

ISOMETRIC MT: NEURAL MACHINE TRANSLATION FOR AUTOMATIC DUBBING

Surafel M. Lakew, Yogesh Virkar, Prashant Mathur, Marcello Federico

Amazon AI

{surafelm|yvvirkar|pramathu|marcfeede}@amazon.com

ABSTRACT

Automatic dubbing (AD) is among the machine translation (MT) use cases where translations should match a given length to allow for synchronicity between source and target speech. For neural MT, generating translations of length close to the source length (e.g. within $\pm 10\%$ in character count), while preserving quality is a challenging task. Controlling MT output length comes at a cost to translation quality, which is usually mitigated with a two step approach of generating N -best hypotheses and then re-ranking based on length and quality. This work introduces a self-learning approach that allows a transformer model to directly learn to generate outputs that closely match the source length, in short *Isometric MT*. In particular, our approach does not require to generate multiple hypotheses nor any auxiliary ranking function. We report results on four language pairs (*English* \rightarrow *French, Italian, German, Spanish*) with a publicly available benchmark. Automatic and manual evaluations show that our method for Isometric MT outperforms more complex approaches proposed in the literature.

Index Terms— Machine Translation, Isometric Translation, Automatic Dubbing

1. INTRODUCTION

Reaching the global audience is a primary factor for audio-visual content localization. Automating the task of localization requires translation of source language speech and a seamless integration of the target language speech with the original visual content [1]. Recent developments in AD [2, 1, 3] have focused on achieving *isochrony*, a form of synchronization at the level of speech utterances [4]. The AD architecture proposed in [1], includes an MT model that translates source transcriptions to a target language, followed by a prosodic alignment (PA) module which segments the translation into phrases and pauses according to the prosody pattern in the source speech. Finally, a text-to-speech (TTS) module synthesizes the target language speech for a final rendering with the original audio background and visual content.

For a TTS module, to generate a natural sounding speech in synchrony with the source utterances, it is important that the length of the translated script should match that of the source.¹ With a goal of generating translations that match the source length, [5, 6] have proposed approaches that bias the output length of MT. In a subsequent work [7], the authors proposed a better approach, where the main idea is to generate translations that fall within a $\pm 10\%$ range of the source length in character count. They also confirmed that generating translations within this range easily allow the PA and the TTS modules to adjust the target speech to a natural sounding speaking rate. [7] is state of the art approach for *MT with length control*,

¹Following [1], length of input text in characters is directly proportional to the duration of the TTS utterance.

but it depends on a two step process *i*) N -best hypotheses (where $N=50$) generation and *ii*) a re-ranking step which interpolates the model scores and length ratios (translation to source) to find the best hypothesis in the N -best space. In this work, we propose a self-learning based approach whose improvements are orthogonal to the re-ranker’s improvements and when combined we achieve state of the art results on a public benchmark.

In statistical MT, self-learning [8] has been investigated to augment training data with pseudo bi-text. In neural MT, the most commonly used self-learning approach is back-translation [9], that leverages a reverse (*target* \rightarrow *source*) model to generate pseudo bi-text from a target language monolingual data, to train the desired *source* \rightarrow *target* direction. Subsequent works [10, 11, 12], have shown variants of self-learning using a backward and forward translation settings.

This work proposes a self-learning approach that applies a controlled generation of pseudo bi-text with a length constraint, to model isometric MT. To assess the effectiveness of the proposed approach, we perform evaluation on a speech translation data [13], in four language directions exhibiting different degree of target to source length ratio: *English* \rightarrow *French, Italian, German, Spanish*. Specifically, our contributions are:

- We propose a self-learning based approaches to learn MT model that can generate isometric translations.
- We compare our proposed approach and show it is on par or better than previous state of the art for controlling MT output length, without using multiple hypothesis generation and re-ranking function.
- We introduce a new MT evaluation mechanism leveraging TTS (*i.e. generating audio from translations for rating by subjects*), and metrics to measure the acceptability rate of isometric translations, particularly for an AD use case.

The rest of this work, discusses MT and output length control approaches §2, followed by a description of our proposed isometric MT §3, experiment and evaluation settings §4, and finally results and discussion of our findings §5.

