub-likelihoods (in our example, $\mathbb{P}_{\theta}(\mathbf{x}_j | \mathbf{x}_i)$). Composite likelihoods are easy to construct and tractable to optimize as they do not require representing the full likelihood, which would involve integrating out variables unobserved in the data (see Appendix A.1).

Johnson et al. (2016) proposed to train a multilingual NMT systems by optimizing a composite likelihood objective (3) while representing all conditional distributions, $\mathbb{P}_{\theta}(\mathbf{x}_j | \mathbf{x}_i)$, with a *shared* encoder and decoder and using language tags, l_t , to distinguish between translation directions:

$$\mathbb{P}(\mathbf{x}_t | \mathbf{x}_s) = \mathbb{P}_{\theta}^{\text{dec}}(\mathbf{x}_t | \mathbf{u}_{st} = f_{\theta}^{\text{enc}}(\mathbf{x}_s, l_t)) \quad (4)$$

This approach has numerous advantages including: (a) simplicity of training and the architecture (by slightly changing the training data, we convert a bilingual NMT into a multilingual one), (b) sharing parameters of the model between different translation tasks that may lead to better and more robust representations. Johnson et al. (2016) also show that resulting models seem to exhibit some degree of zero-shot generalization enabled by parameter sharing. However, since we lack data for zero-shot directions, composite likelihood (3) misses the terms that correspond to the zero-shot models, and hence has no statistical guarantees for performance on zero-shot tasks.²

²In fact, since the objective (3) assumes that the models are independent, plausible zero-shot performance would be more indicative of the limited capacity of the model or artifacts in the data (e.g., presence of multi-parallel sentences) rather than zero-shot generalization.

4 Zero-shot generalization & consistency

Multilingual MT systems can be evaluated in terms of *zero-shot performance*, or quality of translation along the directions they have not been optimized for (*e.g.*, due to lack of data). We formally define zero-shot generalization via consistency.

Definition 1 (Expected Zero-shot Consistency)

Let $\mathcal{E}_s$ and $\mathcal{E}_0$ be supervised and zero-shot tasks, respectively. Let $\ell(\cdot)$ be a non-negative loss function and $\mathcal{M}$ be a model with maximum expected supervised loss bounded by some $\varepsilon > 0$:

$$\max_{(i,j) \in \mathcal{E}_s} \mathbb{E}_{\mathbf{x}_i, \mathbf{x}_j} [\ell(\mathcal{M})] < \varepsilon$$

We call $\mathcal{M}$ zero-shot consistent with respect to $\ell(\cdot)$ if for some $\kappa(\varepsilon) > 0$

$$\max_{(i,j) \in \mathcal{E}_0} \mathbb{E}_{\mathbf{x}_i, \mathbf{x}_j} [\ell(\mathcal{M})] < \kappa(\varepsilon),$$

where $\kappa(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

In other words, we say that a machine translation system is zero-shot consistent if low error on supervised tasks implies a low error on zero-shot tasks in expectation (*i.e.*, the system generalizes). We also note that our notion of consistency somewhat resembles error bounds in the domain adaptation literature (Ben-David et al., 2010).

In practice, it is attractive to have MT systems that are guaranteed to exhibit zero-shot generalization since the access to parallel data is always limited and training is computationally expensive. While the training method of Johnson et al. (2016) does not have guarantees, we show that our proposed approach is provably zero-shot consistent.

5 Approach

We propose a new training objective for multilingual NMT architectures with shared encoders and decoders that avoids the limitations of pure composite likelihoods. Our method is based on the idea of agreement-based learning initially proposed for learning consistent alignments in phrase-based statistical machine translation (SMT) systems (Liang et al., 2006, 2008). In terms of the final objective function, the method ends up being reminiscent of distillation (Kim and Rush, 2016), but suitable for joint multilingual training.

5.1 Agreement-based likelihood

To introduce agreement-based objective, we use the graph from Figure 2 that defines translation tasks between 4 languages (En, Es, Fr, Ru). In particular, consider the composite likelihood objective (3) for a pair of En – Fr sentences, $(\mathbf{x}_{\text{En}}, \mathbf{x}_{\text{Fr}})$:

$$\begin{aligned} \mathcal{L}_{\text{EnFr}}^{\text{ind}}(\theta) &= \log [\mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}} | \mathbf{x}_{\text{En}}) \mathbb{P}_{\theta}(\mathbf{x}_{\text{En}} | \mathbf{x}_{\text{Fr}})] \\ &= \log \left[\sum_{\mathbf{z}'_{\text{Es}}, \mathbf{z}'_{\text{Ru}}} \mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}}, \mathbf{z}'_{\text{Es}}, \mathbf{z}'_{\text{Ru}} | \mathbf{x}_{\text{En}}) \times \right. \\ &\quad \left. \sum_{\mathbf{z}''_{\text{Es}}, \mathbf{z}''_{\text{Ru}}} \mathbb{P}_{\theta}(\mathbf{x}_{\text{En}}, \mathbf{z}''_{\text{Es}}, \mathbf{z}''_{\text{Ru}} | \mathbf{x}_{\text{Fr}}) \right] \end{aligned} \quad (5)$$

where we introduced latent translations into Spanish (Es) and Russian (Ru) and marginalized them out (by virtually summing over all sequences in the corresponding languages). Again, note that this objective assumes independence of $\text{En} \rightarrow \text{Fr}$ and $\text{Fr} \rightarrow \text{En}$ models.

Following Liang et al. (2008), we propose to tie together the single prime and the double prime latent variables, $\mathbf{z}_{\text{Es}}$ and $\mathbf{z}_{\text{Ru}}$, to encourage agreement between $\mathbb{P}_{\theta}(\mathbf{x}_{\text{En}}, \mathbf{z}_{\text{Es}}, \mathbf{z}_{\text{Ru}} | \mathbf{x}_{\text{Fr}})$ and $\mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}}, \mathbf{z}_{\text{Es}}, \mathbf{z}_{\text{Ru}} | \mathbf{x}_{\text{En}})$ on the latent translations. We interchange the sum and the product operations inside the log in (5), denote $\mathbf{z} := (\mathbf{z}_{\text{Es}}, \mathbf{z}_{\text{Ru}})$ to simplify notation, and arrive at the following new objective function:

$$\begin{aligned} \mathcal{L}_{\text{EnFr}}^{\text{agree}}(\theta) &:= \\ \log \sum_{\mathbf{z}} \mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}}, \mathbf{z} | \mathbf{x}_{\text{En}}) \mathbb{P}_{\theta}(\mathbf{x}_{\text{En}}, \mathbf{z} | \mathbf{x}_{\text{Fr}}) \end{aligned} \quad (6)$$

