

# Linguistically-driven Multi-task Pre-training for Low-resource Neural Machine Translation

ZHUOYUAN MAO, Graduate School of Informatics, Kyoto University, Japan

CHENHUI CHU, Graduate School of Informatics, Kyoto University, Japan

SADAO KUROHASHI, Graduate School of Informatics, Kyoto University, Japan

In the present study, we propose novel sequence-to-sequence pre-training objectives for low-resource machine translation (NMT): Japanese-specific sequence to sequence (JASS) for language pairs involving Japanese as the source or target language, and English-specific sequence to sequence (ENSS) for language pairs involving English. JASS focuses on masking and reordering Japanese linguistic units known as bunsetsu, whereas ENSS is proposed based on phrase structure masking and reordering tasks. Experiments on ASPEC Japanese–English & Japanese–Chinese, Wikipedia Japanese–Chinese, News English–Korean corpora demonstrate that JASS and ENSS outperform MASS and other existing language-agnostic pre-training methods by up to +2.9 BLEU points for the Japanese–English tasks, up to +7.0 BLEU points for the Japanese–Chinese tasks and up to +1.3 BLEU points for English–Korean tasks. Empirical analysis, which focuses on the relationship between individual parts in JASS and ENSS, reveals the complementary nature of the subtasks of JASS and ENSS. Adequacy evaluation using LASER, human evaluation, and case studies reveals that our proposed methods significantly outperform pre-training methods without injected linguistic knowledge and they have a larger positive impact on the adequacy as compared to the fluency.

CCS Concepts: • **Computing methodologies** → **Machine translation**.

Additional Key Words and Phrases: low-resource neural machine translation, pre-training, linguistically-driven

## ACM Reference Format:

Zhuoyuan Mao, Chenhui Chu, and Sadao Kurohashi. 2021. Linguistically-driven Multi-task Pre-training for Low-resource Neural Machine Translation. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.* 1, 1 (December 2021), 31 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Introduction

Neural machine translation (NMT) [3, 49] can achieve state-of-the-art performance when large parallel corpora are available for training. However, this prerequisite for parallel corpora limits its usefulness for several language pairs, such as Japanese, Chinese, and Korean, along with domains (history and COVID) for which such large corpora do not exist. Often, these resource-poor language pairs consist of languages that have resource-rich monolingual corpora. Therefore, it is possible to compensate for the lack of parallel corpora by leveraging large monolingual corpora. One popular approach for this is data augmentation, for instance, through back-translation [12, 42]. Another approach involves pre-training the NMT model on tasks that only require monolingual corpora [37, 47].

---

Authors' addresses: Zhuoyuan Mao, [zhuoyuanmao@nlp.ist.i.kyoto-u.ac.jp](mailto:zhuoyuanmao@nlp.ist.i.kyoto-u.ac.jp), Graduate School of Informatics, Kyoto University, Japan; Chenhui Chu, [chu@nlp.ist.i.kyoto-u.ac.jp](mailto:chu@nlp.ist.i.kyoto-u.ac.jp), Graduate School of Informatics, Kyoto University, Japan; Sadao Kurohashi, [kuro@nlp.ist.i.kyoto-u.ac.jp](mailto:kuro@nlp.ist.i.kyoto-u.ac.jp), Graduate School of Informatics, Kyoto University, Japan.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2021 Association for Computing Machinery.

2375-4699/2021/12-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

As a promising technique for leveraging monolingual corpora, pre-training has experienced a surge in popularity in NLP ever since models such as BERT [8] achieved state-of-the-art results in text understanding. However, BERT-like models were not designed to be used for NMT in the sense that they are essentially techniques for pre-training encoders, but not sequence-to-sequence models. To address this, Song et al. [47], Lewis et al. [21] and Liu et al. [23] recently proposed self-supervised language-agnostic pre-training methods, which are sequence-to-sequence pre-training tasks for NMT, have achieved new state-of-the-art results in low-resource scenarios.

Languages that are sufficiently “rich” to have large monolingual corpora often have available tools for linguistic analysis. Meanwhile, usually a low-resource language pair is composed by a resource-rich language and a low-resource language and the linguistic knowledge of the resource-rich language can be easily extracted. In addition, studies such as Sennrich and Haddow [41] and Murthy et al. [27] have demonstrated that linguistic knowledge can improve NMT without using additional corpora. Therefore, it is natural to use monolingual corpora and linguistic tools in bilingual low-resource scenarios. However, the manner in which linguistic knowledge should be provided is not always clear, because NMT models are implemented in an end-to-end scheme. From a technical perspective, it is practical to extract linguistic features on the monolingual side. Therefore, monolingual pre-training provides an ideal framework for leveraging monolingual corpora and injecting linguistic information.

In Mao et al. [24], we proposed a linguistically motivated pre-training approach known as Japanese-specific sequence-to-sequence (JASS), which was inspired by MASS, but focused on syntactic analysis obtained by using a parser. Particularly, we added syntactic constraints to the sentence-masking process of the MASS to obtain the BMASS task.<sup>1</sup> We also proposed the Bunsetsu reordering-based sequence to sequence (BRSS), which is a linguistically motivated reordering task. Several previous studies [21, 39] have provided evidence that “multi-task” pre-training that combines various styles of self-supervised training tasks results in significantly superior results for NMT. We proposed JASS based on a combination of the above-mentioned two tasks and it is tailored for NMT involving Japanese.

In contrast, in this study, we also propose linguistically-driven pre-training methods for English to leverage linguistic-specific information in the pre-training phase.<sup>2</sup> They are referred to as phrase structure-based MASS (PMASS) & head finalization-based sequence-to-sequence (HFSS), and their combination is denoted as English-specific sequence-to-sequence (ENSS).<sup>3</sup> Moreover, unlike the proposed methods for Japanese, the proposed methods for English can be transplanted onto any SVO language. Thus, our proposed ENSS and JASS can be applied to any translation pair involving English or Japanese.<sup>4</sup>

We experimented with ASPEC Japanese–English & Japanese–Chinese [30], Wikipedia Japanese–Chinese [4, 5], and News English–Korean [32] in various pre-training settings for JASS and ENSS.<sup>5</sup> Our results indicate that BMASS, BRSS, and HFSS significantly outperform the state-of-the-art MASS pre-training, whereas PMASS yields marginal improvements. Furthermore, we demonstrate that linguistically-driven multi-task pre-training methods (JASS & ENSS) lead to further improvements

<sup>1</sup>For BMASS, bunsetsus are used as syntactic spans, which is the elementary syntactic component of Japanese. It can be extracted using the KNP. [20, 26]

<sup>2</sup>Although some language pairs involving English are middle- or high-resource scenarios (parallel corpora size over 100k), we deem that it is worth proposing methods for English because a large number of low-resource language pairs involving English are still present.

<sup>3</sup>Head finalization [14] is the technique used to reorder sentences in SVO language to be SOV-like sentences.

<sup>4</sup>According to the reordering task we proposed (specifically, BRSS and HFSS), more significant improvements are expected to observe on English–SOV language or Japanese–SVO language.

<sup>5</sup>In Mao et al. [24], we only conducted experiments on ASPEC Japanese–English and JaRuNC Japanese–Russian for JASS (BLEU results on Japanese–Russian were excessively low for comparison).

of up to +2.9 BLEU points for Japanese to English, +2.7 BLEU points for English to Japanese, +4.3 BLEU points for Japanese to Chinese, +7.0 BLEU points for Chinese to Japanese, +0.5 BLEU points for English to Korean, and +1.3 BLEU points for Korean to English in low-resource scenarios, respectively.

Unlike in our previous study Mao et al. [24], we provide substantial analyses for evaluating the translations generated by JASS and ENSS, which focus on the relationship between different pre-training tasks, and the specific adequacy and fluency of corresponding translations. Specifically, we validate the superior translation adequacy improvement of linguistically-driven methods by implementing automatic adequacy evaluation using LASER, human evaluation, and case study. To confirm the complementary nature between the masked language model and reordering the pre-training task, we performed an evaluation of the pre-training accuracy.

We expect this study to extend the usefulness of linguistically-driven pre-training methods for more low-resource language pairs and compensate for the defects of Mao et al. [24] in terms of the empirical evaluation. The contributions of this study can be summarized as follows.

- (1) **BMASS and BRSS:** Linguistically-driven novel pre-training methods for NMT involving Japanese.
- (2) **PMASS and HFSS:** Linguistically-driven novel pre-training methods for NMT involving English (can be theoretically implemented on any SVO language).
- (3) **Multi-task pre-training (JASS and ENSS):** We demonstrate that multi-task training through the combination of the masked language model and reordering task (BMASS+BRSS & MASS+HFSS) leads to better performance. Particularly, BMASS and BRSS can complement each other more if they are performed based on analogous syntactic units.
- (4) **Empirical evaluation:** Comparisons among MASS, BART, JASS, ENSS and newly added baseline methods (MultiMASS and Deshuffling) for 6 translation directions and 3 different domains in several data size settings to identify situations in which each technique can be the most effective compared to other techniques.
- (5) **Analyses:** Linguistic and statistical analyses of pre-training methods, their inter-relationships, and corresponding translations.

## 2 Related Work

### 2.1 Low-resource Neural Machine Translation

There are mainly three lines of work related to improving NMT in low-resource situations: cross-lingual transfer, data augmentation, and monolingual pre-training. These approaches are potentially complementary. Our work belongs to the monolingual pre-training category.

Cross-lingual transfer addresses the low-resource issue by using data from different language pairs. One can use a richer language pair [60] or several language pairs simultaneously [7, 9]. Murthy et al. [27] also proposed reordering the assisting languages to be similar to a low-resource language.

Data augmentation involves the creation of synthetic bilingual data from monolingual data. In the popular back-translation approach [10, 12, 42], the source side of the data is synthesized using an MT system to back-translate the target side data. Recently, Zhou et al. [58] proposed the creation of this source side through rule-based reordering via word-to-word translation.

In monolingual pre-training approaches, all or part of a model is first trained on tasks that require monolingual data.<sup>6</sup> Pre-training has enjoyed significant success in other NLP tasks with the development of GPT [38], BERT [8], and several others [33, 48, 53].

<sup>6</sup>This is an instance of “transfer learning,” similar to cross-lingual transfer. “Pre-training” often implies that the training task differs from the target task.

Pre-training schemes such as BERT were designed for natural language understanding (NLU) tasks and they are not directly suitable for NMT. Conneau and Lample [6] and Ren et al. [40] proposed multilingual variants. However, they trained the encoder and decoder independently. To address this, Song et al. [47] recently proposed MASS, a new state-of-the-art NMT pre-training task that jointly trains the encoder and decoder. Our approach develops on the initial idea of MASS, but adds more diverse and linguistically-motivated training objectives.

Linguistic information is known to be useful for NMT [41], especially in low-resource scenarios. Outside of pre-training, studies [27, 56, 58] have successfully used a linguistically-motivated reordering similar to that of our BRSS task. Sun et al. [48] used linguistically-motivated pre-training tasks for text understanding. To the best of our knowledge, there are no studies on linguistically-motivated pre-training tasks for NMT.

## 2.2 Pre-training Tasks for Neural Machine Translation

After the appearance of BERT [8], several pre-training methods have been proposed to enhance NMT [6, 21–23, 39, 40, 45–47, 51, 52, 54]. Particularly, Song et al. [47] proposed a random span reconstruction task to pre-train a sequence-to-sequence framework for NMT; Wang et al. [51] first proposed using shuffling, deleting, and replacing operations to implement the denoising pre-training for the NMT system; thereafter, Lewis et al. [21] combined the denoising methods with the masked language model pre-training of Song et al. [47], and provided detailed empirical results for a large number of language pairs; mBART [23] is a multilingual sequence-to-sequence denoising pre-training that is pre-trained through denoising tasks on 25 languages including Japanese, English, Chinese, Russian, and others, and it can be deemed as an extension of Lewis et al. [21]; other studies focus on leveraging the cross-lingual supervision between languages through word alignment [22], phrase alignment [40], sentence-level alignment [6], code-switching technique [54], or assisting languages (shared scripts) [46].