2. BACKGROUND

2.1. Neural Machine Translation

For a language pair with parallel data $\mathcal{D} = \{(s_i, t_i) : i = 1, \dots, N\}$, an MT model parameterized with θ , trains to maximize likelihood on the training sample pairs as,

$$L(\theta) = \arg \max_{\theta} \sum_{i=1}^N \log p(t_i | s_i, \theta) \quad (1)$$

The model considers the source context (s) and previously generated target language tokens ($t_{<k}$), when forming the prediction (t),




Fig. 1. MT output length control approaches; (*left*) current SOTA using length token+ N -best hypotheses re-ranking [7], (*center*) our proposed offline self-learning using a reverse MT (*bottom*) trained in isolation to improve the desired (*top*) forward MT direction, and (*right*) the online self-learning variant using a single bi-directional MT. *Broken* arrows show test time hypotheses generation (*left*), and self-learning training data generation (*center*; *right*). For simplicity, we show embedding lookup function $e()$ for input tokens, where $e(v)$ for length token, $e(l)$ for language token, and $e(vl)$ for combination of length and language token.

$$p(t|s) = \prod_{k=1}^{K+1} p(t_k | t_{<k}, s, \theta) \quad (2)$$

$K+1$ is the length of t with reserved tokens $\langle \text{bos} \rangle$ at $k=0$ and $\langle \text{eos} \rangle$ at $k=K+1$ indicating the beginning and end of the prediction.

2.2. Towards Controlling Output Length of MT

Prior to recent studies on MT output length control [6, 14, 7], several attempts have been made for controlling output attributes: text length in summarization [15, 16], text complexity and formality [17, 18, 19], domain information [20], and languages in multilingual MT [21]. Specific to MT output length control, [6, 14] proposed variants of length-encoding motivated by the positional-encoding of self-attention [22]. In closely related works, [23, 24] proposed approaches to respectively control verbosity of speech transcription and speech translation for subtitling. We will focus on approaches that are particularly suited for the AD use case.

2.2.1. Length Token Embedding

The MT length control approach in [6] first proposed to *classify* the bi-text \mathcal{D} into three classes, based on the target to source character count length ratio (LR) of each sample (s, t) pair. The class labels are defined on LR thresholds: $short < 0.95 \leq normal \leq 1.05 > long$. Then model training is performed by prepending the length token $v \in \{short, normal, long\}$, at the beginning of the source sentence. At time of inference, the desired v is prepended on the input sentence.²

In comparison with a baseline length penalty approach to generate shorter translations [25], and the length-encoding variants of [6, 14], length token has been shown to generate more suitable translations for AD [7].

2.2.2. Hypotheses Re-Ranking

To maximize the suitability of translations for AD, [5] proposed MT hypotheses re-ranking. The approach works by first generating N -best hypotheses (t) for a source s , followed by a re-ranking step based on a scoring function that combines the likelihood of each hypothesis with a so called synchrony score (S_p), as follows:

$$S_d(t, s) = (1 - \alpha) \log P(t | s) + \alpha S_p(t, s), \quad (3)$$

²Both [6] and [7] stated $v = normal$ as the best setting for generating more suitable sentence for AD.

Where S_p is computed as $S_p(t, s) = (1 + |\text{len}(t) - \text{len}(s)|)^{-1}$, α is a parameter tuned on a validation set.

2.2.3. Length Token + Re-Ranking

Recently, [7] proposed to combine the length token of [6] and a variant of Eq. 3 by modifying the synchrony sub-score as, $S_p(t, s) = (1 + \frac{\text{len}(t)}{\text{len}(s)})^{-1}$, which aligns with the overall objective of reducing target to source LR in MT output length control. Fig. 1 (*left*), illustrates the approach of [7] which reported a higher % of dubbing suitable translations and human preference of dubbed videos.

3. ISOMETRIC MT WITH SELF-LEARNING

In this section, we describe our self-learning based output length control for MT. To show the impact of our approach, we take AD architecture proposed in [1].