Next, we factorize each term as:

$$\mathbb{P}(\mathbf{x}, \mathbf{z} | \mathbf{y}) = \mathbb{P}(\mathbf{x} | \mathbf{z}, \mathbf{y}) \mathbb{P}(\mathbf{z} | \mathbf{y})$$

Assuming $\mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}} | \mathbf{z}, \mathbf{x}_{\text{En}}) \approx \mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}} | \mathbf{x}_{\text{En}})$,³ the objective (6) decomposes into two terms:

$$\begin{aligned} \mathcal{L}_{\text{EnFr}}^{\text{agree}}(\theta) & \approx \underbrace{\log \mathbb{P}_{\theta}(\mathbf{x}_{\text{Fr}} | \mathbf{x}_{\text{En}}) + \log \mathbb{P}_{\theta}(\mathbf{x}_{\text{En}} | \mathbf{x}_{\text{Fr}})}_{\text{composite likelihood terms}} + \\ & \underbrace{\log \sum_{\mathbf{z}} \mathbb{P}_{\theta}(\mathbf{z} | \mathbf{x}_{\text{En}}) \mathbb{P}_{\theta}(\mathbf{z} | \mathbf{x}_{\text{Fr}})}_{\text{agreement term}} \end{aligned} \quad (7)$$

³This means that it is sufficient to condition on a sentence in one of the languages to determine probability of a translation in any other language.

We call the expression given in (7) *agreement-based likelihood*. Intuitively, this objective is the likelihood of observing parallel sentences $(\mathbf{x}_{\text{En}}, \mathbf{x}_{\text{Fr}})$ and having sub-models $\mathbb{P}_\theta(\mathbf{z} \mid \mathbf{x}_{\text{En}})$ and $\mathbb{P}_\theta(\mathbf{z} \mid \mathbf{x}_{\text{Fr}})$ agree on all translations into Es and Ru at the same time.

Lower bound. Summation in the agreement term over $\mathbf{z}$ (i.e., over possible translations into Es and Ru in our case) is intractable. Switching back from $\mathbf{z}$ to $(\mathbf{z}_{\text{Es}}, \mathbf{z}_{\text{Ru}})$ notation and using Jensen’s inequality, we lower bound it with cross-entropy:⁴

$$\begin{aligned} \log \sum_{\mathbf{z}} \mathbb{P}_\theta(\mathbf{z} \mid \mathbf{x}_{\text{En}}) \mathbb{P}_\theta(\mathbf{z} \mid \mathbf{x}_{\text{Fr}}) \\ \geq \mathbb{E}_{\mathbf{z}_{\text{Es}} \mid \mathbf{x}_{\text{En}}} [\log \mathbb{P}_\theta(\mathbf{z}_{\text{Es}} \mid \mathbf{x}_{\text{Fr}})] + \\ \mathbb{E}_{\mathbf{z}_{\text{Ru}} \mid \mathbf{x}_{\text{En}}} [\log \mathbb{P}_\theta(\mathbf{z}_{\text{Ru}} \mid \mathbf{x}_{\text{Fr}})] \end{aligned} \quad (8)$$

We can estimate the expectations in the lower bound on the agreement terms by sampling $\mathbf{z}_{\text{Es}} \sim \mathbb{P}_\theta(\mathbf{z}_{\text{Es}} \mid \mathbf{x}_{\text{En}})$ and $\mathbf{z}_{\text{Ru}} \sim \mathbb{P}_\theta(\mathbf{z}_{\text{Ru}} \mid \mathbf{x}_{\text{En}})$. In practice, instead of sampling we use greedy, continuous decoding (with a fixed maximum sequence length) that also makes $\mathbf{z}_{\text{Es}}$ and $\mathbf{z}_{\text{Ru}}$ differentiable with respect to parameters of the model.

5.2 Consistency by agreement

We argue that models produced by maximizing agreement-based likelihood (7) are zero-shot consistent. Informally, consider again our running example from Figure 2. Given a pair of parallel sentences in (En, Fr) , agreement loss encourages translations from En to $\{\text{Es}, \text{Ru}\}$ and translations from Fr to $\{\text{Es}, \text{Ru}\}$ to coincide. Note that $\text{En} \rightarrow \{\text{Es}, \text{Fr}, \text{Ru}\}$ are supervised directions. Therefore, agreement ensures that translations along the zero-shot edges in the graph match supervised translations. Formally, we state it as:

Theorem 2 (Agreement Zero-shot Consistency)

Let L_1, L_2 , and L_3 be a collection of languages with $L_1 \leftrightarrow L_2$ and $L_2 \leftrightarrow L_3$ be supervised while $L_1 \leftrightarrow L_3$ be a zero-shot direction. Let $\mathbb{P}_\theta(\mathbf{x}_j \mid \mathbf{x}_i)$ be sub-models represented by a multilingual MT system. If the expected agreement-based loss, $\mathbb{E}_{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3} [\mathcal{L}_{12}^{\text{agree}}(\theta) + \mathcal{L}_{23}^{\text{agree}}(\theta)]$, is bounded by some $\varepsilon > 0$, then, under some mild technical assumptions on the true distribution of the equivalent translations, the zero-shot cross-entropy

⁴Note that expectations in (8) are conditional on $\mathbf{x}_{\text{En}}$. Symmetrically, we can have a lower bound with expectations conditional on $\mathbf{x}_{\text{Fr}}$. In practice, we symmetrize the objective.