Among the above-mentioned pre-training techniques for NMT, we observe that no study has focused on leveraging specific linguistic features for NMT. Syntactic span-masking [59] and semantic-aware BERT [57] have been proposed using linguistically-driven supervision for language understanding tasks. However, linguistically-driven methods for sequence-to-sequence pre-training should be considered and explored.

Studies have also focused on improving MASS. Siddhant et al. [45] adapted MASS in multilingual scenarios; Qi et al. [36] proposed using an n-stream self-attention mechanism to enhance MASS for language generation tasks. No previous study has attempted to enhance MASS from a linguistic perspective, which will be explored in our study.

Moreover, Wang et al. [52] highlighted that multitask learning can significantly benefit multilingual NMT. In addition to the MT task, the essential jointly-learned tasks should be masked language model task and denoising (reconstruction) task, which are two basic pre-training styles based on which we propose our linguistically-driven methods.

## 3 Preliminary Backgrounds

In this section, we introduce the preliminary backgrounds of pre-training and fine-tuning for NMT and MASS, which serve as the backbone for this study.

### 3.1 Pre-training and Fine-tuning for NMT

We first introduce the pre-training and fine-tuning pipelines for the NMT. As shown in Figure 1 below, we first utilize monolingual corpora to pre-train the initialized sequence-to-sequence model. Subsequently, we use a parallel corpus of languages of interest to fine-tune the pre-trained models.

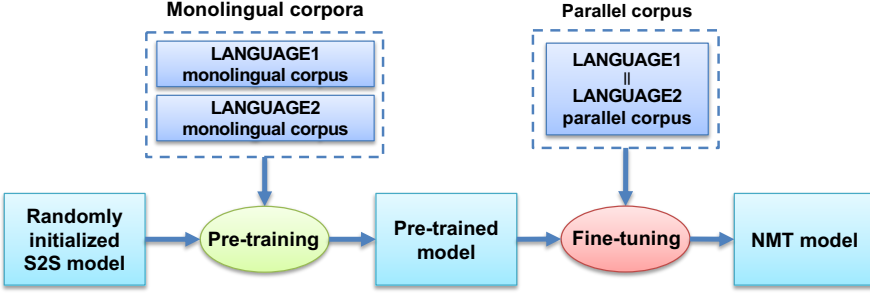


Fig. 1. **Pre-training and fine-tuning for NMT.** “L1” and “L2” refer to two languages involved in the translation task. “S2S” denotes sequence-to-sequence.




Fig. 2. **Sequence-to-sequence structure for MASS.**  $x_i$  represents a token and  $x_3$  to  $x_6$  are consecutive tokens to be masked/predicted.

The fine-tuned model was the final NMT model. All the experiments in this study will be conducted on the basis of this pre-training and fine-tuning pipeline for NMT.

### 3.2 MASS

MASS is a pre-training method for NMT proposed by Song et al. [47]. As shown in Figure 2, in MASS pre-training, the input is a sequence of tokens where a part of the sequence is masked and the output is a sequence where the masking is inverted.

We consider  $x \in \mathcal{X}$ , which is a sequence of tokens where  $\mathcal{X}$  is a monolingual corpus. Additionally, we consider the token span  $C = [p_i, p_j]$ , where  $0 < p_i \leq p_j \leq \text{len}(x)$  and  $\text{len}(x)$  are the number of tokens in sentence  $x$ . We denote the masked sequence by  $x^C$ , where tokens in positions from  $p_i$  to  $p_j$  in  $x$  are replaced by a mask token  $[M]$ .  $x^{!C}$  is the sequence with an inverted mask, that is, where tokens in positions other than the aforementioned fragments are replaced by the mask token  $[M]$ . In MASS, the pre-training objective is to predict the masked fragments in  $x$  using an encoder-decoder model, where  $x^C$  is the input to the encoder and  $x^{!C}$  is the target output of the decoder. The log-likelihood objective function is

$$\mathcal{L}_{\text{mass}}(\mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \log P(x^{!C} | x^C, \theta). \quad (1)$$

where  $\theta$  denotes the model parameters. The number of tokens to be masked is a hyperparameter of the MASS. The NMT model is jointly pre-trained with the MASS task for both the source and target languages.

## 4 Proposed Methods

In this section, we describe JASS and ENSS, which are our proposed pre-training techniques.




Fig. 3. Word and bunsetsu segmentations for a Japanese sentence with meaning “LoveLive is made of three projects.” In word for word English translations, “\_” represents words with no meaningful translations.

#### 4.1 Proposed Methods for Japanese

Our methods are based on the ideas of the original MASS and are improved by jointly learning multiple linguistics-aware tasks. For Japanese, we propose a bunsetsu-based (BMASS) pre-training and bunsetsu reordering-based sequence to sequence (BRSS) pre-training. Their combination, Japanese-specific sequence-to-sequence (JASS) pre-training, is introduced in the following section.

**4.1.1 Bunsetsu** Bunsetsu is the syntactic component of Japanese sentences [20, 26]. It is equivalent to the concepts of noun phrases or verb phrases in English syntax and it constitutes a minimal unit of meaning. The concept of “word” is ambiguous for writing systems such as Japanese where word-separators are not applicable, and Japanese segmenters [20, 26] can segment Japanese sentences either in words or bunsetsus. Therefore, bunsetsu is also more likely to correspond to a well-defined entity or concept than words. Figure 3 illustrates the difference between the word- and bunsetsu-level segmentation. Each bunsetsu contains self-contained information and case markers, which indicate its relation with other bunsetsus. Based on the bunsetsu, we introduce our proposed pre-training techniques for the Japanese.

**4.1.2 BMASS** We propose BMASS, which leverages syntactically parsed Japanese monolingual data for sequence-to-sequence pre-training. MASS pre-trains an NMT model by making it predict random parts of a sentence given their context, whereas BMASS involves making the model predict a set of bunsetsus given the contextual bunsetsus. We expect this will allow the model to learn about bunsetsus and thereby focus on predicting meaningful subsequences instead of random, albeit fluent subsequences.

To perform BMASS, we modify the definition of mask  $C$  in Equation 1:  $C = [[p_{i_1}, p_{j_1}], [p_{i_2}, p_{j_2}], \dots, [p_{i_n}, p_{j_n}]]$ , where  $0 < p_{i_1} \leq p_{j_1} \leq p_{i_2} \leq p_{j_2} \leq \dots p_{i_n} \leq p_{j_n} \leq \text{len}(x)$ . Term  $\text{len}(x)$  denotes the number of tokens in sentence  $x$ . Subsequently, the  $k$ -th position span from  $p_{i_k}$  to  $p_{j_k}$  corresponds to the start and end of a specific bunsetsu in a Japanese sentence. Consequently, we denote the BMASS loss as  $\mathcal{L}_{bmass}$ . The main difference between MASS and BMASS is that in MASS, we mask random token spans, whereas in BMASS, we only mask tokens spans that are complete bunsetsus. The number of bunsetsus to be masked constitutes a hyperparameter for BMASS. Figures 4-b and 4-c provide training pairs for MASS and BMASS.

Note that our BMASS pre-training task differs from the entity masking task of ERNIE [48] and random span masking of SpanBERT [16]. ERNIE and SpanBERT have been proposed without using syntactic units and they are employed in natural language understanding downstream tasks.

**4.1.3 BRSS** Japanese sentences are typically in an SOV word order that can be reordered to SVO to reduce the difficulty of translation to languages with SVO order. We first define a simple process for reordering a (typically SOV) Japanese sentence into a “SVO Japanese” pseudo-sentence that will be used in BRSS. There are several previous studies on reordering a SOV-ordered sentence to a SVO-ordered sentence [13, 19]. In our case, to consistently leverage bunsetsu units in Japanese with BMASS, we propose bunsetsu-based reordering, which is able to generate an SVO-ordered Japanese sentence while retaining syntactic information at the bunsetsu-level. We first define “chunking



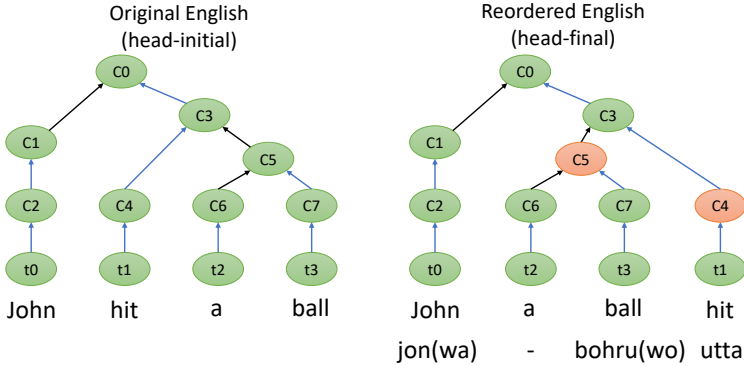


Fig. 5. Example of HPSG parsing result and head finalization. Head finalization reorders an English sentence into a Japanese-like sentence. [14] Blue arrows denote the “head.”

for extracting phrase structures in sentences and applying the following proposed pre-training techniques. Figure 5 (left) shows an instance of parsing an English sentence using HPSG grammar.

**4.2.2 Head Finalization** Using the above-mentioned HPSG, sentences in any language can be characterized using phrase structures. From the definition of a phrase, the “head” of a phrase is subsequently defined as the syntactically determinant part in a phrase. In other words, “head” determines the syntactic category of the phrase and its “dependents.” Particularly, English is referred to as a “head-initial” language because “head” appears before its “dependents,” whereas Japanese is referred to as a “head-final” language because “head” usually appears after “dependents” in a phrase.

The deliberate phrase structures provided by the HPSG parser are utilized in several scenarios in the NLP. Particularly, Isozaki et al. [14] proposed a simple reordering rule for the SVO language (head-initial languages) by using the phrase structure information provided by the HPSG parser. Figure 5 shows an example of reordering an English sentence to be an SOV-like sentence on the basis of the result of HPSG parsing. By reordering sentences in SVO languages such as English to be SOV-like sentences, the performance of statistical machine translation (SMT) is improved. Particularly, Isozaki et al. [14] first proposed head finalization and applied it to English-to-Japanese SMT; Han et al. [11] applied it to Chinese-to-Japanese SMT and obtained significant improvements; more recently, Zhou et al. [58] utilized this reordering technique to generate synthetic parallel sentences in the back-translation phase when translating SOV and SVO languages. In this study, we utilize this reordering rule in the pre-training phase for NMT (see Section 4.2.4).

**4.2.3 PMASS** We propose PMASS by leveraging phrase-span information in an English sentence. In general, we perform PMASS pre-training by limiting the masked tokens in MASS to be an entire phrase span. Thus, for masking plural phrase spans, we denote it as PMASS.P. For masking only a single phrase span, we denote it as PMASS.S. Particularly, the source and target for PMASS.P and PMASS.S pre-training can be generated using our proposed phrase-masking algorithms described in Appendix A. Inspired by MASS, we force the number of masked tokens to be approximately half of the length of the sentence to guarantee the effectiveness of the sequence-to-sequence masked language model. Examples of PMASS.P and PMASS.S are presented in Figure 6-c. We observe that several phrase spans in PMASS.P and a single long phrase span in PMASS.S are masked. We expect such special masking patterns to force the NMT system to extract more phrase-level syntactic information in the pre-training phase.