3.1. Self Learning for Isometric MT

In MT, the main idea of self-learning is to learn a better model using the predictions of the model itself or an auxiliary model output. Self-learning has shown impressive results in unsupervised [26], zero-shot [27, 28], and semi-supervised [9, 10] neural MT modeling. Despite the success of self-learning in MT, existing approaches rely on using external monolingual data and do not evaluate the impact of the pseudo generated data quality before incorporating it into the training stage.

In this work, the self-learning approach is different in the following aspects, *i*) we propose a new synthetic data generation technique based on an output length control criteria, *ii*) we then filter and classify the pseudo data based on predefined length classes. In the sections below, we discuss in detail two type of self-learning approaches.

3.2. Offline Self-Learning

As shown in Fig. 1 (*center*), the offline self-learning approach considers two independent MT models. The first one is a reverse *target* \rightarrow *source* model (MT_R) trained with the length token approach [7], optimizing the objective in Eq. 1. The second model is the desired *source* \rightarrow *target* forward model (MT_F) which is used to generate isometric translations.

Specifically, for parallel training data \mathcal{D} we implement the offline self-learning by first *generating* pseudo source s' from the target t , using MT_R . For inference we prepend $v = long$ on t with the goal of generating more samples with a lower LR . Ultimately,

we aim to address the imbalance of the three length classes in the original bi-text \mathcal{D} .

After performing inference for all t , we then *construct* the pseudo bi-text $\mathcal{D}' = \{(s'_i, t_i): i = 1, \dots, M\}$, with length class label for each sample following the LR and thresholds defined in [7]. We then concatenate the original \mathcal{D} and the pseudo \mathcal{D}' bi-text to fine-tune the desired MT_F model by optimizing Eq. 1.

3.3. Online Self-Learning

Our second approach primarily avoids learning MT_R in isolation, as shown in Fig. 1 (*right*). We implement a bi-directional model with length classes in each direction ($v \in \{\text{short}, \text{normal}, \text{long}\}$). We prepend a token (l) indicating the *source language* on the encoder, and *target language+length* (lv) token on the decoder side. For model training, the objective in Eq. 1 is formalized as,

$$L(\theta) = \arg \max_{\theta} \log p(t|s, \theta) + \arg \max_{\theta} \log p(s|t, \theta) \quad (4)$$

During training, we follow a similar three step procedure of the offline self-learning: *length controlled inference, classification of pseudo bi-text into classes, and model training*. Differently from the offline, the steps are executed on-the-fly at time of model training and the learning procedure is applied both for the reverse and the forward directions.³ We hypothesize modeling a bi-directional isometric MT not only able to perform what normally requires two separate MT models, it also positively re-enforce output length control in both directions, with the addition of the pseudo \mathcal{D}' bi-text to balance the length classes in \mathcal{D} .

For self-learning approaches, based on how the original (\mathcal{D}) and pseudo bi-text (\mathcal{D}') is merged, we have two data configurations, *i) Union* = $\mathcal{D} \cup \mathcal{D}'$, and *ii) Filter* = $\mathcal{D} \cup \mathcal{D}'$ if $\text{LR}(t/s') \leq 1.05$. In other word, for *Filter* we remove the pseudo bi-text data portion labeled $v = \text{long}$.

4. EXPERIMENTS

4.1. Language and Data Processing

We conduct experiments for English (En) \rightarrow Italian (It), French (Fr), German (De), Spanish (Es) pairs. To avoid the effect of training with low-resource data, we pre-train strong MT models on internal data in the magnitude of 10^7 samples. We then apply our approaches as a fine-tuning step using public data of Ted Talks speeches of $\approx 200k$ samples per pair from MuSTC corpus [13]. Across all the approaches compared we use the same pre-training and fine-tuning data configuration. Moses scripts are used to pre-process the raw data, followed by token segmentation using SentencePiece with 32k merge operations.⁴

4.2. Model and Training Configuration

We use a transformer architecture [22] with 6 encoder and decoder layers, self-attention dimension 1024 and 16 heads, and feed-forward sublayers of dimension 4096. Model is optimized with Adam [29], with an initial learning rate of 1×10^{-7} . Dropout

³Although, we can evaluate online isometric MT both for *source* \leftrightarrow *target* directions, for a fair comparison with previous work and the offline setting we focus on the forward direction, and leave the rest for future work.