Algorithm 1 Agreement-based M-NMT training

input Architecture (GNMT), agreement coefficient (γ)

- 1: Initialize: $\theta \leftarrow \theta_0$
- 2: **while** not (converged or step limit reached) **do**
- 3: Get a mini-batch of parallel src-tgt pairs, $(\mathbf{X}_s, \mathbf{X}_t)$
- 4: Supervised loss: $\mathcal{L}^{\text{sup}}(\theta) \leftarrow \log \mathbb{P}_\theta(\mathbf{X}_t \mid \mathbf{X}_s)$
- 5: Auxiliary languages: $L_a \sim \text{Unif}(\{1, \dots, k\})$
- 6: Auxiliary translations:

$$\mathbf{Z}_{a \leftarrow s} \leftarrow \text{Decode}(\mathbf{Z}_a \mid f_\theta^{\text{enc}}(\mathbf{X}_s, L_a))$$

$$\mathbf{Z}_{a \leftarrow t} \leftarrow \text{Decode}(\mathbf{Z}_a \mid f_\theta^{\text{enc}}(\mathbf{X}_t, L_a))$$
- 7: Agreement log-probabilities:

$$\ell_{a \leftarrow s}^t \leftarrow \log \mathbb{P}_\theta(\mathbf{Z}_{a \leftarrow s} \mid \mathbf{X}_t)$$

$$\ell_{a \leftarrow t}^s \leftarrow \log \mathbb{P}_\theta(\mathbf{Z}_{a \leftarrow t} \mid \mathbf{X}_s)$$
- 8: Apply stop-gradients to supervised $\ell_{a \leftarrow s}^t$ and $\ell_{a \leftarrow t}^s$
- 9: Total loss: $\mathcal{L}^{\text{total}}(\theta) \leftarrow \mathcal{L}^{\text{sup}}(\theta) + \gamma(\ell_{a \leftarrow s}^t + \ell_{a \leftarrow t}^s)$
- 10: Update: $\theta \leftarrow \text{optimizer_update}(\mathcal{L}^{\text{total}}, \theta)$
- 11: **end while**

output θ

loss is bounded as follows:

$$\mathbb{E}_{\mathbf{x}_1, \mathbf{x}_3} [-\log \mathbb{P}_\theta(\mathbf{x}_3 \mid \mathbf{x}_1)] \leq \kappa(\varepsilon)$$

where $\kappa(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

For discussion of the assumptions and details on the proof of the bound, see Appendix A.2. Note that Theorem 2 is straightforward to extend from triplets of languages to arbitrary connected graphs, as given in the following corollary.

Corollary 3 *Agreement-based learning yields zero shot consistent MT models (with respect to the cross entropy loss) for arbitrary translation graphs as long as supervised directions span the graph.*

Alternative ways to ensure consistency. Note that there are other ways to ensure zero-shot consistency, e.g., by fine-tuning or post-processing a trained multilingual model. For instance, pivoting through an intermediate language is also zero-shot consistent, but the proof requires stronger assumptions about the quality of the supervised source-pivot model.⁵ Similarly, using model distillation (Kim and Rush, 2016; Chen et al., 2017) would be also provably consistent under the same assumptions as given in Theorem 2, but for a single, pre-selected zero-shot direction. Note that our proposed agreement-based learning framework is provably consistent for *all* zero-shot directions and does not require any post-processing. For discussion of the alternative approaches and consistency proof for pivoting, see Appendix A.3.

⁵Intuitively, we have to assume that source-pivot model does not assign high probabilities to unlikely translations as the pivot-target model may react to those unpredictably.




Figure 3: **A.** Computation graph for the encoder. The representations depend on the input sequence and the target language tag. **B.** Computation graph for the agreement loss. First, encode source and target sequences with the auxiliary language tags. Next, decode $\mathbf{z}_{\text{Es}}$ from both $\mathbf{x}_{\text{En}}$ and $\mathbf{x}_{\text{Fr}}$ using continuous greedy decoder. Finally, evaluate log probabilities, $\log \mathbb{P}_{\theta}(\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}}) | \mathbf{x}_{\text{Fr}})$ and $\log \mathbb{P}_{\theta}(\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}}) | \mathbf{x}_{\text{En}})$, and compute a sample estimate of the agreement loss.

5.3 Agreement-based learning algorithm

Having derived a new objective function (7), we can now learn consistent multilingual NMT models using stochastic gradient method with a couple of extra tricks (Algorithm 1). The computation graph for the agreement loss is given in Figure 3.

Subsampling auxiliary languages. Computing agreement over *all* languages for each pair of sentences at training time would be quite computationally expensive (to agree on k translations, we would need to encode-decode the source and target sequences k times each). However, since the agreement lower bound is a sum over expectations (8), we can approximate it by subsampling: at each training step (and for each sample in the mini-batch), we pick an auxiliary language uniformly at random and compute stochastic approximation of the agreement lower bound (8) for that language only. This stochastic approximation is simple, unbiased, and reduces per step computational overhead for the agreement term from $O(k)$ to $O(1)$.⁶

Overview of the agreement loss computation. Given a pair of parallel sentences, $\mathbf{x}_{\text{En}}$ and $\mathbf{x}_{\text{Fr}}$, and an auxiliary language, say Es , an estimate of the lower bound on the agreement term (8) is computed as follows. First, we concatenate Es language tags to both $\mathbf{x}_{\text{En}}$ and $\mathbf{x}_{\text{Fr}}$ and encode the sequences so that both can be translated into Es (the encoding

process is depicted in Figure 3A). Next, we decode each of the encoded sentences and obtain auxiliary translations, $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}})$ and $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}})$, depicted as blue blocks in Figure 3B. Note that we now can treat pairs $(\mathbf{x}_{\text{Fr}}, \mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}}))$ and $(\mathbf{x}_{\text{En}}, \mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}}))$ as new parallel data for $\text{En} \rightarrow \text{Es}$ and $\text{Fr} \rightarrow \text{Es}$.

Finally, using these pairs, we can compute two log-probability terms (Figure 3B):

$$\begin{aligned} \log \mathbb{P}_{\theta}(\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}}) | \mathbf{x}_{\text{En}}) \\ \log \mathbb{P}_{\theta}(\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}}) | \mathbf{x}_{\text{Fr}}) \end{aligned} \quad (9)$$

using encoding-decoding with teacher forcing (same way as typically done for the supervised directions). Crucially, note that $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}})$ corresponds to a supervised direction, $\text{En} \rightarrow \text{Es}$, while $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}})$ corresponds to zero-shot, $\text{Fr} \rightarrow \text{Es}$. We want each of the components to (i) improve the zero-shot direction while (ii) minimally affecting the supervised direction. To achieve (i), we use continuous decoding, and for (ii) we use stop-gradient-based protection of the supervised directions. Both techniques are described below.