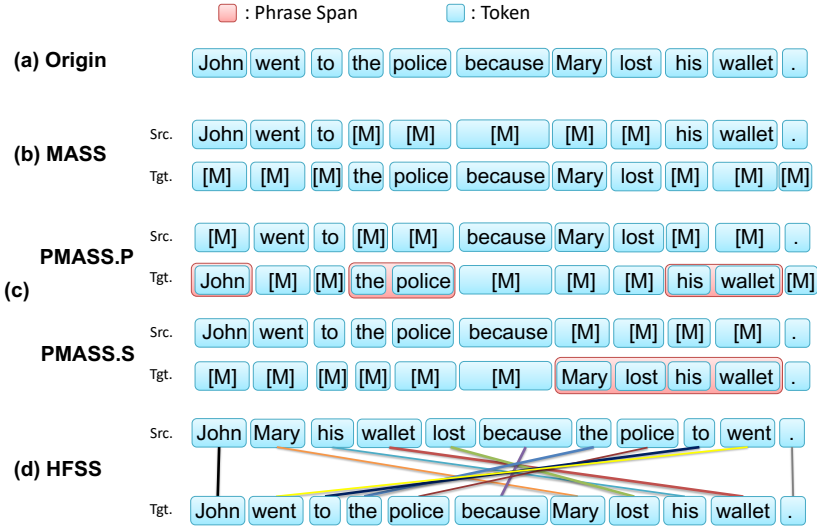


Fig. 6. Example of source and target for MASS, PMASS, and HFSS of a sentence in English.

**4.2.4 HFSS** We propose HFSS using the head finalization technique [14] for pre-training English. As shown in Figure 6-d, the pre-training task is also a reordering task that simulates the translation from SOV languages to English. More precisely, the source sentence for sequence-to-sequence pre-training is the reordered (SOV-like or head-finalized) English sentence, and the target sentence is the original English monolingual sentence. We expect HFSS to help the system learn the word reordering pattern of the translation between head-initial (SVO) and head-final (SOV) languages in advance.

According to the prior experiments for Japanese (see Mao et al. [24] and 6.1), BRSS.F consistently outperforms BRSS.R. In addition, BART [21] also claims that reconstructing the original sentence benefits the language generation tasks. Therefore, we do not distinguish HFSS with HFSS.F and HFSS.R (HFSS.F performs pre-training with the SOV-SVO pattern, whereas HFSS.R performs the reverse pattern).<sup>7</sup> Instead, we directly defined HFSS using the pre-training pattern of HFSS.F. Moreover, HFSS is performed on the basis of head finalization, which utilizes the results from HPSG parsers. This is consistent with PMASS in which we extract phrases using HPSG-parsing results.

We develop our proposal on English through head finalization, whereas for SOV languages such as Japanese, it is unmanageable to reorder SOV sentences to SVO-like sentences [14]. Furthermore, HFSS can be used for all head-initial languages apart from English, as well-developed reordering rules have been proposed and demonstrated to be effective for NMT. However, BRSS can only be implemented for Japanese-involved translation pairs because *bunsetsu* information is required to establish the source and target sentences for sequence-to-sequence pre-training.

<sup>7</sup>More precisely, HFSS.F denotes the source sentence of the head-finalized English sentence and the target sentence of the original English sentence. HFSS.R indicates the source sentence of the original English sentence and the target sentence of the head-finalized English sentence.

### 4.3 Multi-task Pre-training

Multi-task pre-training objectives lead to a robust initial state for NMT systems [21, 39]. Because our proposed methods can also be categorized into two groups of pre-training tasks, we propose a multi-task pre-training task for both Japanese and English.

We define JASS pre-training, which is a combination of the two previous procedures: BMASS and BRSS. Our actual pre-training will consist of the joint execution of these two pre-training sessions. Therefore, the pre-training objective for JASS is

$$\mathcal{L}_{jass}(\mathcal{X}_{ja}) = \mathcal{L}_{bmass}(\mathcal{X}_{ja}) + \mathcal{L}_{brss}(\mathcal{X}_{ja}) \quad (2)$$

where  $\mathcal{X}_{ja}$  represents the monolingual corpus of Japanese, and  $\mathcal{L}_{brss}$  denotes the reordering loss using the forward or reverse variants mentioned in Section 4.1.3. We expect BMASS & BRSS to jointly learn syntactic knowledge and BRSS to learn word ordering knowledge.

For English, we similarly define ENSS pre-training, which combines PMASS and HFSS. More precisely, the training objective is:

$$\mathcal{L}_{enss}(\mathcal{X}_{en}) = \mathcal{L}_{pmass}(\mathcal{X}_{en}) + \mathcal{L}_{hfss}(\mathcal{X}_{en}) \quad (3)$$

where  $\mathcal{X}_{en}$  denotes the monolingual corpus of English,  $\mathcal{L}_{pmass}$  the PMASS.P or PMASS.S loss, and  $\mathcal{L}_{hfss}$  the reordering loss of HFSS.

JASS is specifically designed for Japanese, whereas theoretically, ENSS can be transplanted onto any SVO language as long as we can extract the phrase structure information of the corresponding language from a HPSG parser.

We also mixed JASS pre-training for Japanese with MASS pre-training for the other languages involved in the translation. In practice, we therefore designated using JASS pre-training for Japanese monolingual data with BMASS and BRSS objectives, along with "other languages" monolingual data with the MASS objective. Similarly, for English, ENSS pre-training consists of PMASS & HFSS for English and MASS for "other languages" involved in fine-tuning translation pair.

We also consider attempting the combination of our proposed linguistically-driven methods with a strong baseline pre-training objective, MASS, which we refer to as MASS + JASS (or ENSS) in the subsequent sections. To allow the pre-training model to determine the language and sub-task (MASS, BMASS, BRSS, PMASS, and HFSS) that it should perform, we prepend tags to inputs similar to those used in Johnson et al. [15] (see Section 5.2 for details).

## 5 Experimental Settings

In this section, we evaluate our pre-training methods on simulated low-resource scenarios for ASPEC Japanese–English [30], Japanese–Chinese translations [28], and realistic low-resource scenarios for Wikipedia Japanese–Chinese [4, 5] and News English–Korean [32] translations.

### 5.1 Datasets

We used monolingual data for pre-training and parallel data for fine-tuning. Refer to Table 1 for an overview.

**Monolingual data:** For pre-training, we use monolingual data of 22M lines each for Japanese, English, Chinese, and Korean, randomly sub-sampled from Common Crawl mentioned in the official WMT monolingual training data.<sup>8</sup> <sup>9</sup> For pre-training in Japanese–English and English–Korean, given that these two languages have different scripts and thus have few common words, the pre-training objectives for each language will work separately, even though they are performed jointly

<sup>8</sup><http://www.statmt.org/wmt19/translation-task.html>

<sup>9</sup>Different from Mao et al. [24]. Currently, we unify the monolingual corpus domains for all the languages for fairer comparisons.

	Language	Dataset	Size
Monolingual	Ja	Common Crawl	22M
	Zh	Common Crawl	22M
	En	Common Crawl	22M
	Ko	Common Crawl	22M
Parallel	Ja-En	ASPEC-JE	1M
	Ja-Zh	ASPEC-JC	670k
	Ja-Zh	Wikipedia	258k
	En-Ko	News	94k

Table 1. Overview of training data. “Size” denotes the number of the monolingual sentences or parallel sentences.

for two languages. However, for pre-training in Japanese and Chinese, they share more characters, which indicates that the monolingual pre-training tasks will be run in a pseudo-cross-lingual manner. Thus, we also expect to see whether such pre-training will benefit from more fine-tuning. **Parallel Data:** We use scientific abstracts domain ASPEC parallel corpus for training Japanese–English and Japanese–Chinese models. For Japanese–Chinese fine-tuning, we also utilize the Wikipedia parallel corpus, which is a real low-resource scenario. We use News parallel corpus for English–Korean, which is a low-resource dataset.

For ASPEC, we used the official training, development, and test splits provided by WAT 2019.<sup>10</sup>

<sup>11</sup> For Wikipedia, we used the dataset released by Kyoto University.<sup>12</sup> For News, we use dataset provided by Park et al. [32].<sup>13</sup>

## 5.2 Pre-processing

We tokenize the monolingual data by using the Moses tokenizer for English and Korean,<sup>14</sup> Jumanpp for Japanese,<sup>15</sup> and jieba for Chinese.<sup>16</sup> We obtain the bunsetsu information by using KNP<sup>17</sup> and obtain the HPSG parsing results using enju.<sup>18</sup> Sentences with more than 175 tokens were removed. For each language pair, we constructed a joint vocabulary with 60 000 sub-word units through byte-pair encoding (BPE) [43] on the concatenated monolingual corpora involved during pre-training.<sup>19</sup>, whereas 40 000 BPE merge operations is set for Japanese-English. In the multi-task pre-training, each sentence is prepended with a task token [*MASS*], [*BMASS*], [*BRSS*], [*PMASS*], or [*HFSS*], and a language token [*Ja*], [*En*], [*Zh*], or [*Ko*].<sup>20</sup> This ensures that the model learns to distinguish

<sup>10</sup><http://lotus.kuee.kyoto-u.ac.jp/WAT/WAT2019/index.html#task.html>

<sup>11</sup>For ASPEC Japanese–English, we use the first 1M parallel sentences. Parallel sentences for different fine-tuning size settings were randomly sampled from the selected 1M dataset.

<sup>12</sup><http://nlp.ist.i.kyoto-u.ac.jp/EN/index.php?Wikipedia%20Chinese-Japanese%20Parallel%20Corpus>

<sup>13</sup><https://sites.google.com/site/koreanparalleldata>

<sup>14</sup><https://github.com/moses-smt/mosesdecoder>

<sup>15</sup><https://github.com/ku-nlp/jumanpp>

<sup>16</sup><https://github.com/foxsjy/jieba>

<sup>17</sup><https://github.com/ku-nlp/pyknp>

<sup>18</sup><https://myntp.is.s.u-tokyo.ac.jp/enju/>

<sup>19</sup>Particularly, 30 000 BPE merging operations will lead to a joint vocabulary with a size of approximately 60 000 for Japanese–Chinese and English–Korean

<sup>20</sup>As an implementation trick, we recommend to unify the task tag for the same group of tasks, e.g. use the same tag for [*BRSS*] and [*HFSS*].

between different pre-training objectives and languages. This token can be used when monolingual pre-training is conducted jointly by multiple languages and multiple tasks.

### 5.3 Training and Evaluation Details

In our experiments, we used the open-source OpenNMT [17] implementation of the transformer [50] NMT model.<sup>21</sup> The hyperparameters are set to the transformer-big setting in OpenNMT. Particularly, our model has a 6-layer encoder and decoder, a hidden size of 1024, feed-forward hidden layer size of 4096, batch size of 4096, dropout rate of 0.3, and 16 attention heads. An ADAM optimizer with a learning rate of  $10^{-4}$  was used for both pre-training and fine-tuning. All the pre-training tasks are run until convergence on four TITAN V100 GPU cards occurs, and fine-tuning uses only one GPU. It took approximately two days for each pre-training run. Mixed precision training [25] was used for both pre-training and fine-tuning. For multi-task pre-training, data are randomly shuffled such that even in each mini-batch, different pre-training objectives appear, corresponding to a real joint pre-training. Our proposed pre-training methods converge within the similar training time as compared to that of MASS.

Pre-training tasks are evaluated using perplexity, and the checkpoint with the lowest pre-training perplexity was selected for fine-tuning. We used BLEU [31] for automatic evaluation, adequacy, and fluency for human evaluation. We performed early stopping using 1-gram accuracy and perplexity on the development set. We evaluated the statistical significance of our BLEU scores through bootstrap resampling [18].

### 5.4 Baselines

In addition to MASS, we employ the “text infilling” in BART as another main baseline.<sup>22</sup> We also define two pre-training baselines for comparison with our proposed methods. They are named multi-span-based MASS (MultiMASS) and deshuffling. Moreover, the joint training with MASS and deshuffling was set as the multi-task pre-training baseline. All of the baselines are as follows:

**Baselines without pre-training.** First, we employ the vanilla transformer big as the baseline without pre-training because all of the pre-training methods are based on this model structure. Moreover, following Araabi and Monz [1], we also present the best performance for low-resource NMT by using Transformer model. Hyperparameter details are shown in Appendix B.