⁴Moses: <https://github.com/moses-smt/mosesdecoder>, SentencePiece: <https://github.com/google/sentencepiece>

is uniformly set to 0.1 for all model training. For the online self-learning we share all model parameters between the forward and reverse directions. For model pre-training the best checkpoint is selected based on the lowest loss on the validation set after the training converges. For fine-tuning we train for 5 epochs and take the best checkpoint for evaluation. At inference time we use beam size of 5, except for the re-ranking approach [7] where beam size = 50.

4.3. Model Evaluation

To evaluate both self-learning approaches, we use the human annotated benchmark from [30, 7]. The test set includes 620 samples per language pair, where references are post-edited to match the source length.

4.3.1. Automatic Metrics

We report Sacrebleu scores [31] on the de-tokenized translated segments and reference pair, target to source length ratio (LR) in character count, and the % of length compliant (LC) translations satisfying the $\pm 10\%$ range. We also report a single score that takes into account both translation quality and length compliance, by simply multiplying the BLEU with LC by which we name LCB .

4.3.2. Human Evaluation

To align MT human evaluation with AD, we devise an audio based MT evaluation scheme. We conduct the evaluation by generating audio for a translated segment and its reference using a TTS module. For a fair comparison, a uniform speaking rate is used to generate audio. We then ask annotators to listen to the audio of the translation and the reference and classify them in three categories:

| | |
|---------------------------|--------------------------------------|
| <i>Acceptable</i> (A) | Meaning is similar, fluency is good. |
| <i>Fixable</i> (F) | Meaning is similar, fluency is poor. |
| <i>Wrong</i> (W) | Meaning is different. |

To assess overall MT quality the % of A , F , and W are computed, whereas to identify length compliant translation we report LC for each rating category. As a final MT human evaluation score we compute,

$$HE_{MT} = LC(A) + \frac{1}{2}LC(F) \quad (5)$$

HE_{MT} disregards the wrong translations and considers the % of A and F translations that are isometric. We assume that *acceptable* translations will not require any post-editing, while *fixable* will require minimal post-editing to fix the fluency as such we weigh it with 0.5.

5. RESULT AND DISCUSSION

We compare proposed approaches against a strong *Baseline*, and two MT length control approaches ($v\text{Tok}$ [6] and $v\text{Tok}+\text{Rank}$ [7]). As described in Sec. 2.2, $v\text{Tok}+\text{Rank}$ combines length token ($v\text{Tok}$) and hypotheses re-ranking to generate a higher % of translations suitable for AD [7].

5.1. Automatic Evaluation

MT quality and verbosity control are measured using the automatic metrics in Sec. 4.3. We observe that length ratio of the Baseline