Greedy continuous decoding. In order to make $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}})$ and $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}})$ differentiable with respect to θ (hence, *continuous* decoding), at each decoding step t , we treat the output of the RNN, $\mathbf{h}^t$, as the key and use dot-product attention over the embedding vocabulary, $\mathbf{V}$, to construct $\mathbf{z}_{\text{Es}}^t$:

$$\mathbf{z}_{\text{Es}}^t := \text{softmax} \left\{ (\mathbf{h}^t)^{\top} \mathbf{V} \right\} \mathbf{V} \quad (10)$$

In other words, auxiliary translations, $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{En}})$ and $\mathbf{z}_{\text{Es}}(\mathbf{x}_{\text{Fr}})$, are fixed length sequences of differentiable embeddings computed in a greedy fashion.

⁶In practice, note that there is still a constant factor overhead due to extra encoding-decoding steps to/from auxiliary languages, which is about $\times 4$ when training on a single GPU. Parallelizing the model across multiple GPUs would easily compensate this overhead.

Protecting supervised directions. Algorithm 1 scales agreement losses by a small coefficient γ . We found experimentally that training could be sensitive to this hyperparameter since the agreement loss also affects the supervised sub-models. For example, agreement of $En \rightarrow Es$ (supervised) and $Fr \rightarrow Es$ (zero-shot) may push the former towards a worse translation, especially at the beginning of training. To stabilize training, we apply the `stop_gradient` operator to the log probabilities and samples produced by the supervised sub-models before computing the agreement terms (9), to zero-out the corresponding gradient updates.

6 Experiments

We evaluate agreement-based training against baselines from the literature on three public datasets that have multi-parallel *evaluation data* that allows assessing zero-shot performance. We report results in terms of the BLEU score (Papineni et al., 2002) that was computed using `mteval-v13a.perl`.

6.1 Datasets

UN corpus. Following the setup introduced in Sestorain et al. (2018), we use two datasets, *UNcorpus-1* and *UNcorpus-2*, derived from the United Nations Parallel Corpus (Ziems et al., 2016). *UNcorpus-1* consists of data in 3 languages, En, Es, Fr , where *UNcorpus-2* has Ru as the 4th language. For training, we use parallel corpora between En and the rest of the languages, each about 1M sentences, sub-sampled from the official training data in a way that ensures no multi-parallel training data. The *dev* and *test* sets contain 4,000 sentences and are all multi-parallel.

Europarl v7⁷. We consider the following languages: De, En, Es, Fr . For training, we use parallel data between En and the rest of the languages (about 1M sentences per corpus), preprocessed to avoid multi-parallel sentences, as was also done by Cheng et al. (2017) and Chen et al. (2017) and described below. The *dev* and *test* sets contain 2,000 multi-parallel sentences.

IWSLT17⁸. We use data from the official multilingual task: 5 languages (De, En, It, Nl, Ro), 20 translation tasks of which 4 zero-shot ($De \leftrightarrow Nl$ and $It \leftrightarrow Ro$) and the rest 16 supervised. Note that this dataset has a significant

overlap between parallel corpora in the supervised directions (up to 100K sentence pairs per direction). This implicitly makes the dataset multi-parallel and defeats the purpose of zero-shot evaluation (Dabre et al., 2017). To avoid spurious effects, we also derived **IWSLT17*** dataset from the original one by restricting supervised data to only $En \leftrightarrow \{De, Nl, It, Ro\}$ and removing overlapping pivoting sentences. We report results on both the official and preprocessed datasets.

Preprocessing. To properly evaluate systems in terms of zero-shot generalization, we preprocess Europarl and IWSLT* to avoid multi-lingual parallel sentences of the form *source-pivot-target*, where *source-target* is a zero-shot direction. To do so, we follow Cheng et al. (2017); Chen et al. (2017) and randomly split the overlapping pivot sentences of the original *source-pivot* and *pivot-target* corpora into two parts and merge them separately with the non-overlapping parts for each pair. Along with each parallel training sentence, we save information about source and target tags, after which all the data is combined and shuffled. Finally, we use a shared multilingual subword vocabulary (Sennrich et al., 2015b) on the training data (with 32K merge ops), separately for each dataset.

6.2 Training and evaluation

Additional details on the hyperparameters can be found in Appendix A.4.

Models. We use a smaller version of the GNMT architecture (Wu et al., 2016) in all our experiments: 512-dimensional embeddings (separate for source and target sides), 2 bidirectional LSTM layers of 512 units each for encoding, and GNMT-style, 4-layer, 512-unit LSMT decoder with residual connections from the 2nd layer onward.

Training. We trained the above model using the *standard method* of Johnson et al. (2016) and using our proposed *agreement-based* training (Algorithm 1). In both cases, the model was optimized using Adafactor (Shazeer and Stern, 2018) on a machine with 4 P100 GPUs for up to 500K steps, with early stopping on the dev set.

Evaluation. We focus our evaluation mainly on zero-shot performance of the following methods: (a) **Basic**, which stands for directly evaluating a multilingual GNMT model after standard training (Johnson et al., 2016).

⁷<http://www.statmt.org/europarl/>

⁸<https://sites.google.com/site/iwsltevaluation2017/TED-tasks>

Table 1: Results on UNCorpus-1.

	Sestorain et al. (2018) [†]			Our baselines		
	PBSMT	NMT-0	Dual-0	Basic	Pivot	Agree
En → Es	61.26	51.93	—	56.58	56.58	56.36
En → Fr	50.09	40.56	—	44.27	44.27	44.80
Es → En	59.89	51.58	—	55.70	55.70	55.24
Fr → En	52.22	43.33	—	46.46	46.46	46.17
Supervised (avg.)	55.87	46.85	—	50.75	50.75	50.64
Es → Fr	52.44	20.29	36.68	34.75	38.10	37.54
Fr → Es	49.79	19.01	39.19	37.67	40.84	40.02
Zero-shot (avg.)	51.11	19.69	37.93	36.21	39.47	38.78

[†]Source: <https://openreview.net/forum?id=ByecAoAqK7>.

Table 2: Results on UNCorpus-2.