**MASS.** Using the same settings as in Song et al. [47].

**BART (text infilling).** Different from MASS, BART (text infilling) masks several token spans within a sentence by a single  $[M]$  where span lengths are samples from Poisson distribution and the model is also required to predict the lengths of the masked spans. We use the same settings as in Lewis et al. [21].<sup>23</sup>

**MultiMASS.** MultiMASS is a baseline method added to help demonstrate the effectiveness of masking specific syntactic units such as bunssetsu or phrase spans in a sentence that we propose as BMASS and PMASS.

<sup>21</sup><https://github.com/OpenNMT/OpenNMT-py>

<sup>22</sup>Text infilling has been demonstrated as the most effective pre-training objective for NMT among several objectives in BART [21].

<sup>23</sup>In order to conduct fair comparisons for our proposed methods, we only present the most effective sub-task, text infilling, within BART. The combination of text-infilling and sentence permutation is proven to be the best practice of BART. With regard to sentence permutation, we do not consider it in this paper because it is mainly designed for document NMT. When it comes to multi-sentence pre-training, sentence permutation and other possible patterns of multi-sentence linguistically-driven pre-training tasks should be explored and compared in future work.




Fig. 7. Example of source and target for MultiMASS with the meaning “LoveLive is made of three projects.”




Fig. 8. Example of source and target for deshuffling with the meaning “LoveLive is made of three projects.”

As shown in Figure 7, MultiMASS predicts several randomly masked tokens in a sentence, which differs from the single masked span in MASS, masked busetsu spans in BMASS, several phrase spans in PMASS.P, and a single phrase span in PMASS.S.

**Deshuffling.** Deshuffling denotes the pre-training task of random shuffling-based sentence reconstruction, which is also a crucial pre-training task. We perform this pre-training task as another baseline to confirm the effectiveness of reordering syntactic units in BRSS and the reordering driven by head finalization of HFSS. A pre-training example is presented in Figure 8.

**Multi-task Baseline.** The multi-task baseline is the combination of the respective best baseline methods from the masked language model and reordering pre-training. Thus, the multi-task baseline consists of MASS,<sup>24</sup> and deshuffling. The baseline is formulated as follows:

$$\mathcal{L}(\mathcal{X}) = \mathcal{L}_{mass}(\mathcal{X}) + \mathcal{L}_{deshuffling}(\mathcal{X}) \quad (4)$$

where  $\mathcal{X}$  represents the monolingual corpora.

## 5.5 Pre-trained Models

We pre-trained our NMT models by leveraging the monolingual data of the source and target languages. For Japanese, we can use MASS, BMASS, or BRSS, whereas for English, we can use MASS, PMASS, or HFSS. For Chinese and Korean, we use only the MASS. Particularly, we pre-trained different types of models in Table 2. Note that we use MASS for ENSS because PMASS underperforms MASS by a significant margin (see 6.1).

## 5.6 Fine-tuned NMT Models

We fine-tuned to improve Japanese-English, English-Japanese, Japanese-Chinese, Chinese-Japanese, English-Korean and Korean-English translations. We trained the following NMT models:

- (1) **Ja-En and En-Ja:** Japanese to English and English to Japanese models using from 3k to 50k parallel sentences randomly sampled from **ASPEC** for fine-tuning.
- (2) **Ja-Zh and Zh-Ja:** Japanese to Chinese and Chinese to Japanese models using from 3k to 50k parallel sentences randomly sampled from **ASPEC** and **Wikipedia**, respectively, for fine-tuning.
- (3) **En-Ko and Ko-En:** English to Korean and Korean to English models using 20k (randomly sampled) and 94k (full dataset) parallel sentences from **News** for fine-tuning.

We compared these models with pre-trained model baselines and vanilla baselines, which are fully-supervised NMT models on the same data settings, but without pre-training. In addition,

<sup>24</sup>MASS outperforms MultiMASS, we therefore use MASS rather than MultiMASS. (See 6.1)

#	Pre-trained Model	Details
<i>Main baseline</i>		
1	MASS	Using the same settings as in Song et al. [47].
1*	BART (text infilling)	Using the same settings as in Lewis et al. [21].
<i>Proposed methods for Japanese</i>		
2	BMASS	Similar to MASS, we mask half the number of bunssetsus during pre-training.
3	BRSS	We separately pre-trained on the SVO–SOV (BRSS.F) as well as SOV–SVO (BRSS.R) models.
4	JASS	Multi-task training of BMASS and BRSS.
<i>Combinations of proposed methods with MASS</i>		
5	MASS+BMASS	Multi-task training of MASS and BMASS.
6	MASS+BRSS	Multi-task training of MASS and BRSS.
7	MASS+BMASS+BRSS	Multi-task training of BMASS, BRSS and MASS.
<i>Other baselines for Japanese</i>		
8	MultiMASS (Ja)	Based on MASS pre-training, several random tokens are masked rather than one consecutive span.
9	Deshuffling (Ja)	Random shuffling-based original sentence reconstruction.
10	MASS+Deshuffling (Ja)	Multi-task pre-training baseline for Japanese.
<i>Proposed methods for English</i>		
11	PMASS	Similar to MASS, we mask an entire phrase span based on the head-driven phrase structure grammar. We performed the experiments for PMASS.P and PMASS.S, respectively.
12	HFSS	We train SOV (head finalized)–SVO (original) models for English.
13	ENSS	Multi-task training of MASS and HFSS.
<i>Other baselines for English</i>		
14	MultiMASS (En)	Based on the MASS, several random tokens are masked rather than one consecutive span.
15	Deshuffling (En)	Random shuffling-based original sentence reconstruction.
16	MASS+Deshuffling (En)	Multi-task pre-training baseline for English.
<i>Combination of the proposed methods for English and Japanese</i>		
17	JASS+ENSS	Multi-task training of JASS and ENSS.
<i>Baseline for #17</i>		
18	MASS+Deshuffling	Multi-task pre-training baseline for JASS+ENSS.

Table 2. Settings of pre-trained models.

fine-tuning results under the high-resource scenarios (with more than 50k parallel sentences) are provided and discussed in 6.6.

## 6 Results and Analyses

Tables 3, 4, 5, and 6 contain the NMT BLEU results of our proposed methods for Japanese–English, Japanese–Chinese and English–Korean translation on various translation domains, respectively. Subsequently, we provide in-depth analysis for translation quality in terms of adequacy by using LASER [2], human evaluation scores, specific cases for the real low-resource scenario of Wikipedia Ja-Zh. Finally, we conduct an investigation on the pre-training accuracy to analyze the difference between the pre-trained models and their complementation of each other, and present the results in middle/high-resource scenarios.

#	Model	Ja-En				En-Ja			
		3k	10k	20k	50k	3k	10k	20k	50k
<i>Main baselines</i>									
0	w/o pre, vanilla	0.8	2.1	3.5	16.1	1.1	2.7	5.1	19.4
0*	w/o pre, optimized	2.2	6.8	10.7	19.8	3.3	6.5	13.6	23.7
1	MASS	8.8	13.8	17.2	21.2	9.1	16.0	20.6	25.0
1*	BART (text infilling)	3.1	11.1	15.5	20.7	5.6	14.9	19.8	25.6 <sup>†</sup>
<i>Proposed methods for Japanese</i>									
2	BMASS	8.9	13.9	17.4	21.8	8.7	15.9	20.1	25.4
3	BRSS	8.8	14.9 <sup>†</sup>	18.1 <sup>†</sup>	22.0 <sup>†</sup>	10.0 <sup>†</sup>	17.3 <sup>†</sup>	21.0	26.0 <sup>†</sup>
3 (R)	BRSS.R	8.2	14.3 <sup>†</sup>	17.7 <sup>†</sup>	21.7 <sup>†</sup>	10.0 <sup>†</sup>	17.2 <sup>†</sup>	20.5	25.7 <sup>†</sup>
4	JASS	10.6 <sup>†</sup>	15.7 <sup>†</sup>	<b>18.9<sup>†</sup></b>	<b>22.3<sup>†</sup></b>	<b>11.5<sup>†</sup></b>	17.7 <sup>†</sup>	21.6 <sup>†</sup>	26.5 <sup>†</sup>
<i>Combinations of proposed methods with MASS</i>									
5	1 + 2	9.2	14.8 <sup>†</sup>	17.7 <sup>†</sup>	21.7 <sup>†</sup>	9.7 <sup>†</sup>	16.6 <sup>†</sup>	20.9	25.9 <sup>†</sup>
6	1 + 3	<b>10.9<sup>†</sup></b>	15.9 <sup>†</sup>	18.3 <sup>†</sup>	22.2 <sup>†</sup>	11.0 <sup>†</sup>	17.7 <sup>†</sup>	<b>21.7<sup>†</sup></b>	<b>26.8<sup>†</sup></b>
7	1 + 4	10.5 <sup>†</sup>	15.5 <sup>†</sup>	18.5 <sup>†</sup>	22.0 <sup>†</sup>	<b>11.5<sup>†</sup></b>	<b>17.9<sup>†</sup></b>	<b>21.7<sup>†</sup></b>	26.4 <sup>†</sup>
<i>Other Baselines for Japanese</i>									
8	MultiMASS (Ja)	7.1	12.1	15.1	20.5	6.9	13.0	17.7	24.1
9	Deshuffling (Ja)	6.8	12.7	16.6	21.0	7.8	14.7	19.3	24.9
10	1 + 9	8.2	13.3	17.0	21.4	8.3	15.5	19.5	25.4
<i>Proposed methods for English</i>									
11	PMASS.P	6.8	12.1	15.9	20.7	5.5	13.5	17.8	24.5
11*	PMASS.S	6.5	12.3	16.2	21.2	6.2	13.5	18.2	24.6
12	HFSS	10.5 <sup>†</sup>	<b>16.3<sup>†</sup></b>	<b>18.9<sup>†</sup></b>	<b>22.6<sup>†</sup></b>	9.8 <sup>†</sup>	17.8 <sup>†</sup>	<b>21.7<sup>†</sup></b>	<b>26.8<sup>†</sup></b>
13	ENSS	<b>11.2<sup>†</sup></b>	<b>16.7<sup>†</sup></b>	<b>19.0<sup>†</sup></b>	22.1 <sup>†</sup>	<b>11.7<sup>†</sup></b>	<b>18.7<sup>†</sup></b>	<b>22.5<sup>†</sup></b>	<b>27.0<sup>†</sup></b>
<i>Other baselines for English</i>									
14	MultiMASS (En)	6.9	12.0	15.2	20.1	7.0	12.8	17.5	23.8
15	Deshuffling (En)	6.6	12.5	15.9	20.9	6.8	14.1	19.2	24.7
16	1 + 15	7.7	13.2	16.7	21.0	8.6	15.7	20.4	25.6
<i>Combination of methods for Japanese and English</i>									
17	4 + 13	<b>10.9<sup>†</sup></b>	<b>16.4<sup>†</sup></b>	18.7 <sup>†</sup>	<b>22.3<sup>†</sup></b>	<b>11.9<sup>†</sup></b>	<b>18.4<sup>†</sup></b>	<b>22.0<sup>†</sup></b>	26.5 <sup>†</sup>
18	10 + 16 (baseline)	7.2	12.6	16.4	20.9	8.4	14.8	19.1	25.5

Table 3. BLEU scores for simulated low/high-resource settings for Japanese–English ASPEC translation using from 3k to 50k parallel sentences for fine-tuning. Pre-trained models used for fine-tuning are numbered according to their description in Section 5.5. Results better than MASS with statistical significance  $p < 0.05$  are marked in <sup>†</sup>. Bold denotes the three top scores.