| Pair | Method | SL Data | BLEU | LR | LC | LCB |
|------------------|------------------|---------------|-------------|-------------|------|-------------|
| En-Fr | Baseline | - | 46.1 | 1.16 | 33.1 | 15.2 |
| | <i>vTok</i> | - | 48.4 | 1.06 | 72.7 | 35.2 |
| | <i>vTok+Rank</i> | - | 47.7 | 1.01 | 91.5 | 43.6 |
| | offlineSL | <i>Union</i> | 47.9 | 1.05 | 81.0 | 38.8 |
| | | <i>Filter</i> | 47.3 | 1.03 | 83.4 | 39.4 |
| | onlineSL | <i>Union</i> | 48.3 | 1.04 | 88.9 | 42.9 |
| | | <i>Filter</i> | 48.4 | 1.04 | 88.5 | 42.9 |
| | En-It | Baseline | - | 37.1 | 1.06 | 56.5 |
| <i>vTok</i> | | - | 37.5 | 1.05 | 74.7 | 28.0 |
| <i>vTok+Rank</i> | | - | 38.5 | 1.00 | 90.0 | 34.7 |
| offlineSL | | <i>Union</i> | 38.0 | 1.04 | 80.0 | 30.4 |
| | | <i>Filter</i> | 38.1 | 1.03 | 81.6 | 31.1 |
| onlineSL | | <i>Union</i> | 36.2 | 1.04 | 85.2 | 30.8 |
| | | <i>Filter</i> | 36.5 | 1.03 | 84.7 | 30.9 |
| En-De | | Baseline | - | 33.2 | 1.15 | 29.8 |
| | <i>vTok</i> | - | 34.3 | 1.05 | 79.2 | 27.2 |
| | <i>vTok+Rank</i> | - | 33.9 | 1.00 | 93.1 | 31.5 |
| | offlineSL | <i>Union</i> | 34.3 | 1.01 | 86.5 | 29.7 |
| | | <i>Filter</i> | 34.7 | 1.03 | 92.4 | 32.1 |
| | onlineSL | <i>Union</i> | 34.0 | 1.03 | 90.8 | 30.9 |
| | | <i>Filter</i> | 34.1 | 1.03 | 90.5 | 30.9 |
| | En-Es | Baseline | - | 50.1 | 1.07 | 56.0 |
| <i>vTok</i> | | - | 49.9 | 1.04 | 87.6 | 43.7 |
| <i>vTok+Rank</i> | | - | 49.3 | 1.00 | 95.5 | 47.1 |
| offlineSL | | <i>Union</i> | 50.0 | 1.03 | 90.8 | 45.4 |
| | | <i>Filter</i> | 49.9 | 1.03 | 93.1 | 46.4 |
| onlineSL | | <i>Union</i> | 49.9 | 1.03 | 95.6 | 47.7 |
| | | <i>Filter</i> | 49.8 | 1.04 | 93.9 | 46.7 |

Table 1. Results of self-learning isometric MT approaches (*offlineSL*, *onlineSL*) in comparison with previously proposed MT output length control mechanisms (*vTok* [6], and *vTok+Rank* [7] with N-best hypotheses re-ranking). We measure translation quality (BLEU), length ratio (*LR*), % of length compliant (*LC*), and overall metric *LCB*. Metric in bold shows best performing approach.

models for En-Fr and En-De is significantly higher than that of En-It and En-Es. This is due to the fact that original MuST-C training data is already well balanced for En-It and En-Es in the ratio of 30/35/35 on an average for short/normal/long classes whereas the same ratio is 15/31/54 for En-Fr and En-De. Given that *vTok+Rank* optimizes over N-best candidate translations, it shows better length control over the Baseline and *vTok* approach.

Self-learning approaches (*offlineSL* and *onlineSL*), provides consistent improvements over the baseline and *vTok* in terms of *LCB* metric for both data configuration (*Union*, *Filter*). We attribute the performance gains to training models using the additional pseudo parallel data. In fact, performance of self-learning approaches are similar to *vTok+Rank* for En-De and En-Es which shows how effective controlled self-learning could be, without a hypotheses re-ranking module as in *vTok+Rank*. Although, there’s no clear winning system in terms of *LCB*, we prefer *onlineSL* with *Filter* simply because it has a higher length compliance (*LC*) over *offlineSL* system and *Filter* has higher *LCB* on an average when compared to *Union* approach.

5.2. Human Evaluation

Following the evaluation criteria in Sec. 4.3, Table 2 show results of MT human evaluation. We use 100 randomly selected samples of the test set graded by 40 subjects per language pair and head to head comparison. We compare our proposed approaches against *vTok*

| Pair | <i>vTok</i> | offlineSL | onlineSL | <i>vTok</i> | <i>vTok+Rank</i> |
|-------|-------------|-----------|-------------|-------------|------------------|
| En-Fr | 55.1 | 45.8 | 52.0 | 60.9 | 67.9 |
| En-It | 58.0 | 60.9 | 66.1 | 66.6 | 73.0 |
| En-De | 67.6 | 66.5 | 74.0 | 65.8 | 75.9 |
| En-Es | 69.5 | 67.4 | 72.9 | 66.9 | 70.9 |

Table 2. Results of head to head MT human evaluation for, *vTok* Vs. *offlineSL* Vs. *onlineSL*, and *vTok* Vs. *vTok+Rank*. HE_{MT} score is computed using Eq. 5 that considers acceptable and fixable length compliant translations.