	Sestorain et al. (2018)			Our baselines		
	PBSMT	NMT-0	Dual-0	Basic	Pivot	Agree
En → Es	61.26	47.51	44.30	55.15	55.15	54.30
En → Fr	50.09	36.70	34.34	43.42	43.42	42.57
En → Ru	43.25	30.45	29.47	36.26	36.26	35.89
Es → En	59.89	48.56	45.55	54.35	54.35	54.33
Fr → En	52.22	40.75	37.75	45.55	45.55	45.87
Ru → En	52.59	39.35	37.96	45.52	45.52	44.67
Supervised (avg.)	53.22	40.55	36.74	46.71	46.71	46.27
Es → Fr	52.44	25.85	34.51	34.73	35.93	36.02
Fr → Es	49.79	22.68	37.71	38.20	39.51	39.94
Es → Ru	39.69	9.36	24.55	26.29	27.15	28.08
Ru → Es	49.61	26.26	33.23	33.43	37.17	35.01
Fr → Ru	36.48	9.35	22.76	23.88	24.99	25.13
Ru → Fr	43.37	22.43	26.49	28.52	30.06	29.53
Zero-shot (avg.)	45.23	26.26	29.88	30.84	32.47	32.29

Table 3: Zero-shot results on Europarl. Note that *Soft* and *Distill* are not multilingual systems.

	Previous work		Our baselines		
	Soft [‡]	Distill [†]	Basic	Pivot	Agree
Es → De	—	—	18.23	20.14	20.70
De → Es	—	—	20.28	26.50	22.45
Es → Fr	30.57	33.86	27.99	32.56	30.94
Fr → Es	—	—	27.12	32.96	29.91
De → Fr	23.79	27.03	21.36	25.67	24.45
Fr → De	—	—	18.57	19.86	19.15
Zero-shot (avg.)	—	—	22.25	26.28	24.60

[†]Soft pivoting (Cheng et al., 2017). [‡]Distillation (Chen et al., 2017).

(b) *Pivot*, which performs pivoting-based inference using a multilingual GNMT model (after standard training); often regarded as gold-standard.

(c) *Agree*, which applies a multilingual GNMT model trained with agreement losses directly to zero-shot directions.

To ensure a fair comparison in terms of model capacity, all the techniques above use the same multilingual GNMT architecture described in the previous section. All other results provided in the tables are as reported in the literature.

Implementation. All our methods were implemented using TensorFlow (Abadi et al., 2016) on top of tensorflow2tensor library (Vaswani et al., 2018). Our code will be made publicly available.⁹

⁹www.cs.cmu.edu/~mshediva/code/

Table 4: Results on the official IWSLT17 multilingual task.

	Previous work		Our baselines		
	SOTA [†]	CPG [‡]	Basic	Pivot	Agree
Supervised (avg.)	24.10	19.75	24.63	24.63	23.97
Zero-shot (avg.)	20.55	11.69	19.86	19.26	20.58

[†]Table 2 from Dabre et al. (2017). [‡]Table 2 from Platanios et al. (2018).

Table 5: Results on our proposed IWSLT17*.

	Basic	Pivot	Agree
Supervised (avg.)	28.72	28.72	29.17
Zero-shot (avg.)	12.61	17.68	15.23

6.3 Results on UN Corpus and Europarl

UN Corpus. Tables 1 and 2 show results on the UNCorpus datasets. Our approach consistently outperforms *Basic* and *Dual-0*, despite the latter being trained with additional monolingual data (Sestorain et al., 2018). We see that models trained with agreement perform comparably to *Pivot*, outperforming it in some cases, e.g., when the target is Russian, perhaps because it is quite different linguistically from the English pivot.

Furthermore, unlike *Dual-0*, *Agree* maintains high performance in the supervised directions (within 1 BLEU point compared to *Basic*), indicating that our agreement-based approach is effective as a part of a single multilingual system.

Europarl. Table 3 shows the results on the Europarl corpus. On this dataset, our approach consistently outperforms *Basic* by 2-3 BLEU points but lags a bit behind *Pivot* on average (except on Es → De where it is better). Cheng et al. (2017)¹⁰ and Chen et al. (2017) have reported zero-resource results on a subset of these directions and our approach outperforms the former but not the latter on these pairs. Note that both Cheng et al. (2017) and Chen et al. (2017) train separate models for each language pair and the approach of Chen et al. (2017) would require training $O(k^2)$ models to encompass all the pairs. In contrast, we use a single multilingual architecture which has more limited model capacity (although in theory, our approach is also compatible with using separate models for each direction).

6.4 Analysis of IWSLT17 zero-shot tasks

Table 4 presents results on the original IWSLT17 task. We note that because of the large amount of data overlap and presence of many supervised

¹⁰We only show their best zero-resource result in the table since some of their methods require direct parallel data.




Figure 4: BLEU on the dev set for Agree and the baselines trained on smaller subsets of the Europarl corpus.

translation pairs (16) the vanilla training method (Johnson et al., 2016) achieves very high zero shot performance, even outperforming Pivot. While our approach gives small gains over these baselines, we believe the dataset’s peculiarities make it not reliable for evaluating zero-shot generalization.

On the other hand, on our proposed preprocessed IWSLT17* that eliminates the overlap and reduces the number of supervised directions (8), there is a considerable gap between the supervised and zero-shot performance of Basic. Agree performs better than Basic and is slightly worse than Pivot.

6.5 Small data regime

To better understand the dynamics of different methods in the small data regime, we also trained all our methods on subsets of the Europarl for 200K steps and evaluated on the dev set. The training set size varied from 50 to 450K parallel sentences. From Figure 4, Basic tends to perform extremely poorly while Agree is the most robust (also in terms of variance across zero-shot directions). We see that Agree generally upper-bounds Pivot, except for the (Es, Fr) pair, perhaps due to fewer cascading errors along these directions.

7 Conclusion

In this work, we studied zero-shot generalization in the context of multilingual neural machine translation. First, we introduced the concept of zero-shot consistency that implies generalization. Next, we proposed a provably consistent agreement-based learning approach for zero-shot translation. Empirical results on three datasets showed that agreement-based learning results in up to +3 BLEU zero-shot

improvement over the Johnson et al. (2016) baseline, compares favorably to other approaches in the literature (Cheng et al., 2017; Sestorain et al., 2018), is competitive with pivoting, and does not lose in performance on supervised directions.

We believe that the theory and methodology behind agreement-based learning could be useful beyond translation, especially in multi-modal settings. For instance, it could be applied to tasks such as cross-lingual natural language inference (Conneau et al., 2018), style-transfer (Shen et al., 2017; Fu et al., 2017; Prabhumoye et al., 2018), or multilingual image or video captioning. Another interesting future direction would be to explore different hand-engineered or learned data representations, which one could use to encourage models to agree on during training (*e.g.*, make translation models agree on latent semantic parses, summaries, or potentially other data representations available at training time).