## 6.1 NMT Results

In Tables 3 and 4, where we simulate several low-resource settings for Japanese–English and Japanese–Chinese translations on ASPEC with different pre-training datasets; in Table 5 and 6, where we use realistic low-resource settings for Wikipedia Japanese–Chinese translation and News English–Korean translation, we observe that all settings using pre-training outperform those without pre-training (#0 & #0\*), which indicates the importance of pre-training. The results also indicate that JASS (#4) and ENSS (#13) are generally better than MASS (#1). With regard to two main baselines with pre-training, MASS and BART (text infilling), we observe that MASS

#	Model	Ja-Zh				Zh-Ja			
		3k	10k	20k	50k	3k	10k	20k	50k
<i>Main baselines</i>									
0	w/o pre, vanilla	0.7	3.4	11.5	21.0	1.9	4.5	16.0	28.2
0*	w/o pre, optimized	3.7	12.0	19.5	23.3	6.7	15.8	24.8	31.2
1	MASS	15.7	20.3	22.4	24.7	19.4	25.9	29.4	32.9
1*	BART (text infilling)	13.5	19.0	21.3	24.4	20.3 <sup>†</sup>	25.8	29.1	33.0
<i>Proposed methods</i>									
2	BMASS	16.7 <sup>†</sup>	21.1 <sup>†</sup>	23.0 <sup>†</sup>	25.3 <sup>†</sup>	20.9 <sup>†</sup>	27.2 <sup>†</sup>	30.2 <sup>†</sup>	33.7 <sup>†</sup>
3	BRSS	15.6	21.1 <sup>†</sup>	22.6	24.9	20.7 <sup>†</sup>	26.8 <sup>†</sup>	30.0 <sup>†</sup>	33.3 <sup>†</sup>
4	JASS	<b>17.1<sup>†</sup></b>	<b>22.2<sup>†</sup></b>	<b>23.2<sup>†</sup></b>	25.2 <sup>†</sup>	21.6 <sup>†</sup>	27.5 <sup>†</sup>	<b>30.4<sup>†</sup></b>	<b>33.6<sup>†</sup></b>
<i>Combinations of proposed methods with MASS</i>									
7	1 + 4	17.0 <sup>†</sup>	21.7 <sup>†</sup>	23.1 <sup>†</sup>	<b>25.4<sup>†</sup></b>	<b>21.8<sup>†</sup></b>	<b>27.6<sup>†</sup></b>	30.2 <sup>†</sup>	33.4 <sup>†</sup>
<i>Other baselines</i>									
8	MultiMASS	14.5	20.5	22.3	24.7	19.6	25.7	29.8	33.2
9	Deshuffling	14.1	19.5	21.6	24.3	18.4	25.0	28.7	32.8
10	1 + 9	15.0	20.2	22.1	25.0	18.9	25.9	29.3	33.1

Table 4. BLEU scores for simulated low-resource settings for Japanese–Chinese ASPEC translation using 3k to 50k parallel sentences for fine-tuning. Results better than MASS with statistical significance  $p < 0.05$  are marked in <sup>†</sup>.

#	Model	Ja-Zh				Zh-Ja			
		3k	10k	20k	50k	3k	10k	20k	50k
<i>Main baselines</i>									
0	w/o pre, vanilla	0.9	2.9	2.9	6.0	1.6	2.9	3.9	6.5
0*	w/o pre, optimized	3.3	7.8	11.7	21.9	6.7	12.0	16.2	24.2
1	MASS	7.7	15.4	18.3	23.4	9.6	17.6	23.3	27.1
1*	BART (text infilling)	5.9	14.0	18.0	21.8	8.7	17.8	24.2 <sup>†</sup>	28.5 <sup>†</sup>
<i>Proposed methods</i>									
2	BMASS	10.8 <sup>†</sup>	15.7	<b>20.1<sup>†</sup></b>	24.5 <sup>†</sup>	16.2 <sup>†</sup>	19.4 <sup>†</sup>	25.4 <sup>†</sup>	<b>30.0<sup>†</sup></b>
3	BRSS	11.6 <sup>†</sup>	16.2 <sup>†</sup>	20.0 <sup>†</sup>	24.6 <sup>†</sup>	15.7 <sup>†</sup>	21.6 <sup>†</sup>	25.0 <sup>†</sup>	28.3 <sup>†</sup>
4	JASS	<b>12.0<sup>†</sup></b>	<b>17.0<sup>†</sup></b>	<b>20.1<sup>†</sup></b>	<b>25.0<sup>†</sup></b>	<b>16.6<sup>†</sup></b>	21.2 <sup>†</sup>	<b>26.5<sup>†</sup></b>	29.2 <sup>†</sup>
<i>Combinations of proposed methods with MASS</i>									
7	1 + 4	11.8 <sup>†</sup>	16.8 <sup>†</sup>	<b>20.1<sup>†</sup></b>	24.6 <sup>†</sup>	<b>16.6<sup>†</sup></b>	<b>22.3<sup>†</sup></b>	25.5 <sup>†</sup>	29.6 <sup>†</sup>
<i>Other baselines</i>									
8	MultiMASS	8.2	13.8	18.6	21.5	10.7	17.3	22.0	26.4
9	Deshuffling	9.3	14.2	18.7	22.7	12.4	18.4	23.2	27.4
10	1 + 9	8.7	13.8	19.4	23.2	14.3	18.8	24.8	27.8

Table 5. BLEU scores for simulated low-resource settings for Japanese–Chinese Wikipedia translation using from 3k to 50k parallel sentences for fine-tuning. Results better than MASS with statistical significance  $p < 0.05$  are marked in <sup>†</sup>.

outperforms BART (text infilling) in most cases as shown in Table 3, 4, 5. So we focus on the comparisons with MASS in the following analyses. We also present the results by combing BART



#	Model	En-Ko		Ko-En	
		20k	94k	20k	94k
<i>Main baselines</i>					
0	w/o pre-training	1.3	2.1	2.9	4.5
0*	optimized transformer	2.1	3.7	3.9	8.3
1	MASS	2.9	4.5	5.6	9.6
<i>Proposed methods</i>					
2	PMASS	2.4	4.3	5.2	9.4
3	HFSS	3.1	4.8	7.7 <sup>†</sup>	10.3 <sup>†</sup>
4	ENSS (1 + 3)	<b>3.2</b>	<b>5.0<sup>†</sup></b>	6.8 <sup>†</sup>	<b>10.9<sup>†</sup></b>
<i>Other combinations</i>					
	2 + 3	3.0	4.7	7.0 <sup>†</sup>	10.6 <sup>†</sup>

Table 6. BLEU scores for simulated low-resource settings for English–Korean News translation using 20k and 94k parallel sentences for fine-tuning. Results better than MASS with statistical significance  $p < 0.05$  are marked in †. The BLEU scores are relatively low because English–Korean is a dissimilar language pair. Sennrich and Zhang [44] and Park et al. [32] reported similar BLEU results.

(text infilling) with ours in Appendix C.<sup>25</sup> Without pre-training, we observe that using optimized transformer (#0\*) benefits the low-resource setting, which has been proven by Araabi and Monz [1], Sennrich and Zhang [44]. However, pre-training can further improve the optimized baselines without pre-training.

Particularly, for the Japanese–English translation, BMASS (#2) is comparable to MASS; BRSS (#3 & #3(R)) and their combination, along with JASS (#5) are significantly better than MASS. However, as summarized in Tables 4 and 5, the results for two parallel corpora on different domains for Japanese–Chinese yield significantly better results when using our proposed BMASS and BRSS. We observe that only a few settings on Japanese-to-Chinese BRSS yield lower BLEU results than MASS, whereas other settings using the proposed methods yield better results than MASS by significant margins. Although MASS is better than BMASS for Japanese–English translation, the reverse can be observed for the Japanese–Chinese translation. This indicates that the effects of the proposed span-masking techniques might correlate with specific translation directions and domains. We suppose it is worth exploring the span-masking tricks that are non-sensitive to language pairs and domains in the future.

As summarized in Table 3 and 6, our proposed methods of leveraging linguistic knowledge for English yield significantly higher BLEU results when we perform the reordering pre-training task, HFSS (#12). However, the proposed linguistically-driven masked language model PMASS.P (#11) and PMASS.S (#11\*) yielded comparable results to several other baseline methods such as MultiMASS (#14) and deshuffling (#15). This demonstrates that the syntactical span-based masked language model may merely work on head-final languages such as Japanese.<sup>26</sup> Considering the weak performance of the PMASS, we combined HFSS with MASS for ENSS. The multi-task pre-trained ENSS yielded the highest results on almost all the low-resource settings. We will explore proper

<sup>25</sup>Note that better results from BART than MASS in Lewis et al. [21] are based on the multi-task objectives while we are comparing with the most effective single task within BART here.

<sup>26</sup>The weak performance of PMASS can also be attributed to the discrete nature of the remaining tokens (tokens that are not masked) without constituting complete semantic spans. We will attempt chunking-based masking for PMASS in future work to allow PMASS to be performed in a manner similar to BMASS.

chunking techniques for linguistically-driven span-masking pre-training for languages like English in the future.

However, in Table 3, when performing a universal linguistically-driven pre-training simultaneously for Japanese and English (#17), we did not achieve further significant BLEU improvements. This can be attributed to the increased dependence of NMT on specific linguistic information on a single language side, and the joint pre-training does not allow linguistic knowledge transfer across languages and between dissimilar languages.

In addition to the main baseline MASS, we also performed several other sequence-to-sequence pre-training baselines: MultiMASS (#8 & #14) and deshuffling (#9 & #15) along with their multi-task combinations (#10, #16 & #18) for Japanese and English. As summarized in Tables 3, 4, 5, and 6, we observe that the proposed masked style pre-training task, BMASS, and reordering pre-training tasks, BRSS & HFSS, outperform these baselines by significant margins, thereby indicating that linguistically-driven methods should be superior to self-supervised pre-training without leveraging linguistic features. Moreover, we investigated the percentages of the words of which the position changed. For Japanese pre-training, the percentages are 79.58% for BRSS and 94.72% for deshuffling. For English pre-training, the percentages are 91.97% for HFSS and 95.22% for deshuffling. Although there exists a gap for the percentages between BRSS and deshuffling, we can see that the percentages of deshuffling and HFSS are similar, which demonstrates that the quality of the linguistically generated reordered sentence is much more important than the percentage.

As summarized in Table 3, BRSS-F (English-order to Japanese-order) yielded slightly better results than BRSS-R (vice-versa); thus, we only experimented with BRSS-F for the remaining experiments. We suppose that the reason is that training the decoder with the original sentence is more important than training the encoder with it, which is also the reason why BART pre-training [21] treats the original sentence as the target sentence to be predicted from the decoder.<sup>27</sup> In other words, forcing the decoder to generate a natural sentence leads to a better initialized decoder for NMT. Meanwhile, HFSS pre-training is performed in an analogous manner for the same reason.

As mentioned above, JASS yields the best results when we consider only linguistically driven methods for Japanese. After combining the proposed methods for Japanese with MASS (#5~#7 in Table 3), we observe comparable results as compared to JASS by combining MASS and BRSS. This indicates the effects of combining masked style methods and reordering style methods. In Table 4 and 5, we believe that BMASS is better than MASS for combining with BRSS because of the significant improvements yielded by BMASS.

Moreover, as summarized in Tables 4 and 5, we observe that on the ASPEC domain, JASS improves up to 2.2 BLEU scores, whereas on the Wikipedia domain, JASS achieves up to 7.0 BLEU improvements. This demonstrates the promising performance of the proposed methods. Meanwhile, this indicates that the overlapping of pre-training domain with the fine-tuning domain is directly proportional to the realization of improvements by linguistically-driven pre-training methods.

Finally, by comparing the BLEU results in Table 3 with those reported by Mao et al. [24], we find that the BLEU scores of models pre-trained with News Crawl are better than those pre-trained with the Common Crawl monolingual corpus, which shows that pre-training with a high-quality monolingual dataset leads to superior fine-tuning results.

## 6.2 Adequacy Evaluation

Reference-free MT evaluation evaluates the translation system without using the target reference. Such an evaluation can help circumvent the noise existing in the references of translation targets. After the emergence of multilingual sentence encoders [2], Yankovskaya et al. [55] proposed the

<sup>27</sup>In BART, the original sentences without any noise are treated as target sentences.