| Pair | Method | BLEU | LC | LCB |
|-------|------------------|------|------|-------------|
| En-Fr | <i>vTok+Rank</i> | 47.7 | 91.5 | 43.6 |
| | onlineSL | 48.4 | 88.5 | 42.9 |
| | +Rank | 48.4 | 93.0 | 45.0 |
| En-It | <i>vTok+Rank</i> | 38.5 | 90.0 | 34.7 |
| | onlineSL | 36.5 | 84.7 | 30.9 |
| | +Rank | 37.0 | 91.1 | 33.7 |
| En-De | <i>vTok+Rank</i> | 33.9 | 93.1 | 31.5 |
| | onlineSL | 34.1 | 90.5 | 30.9 |
| | +Rank | 34.4 | 92.1 | 31.7 |
| En-Es | <i>vTok+Rank</i> | 49.3 | 95.5 | 47.1 |
| | onlineSL | 49.8 | 93.9 | 46.7 |
| | +Rank | 49.3 | 98.4 | 48.5 |

Table 3. Ablation study showing improvement for *onlineSL* in all language pairs. The results show re-ranking is additive and these approaches are orthogonal.

and *vTok+Rank*. As it turns out *onlineSL* is better than *vTok* in all languages except En-Fr. This is a very promising result because in one of the language direction (En-Es), we even outperform the *vTok+Rank* system.

5.3. Ablation Study

We also ran an ablation study to check if hypotheses re-ranking complements the self-learning approach. We observe that *onlineSL+Rank* is up to 9% better than the *onlineSL* system in terms of relative *LCB* improvements on En-It as shown in Table 3. In fact, *onlineSL+Rank* beats the state of the art *vTok+Rank* on three language pairs (except En-It) in terms of *LCB*. Concerning the specific metrics there are only two instances where *vTok+Rank* outperformed *onlineSL+Rank*, BLEU for En-It (38.5 vs. 37.0) and LC for En-De (93.1 vs. 92.1). Overall our analysis shows that re-ranking adds on top of the self-learning approach.

6. CONCLUSION

In this work, we propose self-learning based approaches to learn an MT model that can generate isometric translations (*i.e.*, *translations matching the source in character count*). We evaluate the proposed approach on an automatic dubbing use case, where script translation is expected to match the source length to achieve a synchronicity between source and synthetic target language utterances. Our findings both from automatic and subjective human evaluations show the proposed approach can perform better than a strong model with MT output length control and is on par with the current state of the art that requires generating multiple hypothesis and re-ranking.