Acknowledgments

We thank Ian Tenney and Anthony Platanios for many insightful discussions, Emily Pitler for the helpful comments on the early draft of the paper, and anonymous reviewers for careful reading and useful feedback.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283.
- Rohan Anil, Gabriel Pereyra, Alexandre Passos, Robert Ormandi, George E Dahl, and Geoffrey E Hinton. 2018. Large scale distributed neural network training through online distillation. *arXiv preprint arXiv:1804.03235*.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roei Aharoni, Melvin Johnson, and Wolfgang Macherey. 2018. The missing ingredient in zero-shot neural machine translation. *OpenReview.net*.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2017. Unsupervised neural machine translation. *arXiv preprint arXiv:1710.11041*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. 2010. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175.
- Julian Besag. 1975. Statistical analysis of non-lattice data. *The statistician*, pages 179–195.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. Findings of the 2016 conference on machine translation. In *ACL 2016 FIRST CONFERENCE ON MACHINE TRANSLATION (WMT16)*, pages 131–198. The Association for Computational Linguistics.
- 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*.
- Yong Cheng, Qian Yang, Yang Liu, Maosong Sun, and Wei Xu. 2017. Joint training for pivot-based neural machine translation. In *Proceedings of IJCAI*.
- 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*.
- Trevor Cohn and Mirella Lapata. 2007. Machine translation by triangulation: Making effective use of multi-parallel corpora. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 728–735.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- Alexis Conneau, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. *arXiv preprint arXiv:1809.05053*.
- Josep Crego, Jungi Kim, Guillaume Klein, Anabel Rebollo, Kathy Yang, Jean Senellart, Egor Akhanov, Patrice Brunelle, Aurelien Coquard, Yongchao Deng, et al. 2016. Systran’s pure neural machine translation systems. *arXiv preprint arXiv:1610.05540*.
- Raj Dabre, Fabien Cromieres, and Sadao Kurohashi. 2017. Kyoto university mt system description for iwslt 2017. *Proc. of IWSLT, Tokyo, Japan*.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016a. Multi-way, multilingual neural machine translation with a shared attention mechanism. *arXiv preprint arXiv:1601.01073*.
- Orhan Firat, Baskaran Sankaran, Yaser Al-Onaizan, Fatos T Yarman Vural, and Kyunghyun Cho. 2016b. Zero-resource translation with multilingual neural machine translation. *arXiv preprint arXiv:1606.04164*.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2017. Style transfer in text: Exploration and evaluation. *arXiv preprint arXiv:1711.06861*.
- Jonas Gehring, Michael Auli, David Grangier, and Yann N Dauphin. 2016. A convolutional encoder model for neural machine translation. *arXiv preprint arXiv:1611.02344*.
- Thanh-Le Ha, Jan Niehues, and Alexander Waibel. 2017. Effective strategies in zero-shot neural machine translation. *arXiv preprint arXiv:1711.07893*.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tieyan Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *Advances in Neural Information Processing Systems*, pages 820–828.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2014. Dark knowledge. *Presented as the keynote in BayLearn, 2*.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2016. Google’s multilingual neural machine translation system: enabling zero-shot translation. *arXiv preprint arXiv:1611.04558*.
- Marcin Junczys-Dowmunt, Tomasz Dwojak, and Hieu Hoang. 2016. Is neural machine translation ready for deployment? a case study on 30 translation directions. *arXiv preprint arXiv:1610.01108*.
- Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*.
- Philip Koehn. 2017. Europarl: A parallel corpus for statistical machine translation.
- Philipp Koehn. 2009. *Statistical machine translation*. Cambridge University Press.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. *arXiv preprint arXiv:1706.03872*.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. Phrase-based & neural unsupervised machine translation. *arXiv preprint arXiv:1804.07755*.

- Percy Liang, Ben Taskar, and Dan Klein. 2006. Alignment by agreement. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 104–111. Association for Computational Linguistics.
- Percy S Liang, Dan Klein, and Michael I Jordan. 2008. Agreement-based learning. In *Advances in Neural Information Processing Systems*, pages 913–920.
- Bruce G Lindsay. 1988. Composite likelihood methods. *Contemporary mathematics*, 80(1):221–239.
- Yichao Lu, Phillip Keung, Faisal Ladhak, Vikas Bhardwaj, Shaonan Zhang, and Jason Sun. 2018. A neural interlingua for multilingual machine translation. *arXiv preprint arXiv:1804.08198*.
- Minh-Thang Luong, Quoc V Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2015a. Multi-task sequence to sequence learning. *arXiv preprint arXiv:1511.06114*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015b. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Cettolo Mauro, Federico Marcello, Bentivogli Luisa, Niehues Jan, Stüker Sebastian, Sudoh Katsutho, Yoshino Koichiro, and Federmann Christian. 2017. Overview of the iwslt 2017 evaluation campaign. In *International Workshop on Spoken Language Translation*, pages 2–14.
- 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.
- Emmanouil Antonios Platanios, Mrinmaya Sachan, Graham Neubig, and Tom Mitchell. 2018. Contextual parameter generation for universal neural machine translation. *arXiv preprint arXiv:1808.08493*.
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. Style transfer through back-translation. *arXiv preprint arXiv:1804.09000*.
- Sujith Ravi and Kevin Knight. 2011. Deciphering foreign language. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 12–21. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015a. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015b. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Lierni Sestorain, Massimiliano Ciaramita, Christian Buck, and Thomas Hofmann. 2018. Zero-shot dual machine translation. *arXiv preprint arXiv:1805.10338*.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. *arXiv preprint arXiv:1804.04235*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in Neural Information Processing Systems*, pages 6830–6841.
- 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.
- 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.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan N Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, et al. 2018. Tensor2tensor for neural machine translation. *arXiv preprint arXiv:1803.07416*.
- 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.
- Hua Wu and Haifeng Wang. 2007. Pivot language approach for phrase-based statistical machine translation. *Machine Translation*, 21(3):165–181.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Michał Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. The united nations parallel corpus v1. 0. In *LREC*.