#	Model	ASPEC		ASPEC		Wikipedia		News	
		Ja-En	En-Ja	Ja-Zh	Zh-Ja	Ja-Zh	Zh-Ja	En-Ko	Ko-En
*	Reference	80.78		86.10		87.26		73.93	
0	w/o pre-training	52.59	45.89	69.54	67.08	57.55	56.48	59.68	65.38
1	MASS	75.63	76.09	85.52	86.32	81.08	78.52	72.54	73.30
2	BMASS	75.75	76.68	85.42	86.49	80.91	81.36	-	-
3	BRSS	78.34	76.66	85.87	86.54	81.71	<b>84.29</b>	-	-
4	JASS	<b>80.00</b>	77.63	<b>85.96</b>	<b>86.58</b>	<b>85.39</b>	83.08	-	-
11	PMASS	76.08	73.67	-	-	-	-	71.90	74.14
12	HFSS	79.38	79.13	-	-	-	-	73.60	75.59
13	ENSS	79.79	<b>79.64</b>	-	-	-	-	<b>74.13</b>	<b>75.66</b>

Table 7. Adequacy scores evaluated by LASER embedding-based cosine similarity for ASPEC Japanese–English, Japanese–Chinese, Wikipedia Japanese–Chinese and News English–Korean translations, respectively, using 10k sentences for fine-tuning (using 94k sentences for English–Korean). Reference (\*) is the cosine similarity between test sets in two languages.

#	Model	BLEU		Adequacy		Fluency	
		Ja-Zh	Zh-Ja	Ja-Zh	Zh-Ja	Ja-Zh	Zh-Ja
0	w/o pre-training	2.9	2.9	1.22	1.05	3.90	3.99
1	MASS	15.4	17.6	2.72	2.33	4.11	4.09
2	BMASS	15.7	19.4	3.12	2.88	4.34	4.32
3	BRSS	16.2	<b>21.6</b>	3.30	3.35	4.30	<b>4.40</b>
4	JASS	<b>17.0</b>	21.2	<b>3.79</b>	<b>3.44</b>	<b>4.47</b>	4.36

Table 8. Adequacy and fluency of Wikipedia Japanese–Chinese translations using 10k sentences for fine-tuning.

use of multilingual sentence embeddings encoded by LASER to implement the reference-free MT evaluation. More precisely, we first apply LASER to encode the source sentence and the translated sentence, respectively. Thereafter, the cosine value of those two embeddings is used to evaluate the similarity between the source and translation. This cosine value is thus the metric used to evaluate translation adequacy. This approach has two advantages. The first advantage is that target references are not required, as mentioned above. The other advantage is that every two translation directions can be compared with each other because language-agnostic embedding is used for evaluation.

We report the adequacies in Table 7. First, we observe that methods with pre-training can yield more semantically correct translations than those without pre-training. Second, our proposed methods can significantly obtain higher LASER similarity scores than the MASS baseline, particularly the results on ASPEC Japanese–English, Wikipedia Chinese–Japanese and News English–Korean translations. Moreover, we can observe that the adequacy results obtained from the LASER embedding-based cosine similarity scores are consistent with the BLEU results.

#	Model	BLEU		Adequacy		Fluency	
		Ja-En	En-Ja	Ja-En	En-Ja	Ja-En	En-Ja
0	w/o pre-training	2.1	2.7	1.08	1.08	2.56	3.60
1	MASS	13.8	16.0	2.61	3.03	3.40	4.17
2	PMASS	12.1	13.5	2.40	2.76	3.24	4.07
3	HFSS	16.3	17.8	3.24	4.00	3.60	4.31
4	ENSS	<b>16.7</b>	<b>18.7</b>	<b>3.72</b>	<b>4.11</b>	<b>3.76</b>	<b>4.42</b>

Table 9. Adequacy and fluency of ASPEC Japanese–English translations using 10k sentences for fine-tuning.

### 6.3 Human Evaluation

Following Nakazawa et al. [29], we performed adequacy and fluency evaluations for the Japanese–Chinese and Japanese–English translations when 10k Wikipedia parallel sentences and 10k ASPEC parallel sentences were used for fine-tuning the pre-trained models. We randomly sampled 100 test-set English sentences and blindly evaluated their translations across various models. Each sentence was scored on a scale of 1 to 5, with 1 representing the worst score. The higher the score, the more adequate (meaningful) or fluent (well-formed) the sentence is. The final score was the average of the scores of 100 sentences. We did not consider the references, but only considered the sources for our evaluation.

In Table 8 and 9, we can observe that NMT models, even without pre-training, are capable of generating rather fluent sentences, and the lack of parallel sentences (low-resource scenario) will mainly influence the translation adequacy (refer to the extremely low adequacy of models without pre-training). Meanwhile, we can observe that our proposed BMASS, BRSS, JASS, HFSS, and ENSS result in large improvements in adequacy and moderate improvements in fluency, for both translation directions, whereas PMASS yielded marginal improvements. The improved performance of adequacy compared with that of MASS demonstrates the effectiveness of linguistically-driven pre-training methods. Moreover, we can observe that the results of human evaluation are almost consistent with those of BLEU.

### 6.4 Case Study

We conducted case studies on Japanese-to-English translation fine-tuned using 10k ASPEC parallel sentences and Chinese-to-Japanese translation fine-tuned using 10k Wikipedia parallel sentences to make improvements shown by BLEU score evaluations visible. As summarized in Tables 10 and 11, we find that the vanilla NMT system trained using 10k parallel sentences without pre-training can hardly implement the translation. With regard to models with pre-training, we observed that MASS and other baseline models generated several incorrect tokens in terms of semantics, whereas the entire sentence seemed fluent. However, our proposed methods can generate sentences with superior adequacy and fluency, where fewer missing keywords are observed.

### 6.5 Pre-training Accuracy

Pre-training accuracy is the accuracy of the monolingual pre-training tasks, and it can be an indicator of task complexity and pre-training objective performance. Tables 12 and 13 summarize the component-wise and overall pre-training accuracies for various models, respectively, on the ASPEC Japanese and English development set sentences. Regarding individual component methods, it can be observed that MASS and PMASS are the harder tasks, given their low accuracy, whereas

#	Reference-Ja	水の性質の多様性について、まず、水分子同士の間にも働く力である水素結合と、そのネットワーク構造について解説した。
	Reference-En	Various properties of water were explained on hydrogen bonds in which the force works among the water molecules and the network structure.
0	w/o pre-training	This paper introduces the outline of the development of the system, and it is described.
1	MASS	The network structure of the water, hydrogen combination as the power of the water, and the network structure are explained.
2	BMASS	On the basis of the water properties, hydrogen coupling and the network structure are explained in the first stage of water.
3	BRSS	On the formation of the water, this paper explains hydrogen bond and hydrogen bond, which is connected between the water vapor man fellows.
4	JASS	This paper explains the development of the properties of water and hydrogen combination, which is the power between the moisture man fellows and the network structure.
8	MultiMASS (Ja)	This paper explains the rich characteristics of the water and also explains the network structure of the hydrogen joining with the network structure.
9	Deshuffling (Ja)	This paper explains the potential of the water in the water, and the network structure that is connected between the water and hydrogen joining.
10	Multi-task baseline (Ja)	The active properties of water are explained, and hydrogen combination that is connected to the network structure and the power of the water are explained.
11	PMASS	The importance of the properties of water and the network structure, which is the active component of the water, are explained.
12	HFSS	The formation of the properties of water is first explained, then hydrogen combination and the network structure between the moisture man.
13	ENSS	The importance of the property of the water is first explained: hydrogen combination and the network structure, which is the power for the entire body of the water.
14	MultiMASS (En)	The growth of the water is explained, and the network structure and structure are explained through the hydrogen combination and network structure.
15	Deshuffling (En)	The network structure of the water properties is explained, and the network structure with hydrogen in the water is described.
16	Multi-task baseline (En)	This paper explains the growth of the water properties, and it also explains hydrogen bonding and its network structure with the ability to develop between the water molecules.

Table 10. Japanese-English translation examples fine-tuned using 10k ASPEC parallel sentences.

#	Reference-Ja	翌嘉永2年(1849年)11月22日に婚姻し、御廉中様(將軍世子の正室)と称された。
	Reference-Zh	翌嘉永2年(1849年)11月22日成婚、被称作御廉中样(将军世子的正室)。(He married on November 22nd in the second year of Kaei (1849), and was called Gorenju (Seishitsu of the Shogun's son).)
0	w/o pre-training	慶長2年(1583年)、劉宮將軍となり、幕府の命を率いた。(He became 劉宮 Shogun in the second year of Keicho, and lived the life of Shogun.)
1	MASS	翌嘉永2年(1849年)5月27日に成婚。(He married on May 27th in the second year of Kaei (1849).)
2	BMASS	翌嘉永2年(1849年)11月23日に結婚。(He married on November 23rd in the second year of Kaei (1849).)
3	BRSS	翌嘉永2年(1849年)11月23日に成婚し、御廉の中(將軍世子の正室)と呼ばれる。(He married on November 23rd in the second year of Kaei (1849), and was called Naka of Goren (Seishitsu of the Shogun's son).)
4	JASS	翌嘉永2年(1849年)11月22日成婚し、御廉中様(將軍世子の正室)と称した。(He married on November 22nd in the second year of Kaei (1849), and was called Gorenju (Seishitsu of the Shogun's son).)
8	MultiMASS	翌嘉永2年(1849年)1月29日に成婚した。(He married on January 29th in the second year of Kaei (1849).)
9	Deshuffling	翌嘉永2年(1849年)11月23日、御廉中の正室(將軍世子の正室)と称された。(He was called Seishitsu of Gorenju (Seishitsu of the Shogun's son) on November 22nd in the second year of Kaei (1849).)
10	Multi-task baseline	翌嘉永2年(1849年)11月22日成婚し、御廉の代わりに義子と呼ばれる。(He married on November 22nd in the second year of Kaei (1849), and was called Gishi instead of Goren.)

Table 11. Japanese-Chinese translation examples fine-tuned using 10k Wikipedia parallel sentences. Sentences in brackets correspond to English sentences of the above Japanese translations.

BRSS and HFSS are the easier tasks. Moreover, for Japanese, the accuracy of MASS and BRSS

#	Model	Overall	MASS	BMASS	BRSS
1	MASS	69.75	69.75	-	-
2	BMASS	77.32	-	77.32	-
3	BRSS	87.90	-	-	95.90
4	JASS	85.15	-	77.34	<b>97.89</b>
5	1 + 2	74.53	<b>70.17</b>	78.59	-
6	1 + 3	81.58	69.72	-	94.43
7	1 + 4	80.81	<b>70.22</b>	77.90	<b>97.73</b>

Table 12. Component-wise and overall pre-training accuracies on ASPEC Japanese development sentences. Column names “MASS,” “BMASS,” and “BRSS” denote the pre-training components in the respective model. Note the boost of BRSS accuracy in multitask settings, although the opposite could have been expected.

#	Model	Overall	MASS	PMASS	HFSS
1	MASS	70.97	70.97	-	-
2	PMASS	71.04	-	71.04	-
3	HFSS	96.48	-	-	96.48
4	ENSS	84.97	<b>71.24</b>	-	<b>98.05</b>

Table 13. Component-wise and overall pre-training accuracies on ASPEC English development sentences. Column names “MASS,” “PMASS,” and “HFSS” denote the pre-training components in the respective model. Note the boost of the HFSS accuracy in multitask settings, although the opposite could have been expected.

improves when coupled with BMASS, whereas for English, the accuracy of HFSS and MASS improves when they are combined with each other. Cross-referencing these accuracies with the BLEU scores in Table 3, we observe that an increase in BLEU scores has no significant relationship with the pre-training accuracy. However, masked language model-based pre-training methods (MASS & BMASS) seem to act as an accuracy improving catalyst for BRSS and HFSS, and this in turn has a positive impact on the translation quality.