7. REFERENCES

- [1] Marcello Federico, Robert Enyedi, Roberto Barra-Chicote, Ritwik Giri, Umut Isik, Arvindh Krishnaswamy, and Hassan Sawaf, “From Speech-to-Speech Translation to Automatic Dubbing,” in *Proc. of IWSLT*, Online, July 2020, pp. 257–264, ACL.
- [2] Alp Öktem, Mireia Farrùs, and Antonio Bonafonte, “Prosodic Phrase Alignment for Machine Dubbing,” in *Proc. Interspeech*, 2019.
- [3] Marcello Federico, Yogesh Virkar, Robert Enyedi, and Roberto Barra-Chicote, “Evaluating and optimizing prosodic alignment for automatic dubbing,” in *Proc. of Interspeech*, 2020, p. 5.
- [4] Frederic Chaume, “Synchronization in dubbing: A translation approach,” in *Topics in Audiovisual Translation*, Pilar Orero, Ed. 2004, pp. 35–52, John Benjamins B.V.
- [5] Ashutosh Saboo and Timo Baumann, “Integration of Dubbing Constraints into Machine Translation,” in *Proc. of WMT*, Florence, Italy, Aug. 2019, pp. 94–101, ACL.
- [6] Surafel Melaku Lakew, Mattia Di Gangi, and Marcello Federico, “Controlling the output length of neural machine translation,” in *Proc. IWSLT*, 2019.
- [7] Surafel Lakew, Marcello Federico, Yue Wang, Cuong Hoang, Yogesh Virkar, Roberto Barra-Chicote, and Robert Enyedi, “Machine translation verbosity control for automatic dubbing,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [8] Nicola Ueffing, Gholamreza Haffari, and Anoop Sarkar, “Semi-supervised model adaptation for statistical machine translation,” *Mach. Transl.*, vol. 21, no. 2, pp. 77–94, 2007.
- [9] Rico Sennrich, Barry Haddow, and Alexandra Birch, “Improving neural machine translation models with monolingual data,” *arXiv preprint arXiv:1511.06709*, 2015.
- [10] Sergey Edunov, Myle Ott, Michael Auli, and David Grangier, “Understanding back-translation at scale,” *arXiv preprint arXiv:1808.09381*, 2018.
- [11] Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato, “Revisiting self-training for neural sequence generation,” *arXiv preprint arXiv:1909.13788*, 2019.
- [12] Isaac Caswell, Ciprian Chelba, and David Grangier, “Tagged back-translation,” *arXiv preprint arXiv:1906.06442*, 2019.
- [13] Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi, “MuST-C: a Multilingual Speech Translation Corpus,” in *Proc. NAACL*, 2019, pp. 2012–2017.
- [14] Jan Niehues, “Machine Translation with Unsupervised Length-Constraints,” *Proc. of AMTA*, Apr. 2020.
- [15] Angela Fan, David Grangier, and Michael Auli, “Controllable Abstractive Summarization,” *Proc. of the 2nd Workshop on Neural Machine Translation and Generation*, Nov. 2017.
- [16] Sho Takase and Naoaki Okazaki, “Positional Encoding to Control Output Sequence Length,” *Proc. of NAACL*, Apr. 2019.
- [17] Sweta Agrawal and Marine Carpuat, “Controlling Text Complexity in Neural Machine Translation,” in *Proc. of EMNLP-IJCNLP*, Hong Kong, China, Nov. 2019, pp. 1549–1564, ACL.
- [18] Rico Sennrich, Barry Haddow, and Alexandra Birch, “Controlling politeness in neural machine translation via side constraints,” in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016, pp. 35–40.
- [19] Xing Niu and Marine Carpuat, “Controlling neural machine translation formality with synthetic supervision,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, vol. 34, pp. 8568–8575.
- [20] Cong Duy Vu Hoang, Gholamreza Haffari, and Trevor Cohn, “Improved Neural Machine Translation using Side Information,” in *Proc. of the Australasian Language Technology Association Workshop 2018*, Dunedin, New Zealand, Dec. 2018, pp. 6–16.
- [21] Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al., “Google’s multilingual neural machine translation system: Enabling zero-shot translation,” *Trans. of the ACL*, vol. 5, pp. 339–351, 2017.
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Proc. NIPS*, 2017, pp. 5998–6008.
- [23] Danni Liu, Jan Niehues, and Gerasimos Spanakis, “Adapting end-to-end speech recognition for readable subtitles,” *arXiv preprint arXiv:2005.12143*, 2020.
- [24] Alina Karakanta, Matteo Negri, and Marco Turchi, “Is 42 the answer to everything in subtitling-oriented speech translation?,” in *Proceedings of the 17th International Conference on Spoken Language Translation (IWSLT)*, 2020.
- [25] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al., “Google’s neural machine translation system: Bridging the gap between human and machine translation,” *arXiv preprint arXiv:1609.08144*, 2016.
- [26] Guillaume Lample and Alexis Conneau, “Cross-lingual language model pretraining,” *arXiv preprint arXiv:1901.07291*, 2019.
- [27] Surafel M Lakew, Quintino F Lotito, Matteo Negri, Marco Turchi, and Marcello Federico, “Improving zero-shot translation of low-resource languages,” *arXiv preprint arXiv:1811.01389*, 2018.
- [28] Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich, “Improving massively multilingual neural machine translation and zero-shot translation,” *arXiv preprint arXiv:2004.11867*, 2020.
- [29] Diederik Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [30] Yogesh Virkar, Marcello Federico, Robert Enyedi, and Roberto Barra-Chicote, “Improvements to prosodic alignment for automatic dubbing,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [31] Matt Post, “A call for clarity in reporting BLEU scores,” in *Proc. of the Third Conference on Machine Translation: Research Papers*, Belgium, Brussels, Oct. 2018, pp. 186–191, Association for Computational Linguistics.