A Appendices

A.1 Complete likelihood

Given a set of conditional models, $\{\mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)\}$, we can write out the full likelihood over equivalent translations, $(\mathbf{x}_1, \dots, \mathbf{x}_k)$, as follows:

$$\mathbb{P}_\theta(\mathbf{x}_1, \dots, \mathbf{x}_k) := \frac{1}{Z} \prod_{i,j \in \mathcal{E}} \mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i) \quad (11)$$

where $Z := \sum_{\mathbf{x}_1, \dots, \mathbf{x}_k} \prod_{i,j \in \mathcal{E}} \mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)$ is the normalizing constant and $\mathcal{E}$ denotes *all* edges in the graph (Figure 5). Given only bilingual parallel corpora, $\mathcal{C}_{ij}$ for $i, j \in \mathcal{E}_s$, we can observe only certain pairs of variables. Therefore, the log-likelihood of the data can be written as:

$$\begin{aligned} \mathcal{L}(\theta) := & \sum_{i,j \in \mathcal{E}_s} \sum_{\mathbf{x}_i, \mathbf{x}_j \in \mathcal{C}_{ij}} \log \sum_{\mathbf{z}} \mathbb{P}_\theta(\mathbf{x}_1, \dots, \mathbf{x}_k) \end{aligned} \quad (12)$$

Here, the outer sum iterates over available corpora. The middle sum iterates over parallel sentences in a corpus. The most inner sum marginalizes out unobservable sequences, denoted $\mathbf{z} := \{\mathbf{x}_l\}_{l \neq i,j}$, which are sentences equivalent under this model to $\mathbf{x}_i$ and $\mathbf{x}_j$ in languages other than L_i and L_j . Note that due to the inner-most summation, computing the log-likelihood is intractable.

We claim the following.

Claim 4 *Maximizing the full log-likelihood yields zero-shot consistent models (Definition 1).*

Proof. To better understand why this is the case, let us consider example in Figure 5 and compute the log-likelihood of $(\mathbf{x}_1, \mathbf{x}_2)$:

$$\begin{aligned} \log \mathbb{P}_\theta(\mathbf{x}_1, \mathbf{x}_2) &= \log \sum_{\mathbf{x}_3, \mathbf{x}_4} \mathbb{P}_\theta(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4) \\ &\propto \log \mathbb{P}_\theta(\mathbf{x}_1 | \mathbf{x}_2) + \log \mathbb{P}_\theta(\mathbf{x}_2 | \mathbf{x}_1) + \\ &\quad \log \sum_{\mathbf{x}_3, \mathbf{x}_4} \mathbb{P}_\theta(\mathbf{x}_1 | \mathbf{x}_3) \mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_1) \times \\ &\quad \mathbb{P}_\theta(\mathbf{x}_2 | \mathbf{x}_3) \mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_2) \times \\ &\quad \mathbb{P}_\theta(\mathbf{x}_1 | \mathbf{x}_4) \mathbb{P}_\theta(\mathbf{x}_4 | \mathbf{x}_1) \times \\ &\quad \mathbb{P}_\theta(\mathbf{x}_2 | \mathbf{x}_4) \mathbb{P}_\theta(\mathbf{x}_4 | \mathbf{x}_2) \times \\ &\quad \mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_4) \mathbb{P}_\theta(\mathbf{x}_4 | \mathbf{x}_3) \end{aligned}$$

Note that the terms that encourage agreement on the translation into L_3 are colored in **green** (similarly, terms that encourage agreement on the translation into L_4 are colored in **blue**). Since all other terms are probabilities and bounded by 1, we have:




Figure 5: Probabilistic graphical model for a multilingual system with four languages (L_1, L_2, L_3, L_4). Variables can only be observed only in pairs (shaded in the graph).

$$\begin{aligned} &\log \mathbb{P}_\theta(\mathbf{x}_1, \mathbf{x}_2) + \log Z \\ &\leq \log \mathbb{P}_\theta(\mathbf{x}_1 | \mathbf{x}_2) + \log \mathbb{P}_\theta(\mathbf{x}_2 | \mathbf{x}_1) + \\ &\quad \log \sum_{\mathbf{x}_3, \mathbf{x}_4} \mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_1) \mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_2) \times \\ &\quad \mathbb{P}_\theta(\mathbf{x}_4 | \mathbf{x}_1) \mathbb{P}_\theta(\mathbf{x}_4 | \mathbf{x}_2) \\ &\equiv \mathcal{L}^{\text{agree}}(\theta) \end{aligned}$$

In other words, the full log likelihood lower-bounds the agreement objective (up to a constant $\log Z$). Since optimizing for agreement leads to consistency (Theorem 2), and maximizing the full likelihood would necessarily improve the agreement, the claim follows. ■

Remark 5 *Note that the other terms in the full likelihood also have a non-trivial purpose: (a) the terms $\mathbb{P}_\theta(\mathbf{x}_1 | \mathbf{x}_3)$, $\mathbb{P}_\theta(\mathbf{x}_1 | \mathbf{x}_4)$, $\mathbb{P}_\theta(\mathbf{x}_2 | \mathbf{x}_3)$, $\mathbb{P}_\theta(\mathbf{x}_2 | \mathbf{x}_4)$, encourage the model to correctly re-construct $\mathbf{x}_1$ and $\mathbf{x}_2$ when back-translating from unobserved languages, L_3 and L_4 , and (b) terms $\mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_4)$, $\mathbb{P}_\theta(\mathbf{x}_4 | \mathbf{x}_3)$ enforce consistency between the latent representations. In other words, full likelihood accounts for a combination of agreement, back-translation, and latent consistency.*

A.2 Proof of agreement consistency

The statement of Theorem 2 mentions an assumption on the true distribution of the equivalent translations. The assumption is as follows.

Assumption 6 *Let $\mathbb{P}(\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k)$ be the ground truth conditional distribution that specifies the probability of $\mathbf{x}_i$ to be a translation of $\mathbf{x}_j$ and $\mathbf{x}_k$ into language L_i , given that $(\mathbf{x}_j, \mathbf{x}_k)$ are correct translations of each other in languages L_j and L_k , respectively. We assume:*

$$0 \leq \delta \leq \mathbb{E}_{\mathbf{x}_k | \mathbf{x}_i, \mathbf{x}_j} [\mathbb{P}(\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k)] \leq \xi \leq 1$$

This assumption means that, even though there might be multiple equivalent translations, there must be not too many of them (implied by the δ lower bound) and none of them must be much more preferable than the rest (implied by the ξ upper bound). Given this assumption, we can prove the following simple lemma.