One possible reason for this is that multi-task training of different pre-training methods helps boost the performance of individual methods. This is in accordance with several previous studies on multi-task training for NMT [9, 21, 23, 39]. Therefore, we recommend that such an analysis of multi-objective pre-training methods can help isolate the importance of individual pre-training objectives. Nevertheless, our analyses reveal that the components of JASS, BMASS, and BRSS, and the components of ENSS, MASS, and HFSS are certainly responsible for improving translation quality for Japanese-involved or English-language pairs.

## 6.6 Results in Middle/High-resource Scenarios

As summarized in Table 14, we report that BLEU leads to middle/high-resource scenarios. The fine-tuning is performed by more than 200k parallel sentences on the respective language pair and domain. By comparing with models without pre-training, we find that pre-training can still lead to some improvements, but much less than those in low-resource scenarios. Second, we observe that most pre-training methods obtained comparable BLEU results regardless of whether they were linguistically-driven methods or not. This indicates that in middle/high-resource scenarios, our

#	Model	Ja-En		En-Ja		Ja-Zh		Zh-Ja		Wiki	
		ASP 200k	ASP 1M	ASP 200k	ASP 1M	ASP 200k	ASP 672k	Wiki 258k	ASP 200k		ASP 672k
<i>Main baselines</i>											
0	w/o pre-training	26.1	27.5	33.4	35.8	27.0	31.2	24.6	36.7	42.4	30.4
1	MASS	26.5	28.8	33.7	37.6	27.2	<b>33.0</b>	29.4	36.7	<b>44.8</b>	34.6
<i>Proposed methods for Japanese</i>											
2	BMASS	26.5	<b>28.9</b>	33.8	37.8	<b>27.8</b>	32.6	29.8	36.4	44.7	<b>35.6</b>
3	BRSS	<b>26.8</b>	28.4	<b>34.0</b>	37.4	27.2	32.7	30.8	36.8	44.5	35.0
4	JASS	26.7	28.8	33.2	37.5	27.2	32.7	<b>31.1</b>	<b>37.4</b>	<b>44.8</b>	35.4
<i>Proposed methods for English</i>											
11	PMASS.P	26.3	28.0	33.5	37.0	-	-	-	-	-	-
11*	PMASS.S	26.2	28.9	33.2	37.8	-	-	-	-	-	-
12	HFSS	26.5	28.6	33.9	37.7	-	-	-	-	-	-
13	ENSS	26.3	28.8	33.9	<b>37.9</b>	-	-	-	-	-	-

Table 14. BLEU scores in middle/high-resource scenarios. “ASP” and “Wiki” denote ASPEC and Wikipedia parallel corpus, respectively.

proposed methods might be limited, which also shows that linguistically-driven supervision can be utilized to compensate for the lack of parallel sentences.

## 7 Conclusion

In this study, we proposed JASS and ENSS pre-training methods that leverage information from syntactic structures of sentences on the basis of language-agnostic pre-training schemes such as MASS for NMT. Our work leveraged abundant monolingual data and syntactic analysis such that the pre-training phase became aware of specific language structures. Our experiments on ASPEC Japanese–English, Japanese–Chinese, Wikipedia Japanese–Chinese, and News English–Korean translations demonstrated that JASS and ENSS outperform MASS and other language-agnostic pre-training methods in most low-resource settings. This demonstrates the importance of injecting language-specific information into the pre-training objective, as well as the benefit of multi-task pre-training with masked style and reordering objectives. Our adequacy evaluation through LASER, human evaluation, and case study also demonstrated that our methods resulted in a significant improvement in terms of the adequacy and fluency of translations. The analyses of pre-training accuracy reveal the complementary nature of individual tasks within JASS and ENSS.

Our future work will focus on implementing linguistic-aware multilingual pre-training using more languages for more robust pre-trained models. We also note that Raffel et al. [39] demonstrated that several NLP tasks such as text understanding can be reformulated as text-to-text tasks. This broadens the domain of usefulness of sequence-to-sequence pre-training tasks including ours, and we will be interested in evaluating our approach on various NLP tasks.

## Acknowledgments

We sincerely thank Dr. Raj Dabre, Dr. Fabien Cromieres, and Mr. Haiyue Song for their support and insightful comments on the JASS part of this work. This work was supported by Grant-in-Aid for Young Scientists #19K20343, JSPS and Information/AI/Data Science Doctoral Fellowship of Kyoto University.

## References

- [1] Ali Araabi and Christof Monz. 2020. Optimizing Transformer for Low-Resource Neural Machine Translation. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 3429–3435.
- [2] Mikel Artetxe and Holger Schwenk. 2019. Massively Multilingual Sentence Embeddings for Zero-Shot Cross-Lingual Transfer and Beyond. *Transactions of the Association for Computational Linguistics* 7 (March 2019), 597–610.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. San Diego, USA.
- [4] Chenhui Chu, Toshiaki Nakazawa, and Sadao Kurohashi. 2014. Constructing a Chinese–Japanese Parallel Corpus from Wikipedia. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*. European Language Resources Association (ELRA), Reykjavik, Iceland, 642–647.
- [5] Chenhui Chu, Toshiaki Nakazawa, and Sadao Kurohashi. 2016. Integrated Parallel Sentence and Fragment Extraction from Comparable Corpora: A Case Study on Chinese–Japanese Wikipedia. *ACM Trans. Asian Low Resour. Lang. Inf. Process.* 15, 2 (2016), 10:1–10:22.
- [6] Alexis Conneau and Guillaume Lample. 2019. Cross-lingual Language Model Pretraining. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*. 7057–7067.
- [7] Raj Dabre, Atsushi Fujita, and Chenhui Chu. 2019. Exploiting Multilingualism through Multistage Fine-Tuning for Low-Resource Neural Machine Translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 1410–1416.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186.
- [9] Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. Multi-Task Learning for Multiple Language Translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Beijing, China, 1723–1732.
- [10] Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding Back-Translation at Scale. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 489–500.
- [11] Dan Han, Katsuhito Sudoh, Xianchao Wu, Kevin Duh, Hajime Tsukada, and Masaaki Nagata. 2012. Head Finalization Reordering for Chinese-to-Japanese Machine Translation. In *Proceedings of the Sixth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Association for Computational Linguistics, Jeju, Republic of Korea, 57–66.
- [12] Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative Back-Translation for Neural Machine Translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*. Association for Computational Linguistics, Melbourne, Australia, 18–24.
- [13] Sho Hoshino, Yusuke Miyao, Katsuhito Sudoh, and Masaaki Nagata. 2013. Two-Stage Pre-ordering for Japanese-to-English Statistical Machine Translation. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*. Asian Federation of Natural Language Processing, Nagoya, Japan, 1062–1066.
- [14] Hideki Isozaki, Katsuhito Sudoh, Hajime Tsukada, and Kevin Duh. 2010. Head Finalization: A Simple Reordering Rule for SOV Languages. In *Proceedings of the Joint Fifth Workshop on Statistical Machine Translation and MetricsMATR*. Association for Computational Linguistics, Uppsala, Sweden, 244–251.
- [15] Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. Google’s Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation. *Transactions of the Association for Computational Linguistics* 5 (2017), 339–351.
- [16] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving Pre-training by Representing and Predicting Spans. *Transactions of the Association for Computational Linguistics* 8 (2020), 64–77.
- [17] Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. OpenNMT: Open-Source Toolkit for Neural Machine Translation. In *Proceedings of ACL 2017, System Demonstrations*. Association for Computational Linguistics, Vancouver, Canada, 67–72.
- [18] Philipp Koehn. 2004. Statistical Significance Tests for Machine Translation Evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Barcelona,



- Spain, 388–395.
- [19] Mamoru Komachi, Yuji Matsumoto, and Masaaki Nagata. 2006. Phrase reordering for statistical machine translation based on predicate-argument structure. In *2006 International Workshop on Spoken Language Translation, IWSLT 2006, Keihanna Science City, Kyoto, Japan, November 27-28, 2006*. 77–82.
  - [20] Sadao Kurohashi, Toshihisa Nakamura, Yuji Matsumoto, and Makoto Nagao. 1994. Improvements of Japanese morphological analyzer JUMAN. In *Proceedings of the International Workshop on Sharable Natural Language Resources*. 22–28.
  - [21] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 7871–7880.
  - [22] Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. 2020. Pre-training Multilingual Neural Machine Translation by Leveraging Alignment Information. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 2649–2663.
  - [23] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual Denoising Pre-training for Neural Machine Translation. *Transactions of the Association for Computational Linguistics* 8 (2020), 726–742.
  - [24] Zhuoyuan Mao, Fabien Cromieres, Raj Dabre, Haiyue Song, and Sadao Kurohashi. 2020. JASS: Japanese-specific Sequence to Sequence Pre-training for Neural Machine Translation. In *Proceedings of the 12th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 3683–3691.
  - [25] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Diamos, Erich Elsen, David García, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. Mixed Precision Training. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
  - [26] Hajime Morita, Daisuke Kawahara, and Sadao Kurohashi. 2015. Morphological Analysis for Unsegmented Languages using Recurrent Neural Network Language Model. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, 2292–2297.
  - [27] Rudra Murthy, Anoop Kunchukuttan, and Pushpak Bhattacharyya. 2019. Addressing word-order Divergence in Multilingual Neural Machine Translation for extremely Low Resource Languages. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 3868–3873.
  - [28] Toshiaki Nakazawa, Hideya Mino, Isao Goto, Graham Neubig, Sadao Kurohashi, and Eiichiro Sumita. 2015. Overview of the 2nd Workshop on Asian Translation. In *Proceedings of the 2nd Workshop on Asian Translation (WAT2015)*. Workshop on Asian Translation, Kyoto, Japan, 1–28.
  - [29] Toshiaki Nakazawa, Katsuhito Sudoh, Shohei Higashiyama, Chenchen Ding, Raj Dabre, Hideya Mino, Isao Goto, Win Pa Pa, Anoop Kunchukuttan, and Sadao Kurohashi. 2018. Overview of the 5th Workshop on Asian Translation. In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation: 5th Workshop on Asian Translation: 5th Workshop on Asian Translation*. Association for Computational Linguistics, Hong Kong.
  - [30] Toshiaki Nakazawa, Manabu Yaguchi, Kiyotaka Uchimoto, Masao Utiyama, Eiichiro Sumita, Sadao Kurohashi, and Hitoshi Isahara. 2016. ASPEC: Asian Scientific Paper Excerpt Corpus. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*. European Language Resources Association (ELRA), Portorož, Slovenia, 2204–2208.
  - [31] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318.
  - [32] Jungyeul Park, Jeon-Pyo Hong, and Jeong-Won Cha. 2016. Korean Language Resources for Everyone. In *Proceedings of the 30th Pacific Asia Conference on Language, Information and Computation: Oral Papers*. Seoul, South Korea, 49–58.
  - [33] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 2227–2237.
  - [34] Carl Pollard and Ivan A. Sag. 1988. *Information-Based Syntax and Semantics: Vol. 1: Fundamentals*. Center for the Study of Language and Information, USA.
  - [35] Carl Pollard and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. The University of Chicago Press, Chicago.
  - [36] Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. ProphetNet: Predicting Future N-gram for Sequence-to-Sequence Pre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 2401–2410.