Lemma 7 *Let $L_i \rightarrow L_j$ be one of the supervised directions, $\mathbb{E}_{\mathbf{x}_i, \mathbf{x}_j} [-\log \mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)] \leq \varepsilon$. Then the following holds:*

$$\mathbb{E}_{\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k} \left[\frac{\mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)}{\mathbb{P}(\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k)} \right] \geq \log \frac{1}{\xi} - \varepsilon \delta$$

Proof. First, using Jensen’s inequality, we have:

$$\begin{aligned} \log \mathbb{E}_{\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k} \left[\frac{\mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)}{\mathbb{P}(\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k)} \right] &\geq \\ \mathbb{E}_{\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k} [\log \mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i) - \log \mathbb{P}(\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k)] &\end{aligned}$$

The bound on the supervised direction implies that

$$\mathbb{E}_{\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k} [-\log \mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)] \geq -\varepsilon \delta$$

To bound the second term, we use Assumption 6:

$$\mathbb{E}_{\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k} [-\log \mathbb{P}(\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k)] \geq \log \frac{1}{\xi}$$

Putting these together yields the bound. $\blacksquare$

Now, using Lemma 7, we can prove Theorem 2.

Proof. By assumption, the agreement-based loss is bounded by ε . Therefore, expected cross-entropy on all supervised terms, $L_1 \leftrightarrow L_2$, is bounded by ε . Moreover, the agreement term (which is part of the objective) is also bounded:

$$-\mathbb{E}_{\mathbf{x}_i, \mathbf{x}_j} \left[\sum_{\mathbf{x}_k} \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j) \log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_i) \right] \leq \varepsilon$$

Expanding this expectation, we have:

$$\begin{aligned} &\sum_{\mathbf{x}_i, \mathbf{x}_j} \mathbb{P}(\mathbf{x}_i, \mathbf{x}_j) \sum_{\mathbf{x}_k} \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j) \log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_i) \\ &= \sum_{\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k} \mathbb{P}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) \times \\ &\quad \frac{\mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j)}{\mathbb{P}(\mathbf{x}_k | \mathbf{x}_i, \mathbf{x}_j)} \log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_i) \\ &= \sum_{\mathbf{x}_i, \mathbf{x}_k} \mathbb{E}_{\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k} \left[\frac{\mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j)}{\mathbb{P}(\mathbf{x}_k | \mathbf{x}_i, \mathbf{x}_j)} \right] \times \\ &\quad \mathbb{P}(\mathbf{x}_i, \mathbf{x}_k) \log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_i) \end{aligned}$$

Combining that with Lemma 7, we have:

$$\mathbb{E}_{\mathbf{x}_i, \mathbf{x}_k} [-\log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_i)] \leq \frac{\varepsilon}{\log \frac{1}{\xi} - \delta \varepsilon} \equiv \kappa(\varepsilon)$$

Since by Assumption 6, δ and ξ are some constants, $\kappa(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$. $\blacksquare$

A.3 Consistency of distillation and pivoting

As we mentioned in the main text of the paper, distillation (Chen et al., 2017) and pivoting yield zero-shot consistent models. Let us understand why this is the case.

In our notation, given $L_1 \rightarrow L_2$ and $L_2 \rightarrow L_3$ as supervised directions, distillation optimizes a KL-divergence between $\mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_2)$ and $\mathbb{P}_\theta(\mathbf{x}_3 | \mathbf{x}_1)$, where the latter is a zero-shot model and the former is supervised. Noting that KL-divergence lower-bounds cross-entropy, it is a looser bound on the agreement loss. Hence, by ensuring that KL is low, we also ensure that the models agree, which implies consistency (a more formal proof would exactly follow the same steps as the proof of Theorem 2).

To prove consistency of pivoting, we need an additional assumption on the quality of the source-pivot model.

Assumption 8 *Let $\mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)$ be the source-pivot model. We assume the following bound holds for each pair of equivalent translations, $(\mathbf{x}_j, \mathbf{x}_k)$:*

$$\mathbb{E}_{\mathbf{x}_i | \mathbf{x}_j, \mathbf{x}_k} \left[\frac{\mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i)}{\mathbb{P}(\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k)} \right] \leq C$$

where $C > 0$ is some constant.

Theorem 9 (Pivoting consistency) *Given the conditions of Theorem 2 and Assumption 8, pivoting is zero-shot consistent.*

Proof. We can bound the expected error on pivoting as follows (using Jensen’s inequality and the conditions from our assumptions):

$$\begin{aligned} &\mathbb{E}_{\mathbf{x}_i, \mathbf{x}_k} \left[-\log \sum_{\mathbf{x}_j} \mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i) \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j) \right] \\ &\leq \mathbb{E}_{\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k} [-\mathbb{P}_\theta(\mathbf{x}_j | \mathbf{x}_i) \log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j)] \\ &\leq \sum_{\mathbf{x}_i, \mathbf{x}_k} \mathbb{E}_{\mathbf{x}_j | \mathbf{x}_i, \mathbf{x}_k} \left[\frac{\mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_j)}{\mathbb{P}(\mathbf{x}_k | \mathbf{x}_i, \mathbf{x}_j)} \right] \times \\ &\quad \mathbb{P}(\mathbf{x}_i, \mathbf{x}_k) \log \mathbb{P}_\theta(\mathbf{x}_k | \mathbf{x}_i) \\ &\leq C \varepsilon \end{aligned}$$

$\blacksquare$

A.4 Details on the models and training

Architecture. All our NMT models used the GNMT (Wu et al., 2016) architecture with Luong attention (Luong et al., 2015b), 2 bidirectional encoder, and 4-layer decoder with residual connections. All hidden layers (including embeddings) had 512 units. Additionally, we used separate embeddings on the encoder and decoder sides as well as tied weights of the softmax that produced logits with the decoder-side (*i.e.*, target) embeddings. Standard dropout of 0.2 was used on all hidden layers. Most of the other hyperparameters we set to default in the T2T (Vaswani et al., 2018) library for the text2text type of problems.

Training and hyperparameters. We scaled agreement terms in the loss by $\gamma = 0.01$. The training was done using Adafactor (Shazeer and Stern, 2018) optimizer with 10,000 burn-in steps at 0.01 learning rate and further standard square root decay (with the default settings for the decay from the T2T library). Additionally, implemented agreement loss as a subgraph as a loss was not computed if γ was set to 0. This allowed us to start training multilingual NMT models in the burn-in mode using the composite likelihood objective and then switch on agreement starting some point during optimization (typically, after the first 100K iterations; we also experimented with 0, 50K, 200K, but did not notice any difference in terms of final performance). Since the agreement subgraph was not computed during the initial training phase, it tended to accelerate training of agreement models.