- [37] Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. When and Why Are Pre-Trained Word Embeddings Useful for Neural Machine Translation?. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 529–535.
- [38] Alec Radford. 2018. Improving Language Understanding by Generative Pre-Training.
- [39] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.* 21 (2020), 140:1–140:67.
- [40] Shuo Ren, Yu Wu, Shujie Liu, Ming Zhou, and Shuai Ma. 2019. Explicit Cross-lingual Pre-training for Unsupervised Machine Translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 770–779.
- [41] Rico Sennrich and Barry Haddow. 2016. Linguistic Input Features Improve Neural Machine Translation. In *Proceedings of the First Conference on Machine Translation: Volume 1, Research Papers*. Association for Computational Linguistics, Berlin, Germany, 83–91.
- [42] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 86–96.
- [43] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1715–1725.
- [44] Rico Sennrich and Biao Zhang. 2019. Revisiting Low-Resource Neural Machine Translation: A Case Study. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, 211–221.
- [45] Aditya Siddhant, Ankur Bapna, Yuan Cao, Orhan Firat, Mia Chen, Sneha Kudugunta, Naveen Arivazhagan, and Yonghui Wu. 2020. Leveraging Monolingual Data with Self-Supervision for Multilingual Neural Machine Translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 2827–2835.
- [46] Haiyue Song, Raj Dabre, Zhuoyuan Mao, Fei Cheng, Sadao Kurohashi, and Eiichiro Sumita. 2020. Pre-training via Leveraging Assisting Languages for Neural Machine Translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*. Association for Computational Linguistics, Online, 279–285.
- [47] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. MASS: Masked Sequence to Sequence Pre-training for Language Generation. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*. 5926–5936.
- [48] Yu Sun, Shuohuan Wang, Yu-Kun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. 2020. ERNIE 2.0: A Continual Pre-Training Framework for Language Understanding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 8968–8975.
- [49] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Proceedings of the 27th Neural Information Processing Systems Conference (NIPS)*. Montréal, Canada, 3104–3112.
- [50] 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 *Proceedings of the 30th Neural Information Processing Systems Conference (NIPS)*. Long Beach, USA, 5998–6008.
- [51] Liang Wang, Wei Zhao, Ruoyu Jia, Sujian Li, and Jingming Liu. 2019. Denoising based Sequence-to-Sequence Pre-training for Text Generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 4003–4015.
- [52] Yiren Wang, ChengXiang Zhai, and Hany Hassan. 2020. Multi-task Learning for Multilingual Neural Machine Translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 1022–1034.
- [53] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. 5754–5764.

- [54] Zhen Yang, Bojie Hu, Ambyera Han, Shen Huang, and Qi Ju. 2020. CSP:Code-Switching Pre-training for Neural Machine Translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 2624–2636.
- [55] Elizaveta Yankovskaya, Andre Tättar, and Mark Fishel. 2019. Quality Estimation and Translation Metrics via Pre-trained Word and Sentence Embeddings. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*. Association for Computational Linguistics, Florence, Italy, 101–105.
- [56] Jiajun Zhang and Chengqing Zong. 2016. Exploiting Source-side Monolingual Data in Neural Machine Translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, 1535–1545.
- [57] Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020. Semantics-Aware BERT for Language Understanding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 9628–9635.
- [58] Chunting Zhou, Xuezhe Ma, Junjie Hu, and Graham Neubig. 2019. Handling Syntactic Divergence in Low-resource Machine Translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 1388–1394.
- [59] Junru Zhou, Zhuosheng Zhang, Hai Zhao, and Shuailiang Zhang. 2020. LIMIT-BERT : Linguistics Informed Multi-Task BERT. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 4450–4461.
- [60] Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer Learning for Low-Resource Neural Machine Translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, 1568–1575.

## A Algorithms for PMASS

In this section, we introduce Algorithms 1 and 2 for PMASS.S and PMASS.P respectively. We utilize the HPSG parsing result (Figure 5 (left)) to detect phrase spans to be masked. For PMASS.S, we can rapidly detect an entire phrase span to be masked. For PMASS.P, we start from the root of the HPSG parsing tree and stochastically mask the left child or the right child; then shift to the unmasked child node to find the next masking candidate. We implement this in a recursive manner.

## B Hyperparameters for Optimized Transformer

Following Araabi and Monz [1], we use the hyperparameter settings shown in Table 15 for training optimized transformer on different parallel data settings. Although optimized hyperparameter settings can significantly improve low-resource NMT, they require laborious grid search for the optimal setting while fine-tuning NMT based on pre-trained models do not.

## C Results of Combining BART with Ours

In Table 16, 17 and 18, we report the results of combining BART and our proposed methods for Japanese–English and Japanese–Chinese translations. We observe that BART (text infilling) can not further improve our proposed methods, which indicates that BART (text infilling) does not have a complement nature with our linguistically-driven multi-task pre-training methods.

---

**Algorithm 1:** Algorithm for determining masked phrase spans for PMASS.S.

---

**Input:** Length of the sentence  $L$ , tree of HPSG parsing result for the sentence  $T$ .

**Output:** Token List  $M$  consisting of all the tokens on  $N$ . (to be masked)

```
1 Initialize Current Node  $N$  by ROOT of  $T$ ;  
2 while number of tokens on  $N$  >  $\text{int}(L/2)$  do  
3   if number of tokens on left child of  $N$  > number of tokens on right child of  $N$  then  
4     |  $N \leftarrow$  left child of  $N$ ;  
5   else  
6     |  $N \leftarrow$  right child of  $N$ ;  
7   end  
8 end
```

---

**Algorithm 2:** Algorithm for determining masked phrase spans for PMASS.P.**Input:** Length of the sentence  $L$ , tree of HPSG parsing result for the sentence  $T$ .**Output:**  $P_{\text{mass}}(N=\text{ROOT of } T, L, l=0, M=\text{empty list})$  (tokens to be masked)

```

1 Function Pmass( $N, L, l, M$ ):
2   if tag of  $N$  is sentence then
3     | return Pmass(child of  $N, L, l, M$ )
4   else if tag of  $N$  is tok then
5     | if  $\text{int}(L/2) - l > 0$  then
6       |  $M.\text{append}(\text{token on } N)$ 
7     | return  $M$ 
8   else if  $N$  only has one child and  $N.\text{tag}$  is cons then
9     | return Pmass(child of  $N, L, l, M$ )
10  else
11    |  $ll \leftarrow$  number of tokens on the left child of  $N$ ;
12    |  $lr \leftarrow$  number of tokens on the right child of  $N$ ;
13    | if  $ll$  is 1 and  $lr$  is 1 then
14      | if  $\text{int}(L/2) - l > 1$  then
15        |  $M.\text{append}(\text{token on } N)$ 
16      | return  $M$ 
17    | else if  $\text{int}(L/2) \leq l$  then
18      | return  $M$ 
19    | else if  $ll \leq \text{int}(L/2) - l$  and  $lr > \text{int}(L/2) - l$  then
20      | if random  $p < 0.5$  then
21        |  $M \leftarrow M + \text{tokens on the left child of } N$ ;
22        |  $l \leftarrow l + ll$ ;
23        | return Pmass(right child of  $N, L, l, M$ )
24      | else
25        | return Pmass(right child of  $N, L, l, M$ )
26      | end
27    | else if  $lr \leq \text{int}(L/2) - l$  and  $ll > \text{int}(L/2) - l$  then
28      | if random  $p < 0.5$  then
29        |  $M \leftarrow M + \text{tokens on the right child of } N$ ;
30        |  $l \leftarrow l + lr$ ;
31        | return Pmass(left child of  $N, L, l, M$ )
32      | else
33        | return Pmass(left child of  $N, L, l, M$ )
34      | end
35    | else if  $ll > \text{int}(L/2) - l$  and  $lr > \text{int}(L/2) - l$  then
36      | if random  $p < 0.5$  then
37        | return Pmass(left child of  $N, L, l, M$ )
38      | else
39        | return Pmass(right child of  $N, L, l, M$ )
40      | end
41    | else
42      |  $M \leftarrow M + \text{tokens on the left child of } N$ ;
43      |  $l \leftarrow l + ll$ ;
44      | return Pmass(right child of  $N, L, l, M$ )
45 End Function

```

Hyperparameters	Default	3k	10k	20k	50k	94k
BPE operations	30k	5k	10k	10k	12k	15k
Encoder/decoder layers	6	2	2	2	2	2
Attention heads	16	2	2	2	2	2
Embedding dimension	1024	512	512	512	512	512
Feed forward dimension	4096	512	1024	1024	2048	2048
Dropout	0.3	0.3	0.3	0.3	0.3	0.3
Label smoothing	0.1	0.6	0.5	0.5	0.5	0.4
Batch size	4096	4096	4096	4096	4096	8192

Table 15. Hyperparameters for optimized transformer. “Default” denotes the setting of transformer-big. For English-Japanese, BPE operations for “Vanilla transformer-big” is 40k.

Model	Ja-En				En-Ja			
	3k	10k	20k	50k	3k	10k	20k	50k
MASS	8.8	13.8	17.2	21.2	9.1	16.0	20.6	25.0
ENSS	<b>11.2<sup>†</sup></b>	<b>16.7<sup>†</sup></b>	<b>19.0<sup>†</sup></b>	22.1 <sup>†</sup>	<b>11.7<sup>†</sup></b>	<b>18.7<sup>†</sup></b>	<b>22.5<sup>†</sup></b>	27.0 <sup>†</sup>
BART (text infilling)	3.1	11.1	15.5	20.7	5.6	14.9	19.8	25.6 <sup>†</sup>
BART + ENSS	10.7 <sup>†</sup>	15.8 <sup>†</sup>	18.9 <sup>†</sup>	<b>22.4<sup>†</sup></b>	10.6 <sup>†</sup>	17.6 <sup>†</sup>	21.2 <sup>†</sup>	<b>27.2<sup>†</sup></b>

Table 16. BLEU scores compared with BART for simulated low/high-resource settings for Japanese–English ASPEC translation using from 3k to 50k parallel sentences for fine-tuning. Results better than MASS with statistical significance  $p < 0.05$  are marked in <sup>†</sup>.

Model	Ja-Zh				Zh-Ja			
	3k	10k	20k	50k	3k	10k	20k	50k
MASS	15.7	20.3	22.4	24.7	19.4	25.9	29.4	32.9
JASS	<b>17.1<sup>†</sup></b>	<b>22.2<sup>†</sup></b>	<b>23.2<sup>†</sup></b>	<b>25.2<sup>†</sup></b>	21.6 <sup>†</sup>	<b>27.5<sup>†</sup></b>	<b>30.4<sup>†</sup></b>	<b>33.6<sup>†</sup></b>
BART (text infilling)	13.5	19.0	21.3	24.4	20.3 <sup>†</sup>	25.8	29.1	33.0
BART + JASS	<b>17.1<sup>†</sup></b>	21.3 <sup>†</sup>	23.1 <sup>†</sup>	25.0	<b>21.9<sup>†</sup></b>	<b>27.5<sup>†</sup></b>	<b>30.4<sup>†</sup></b>	<b>33.6<sup>†</sup></b>

Table 17. BLEU scores compared with BART for simulated low-resource settings for Japanese–Chinese ASPEC translation using 3k to 50k parallel sentences for fine-tuning. Results better than MASS with statistical significance  $p < 0.05$  are marked in <sup>†</sup>.

Model	Ja-Zh				Zh-Ja			
	3k	10k	20k	50k	3k	10k	20k	50k
MASS	7.7	15.4	18.3	23.4	9.6	17.6	23.3	27.1
JASS	<b>12.0</b> <sup>†</sup>	<b>17.0</b> <sup>†</sup>	<b>20.1</b> <sup>†</sup>	<b>25.0</b> <sup>†</sup>	<b>16.6</b> <sup>†</sup>	21.2 <sup>†</sup>	<b>26.5</b> <sup>†</sup>	29.2 <sup>†</sup>
BART (text infilling)	5.9	14.0	18.0	21.8	8.7	17.8	24.2 <sup>†</sup>	28.5 <sup>†</sup>
BART + JASS	11.4 <sup>†</sup>	16.5 <sup>†</sup>	19.4 <sup>†</sup>	24.3 <sup>†</sup>	16.2 <sup>†</sup>	<b>22.5</b> <sup>†</sup>	26.2 <sup>†</sup>	<b>30.0</b> <sup>†</sup>

Table 18. BLEU scores compared with BART for simulated low-resource settings for Japanese–Chinese Wikipedia translation using from 3k to 50k parallel sentences for fine-tuning. Results better than MASS with statistical significance  $p < 0.05$  are marked in †.