

# Implicit Word Reordering with Knowledge Distillation for Cross-Lingual Dependency Parsing

Zhuoran Li, Chunming Hu, Junfan Chen, Zhijun Chen, Richong Zhang\*

SKLSDE, Beihang University, Beijing, China

{lizhuoranget, hucm, zhijunchen}@buaa.edu.cn, {chenjif, zhangrc}@act.buaa.edu.cn

## Abstract

Word order difference between source and target languages is a major obstacle to cross-lingual transfer, especially in the dependency parsing task. Current works are mostly based on order-agnostic models or word reordering to mitigate this problem. However, such methods either do not leverage grammatical information naturally contained in word order or are computationally expensive as the permutation space grows exponentially with the sentence length. Moreover, the reordered source sentence with an unnatural word order may be a form of noising that harms the model learning. To this end, we propose an Implicit Word Reordering framework with Knowledge Distillation (**IWR-KD**). This framework is inspired by that deep networks are good at learning feature linearization corresponding to meaningful data transformation, e.g. word reordering. To realize this idea, we introduce a knowledge distillation framework composed of a word-reordering teacher model and a dependency parsing student model. We verify our proposed method on Universal Dependency Treebanks across 31 different languages and show it outperforms a series of competitors, together with experimental analysis to illustrate how our method works towards training a robust parser.

## Introduction

Dependency parsing is a fundamental task which aims to extract the low-level grammatical relationships between words in a sentence (McDonald 2006; Nivre 2008; Kiperwasser and Goldberg 2016), such as subject-verb relationships. Recently, cross-lingual dependency parsing has attracted considerable attention from academic and industrial communities, for which a parser is trained on a source language and directly applied in the target language of interest. Multilingual pre-trained language models (mPLMs) demonstrate exceptional performance in cross-lingual dependency parsing (Devlin et al. 2019; Conneau et al. 2020). However, these mPLMs inevitably encode word order features to model contextual representations, referred to as *order-sensitive* (Ahmad et al. 2019; Liu et al. 2021). Since word order is inherently different across languages, there is a risk of over-fitting into the word order of the source language that could hurt the performance in the target languages.

\*Corresponding author.

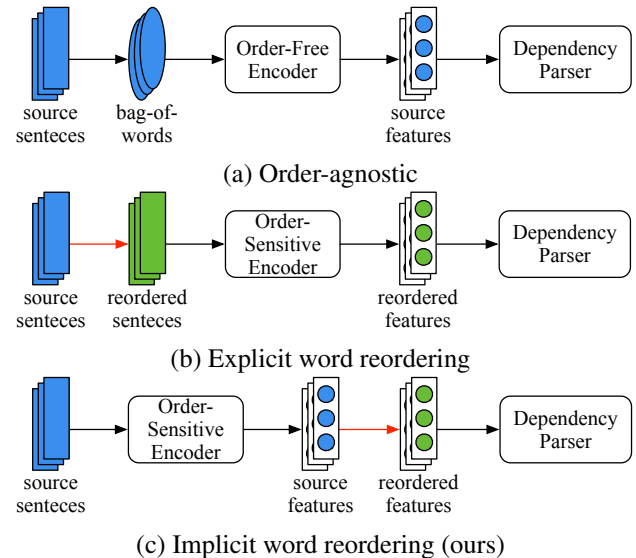


Figure 1: Comparison between different methods for Word Ordering Difference in cross-lingual dependency parsing. (a) Removing word order information. (b) Permuting the words in a source sentence to resemble the word order of a target language. (c) Our method adapts the word order in the feature space. Red arrows indicate reordering steps.

Therefore, a lot of works pay attention to the word order difference problem in cross-lingual dependency parsing. The methods mainly fall into two categories: First, as shown in Figure 1 (a), order-agnostic based methods, which utilize order-free encoders, e.g. Self-Attention (Ahmad et al. 2019), Frozen Position Embedding (Liu et al. 2021), Bag-of-Words (Ji et al. 2021), to maintain robust performance to the change of word order. Second, as shown in Figure 1 (b), word reordering-based methods (Rasooli and Collins 2019; Liu et al. 2020; Arviv et al. 2023), which generally first generate a large set of new sentences syntactically similar to the target language via rearranging words in the source sentence, and then select generated high-quality sentences to train a target-language dependency parser. Since these word reordering methods essentially rearrange words in the source language, we call this category of methods as *explicit word*

reordering (EWR).

However, there are limitations to the previous methods. The order-agnostic methods could weaken its representation capability due to the lack of word order information, resulting in a drop in dependency parsing. The EWR methods can be computationally expensive, as the permutation space grows exponentially with the sentence length. Moreover, the explicit word rearrangement may introduce linguistic adversity, i.e. unnatural word order in the source language, and therefore can be seen as a form of noising that harms the model learning (Wei et al. 2021; Arviv et al. 2023).

To address the limitations, we explore a new attempt, *implicit word reordering* and propose an Implicit Word Reordering method with Knowledge Distillation (IWR-KD) for cross-lingual dependency parsing. Motivated by that deep learning are surprisingly good at features linearization (Bengio et al. 2013; Upchurch et al. 2017; Wang et al. 2019), IWR-KD learns to implicitly adapt the word order relationship in the word representation space rather than really permuting the words given a sentence. Figure 1 (c) shows the differences between the proposed IWR-KD method and the typical order-agnostic and EWR-based methods.

Specifically, we train a word order model using the target-language part-of-speech (POS) data<sup>1</sup>, which is much easier to annotate than the parsing tree and often reflects the syntactic structure of a language (Liu et al. 2020). This trained target-language word order prediction model is then used as a teacher model to decide the new order between two POS tags. Finally, we train a student model to mimic the order prediction of the teacher and the dependency parsing using the labelled source language training data. Once the student model is ready, it can not only on-the-fly generate representations that correspond to the target-language word order relationship, but also preserve the essential linguistic structure of the original input, thus avoiding the mentioned limitations of previous works.

In summary, the contributions of this paper are as follows:

- We propose an implicit word reordering method for cross-lingual dependency parsing, addressing the limitations of previous approaches related to word order representations and high computational costs.
- We introduce a word reordering teacher model to effectively incorporate the target word order knowledge into the dependency parsing student model.
- We conduct extensive experiments demonstrating that IWR-KD outperforms several competing baselines on Universal Dependency Treebanks across 31 languages.

## Preliminaries

In this section, we first briefly review the cross-lingual dependency parsing task, backbone and prior word reordering work. These are the foundations applied to our approach.

<sup>1</sup>In this paper, we assume that our corpora are annotated with gold POS tags, even in target languages that lack gold training trees. While this is an idealized setting, it is sufficient to achieve good results. This approach has often been used in studies on unsupervised and cross-lingual transfer. Our ablation study explores a potential alternative that does not rely on gold tags.





Figure 2: An example of an English sentence that is explicitly reordered to resemble to the Estonian syntactic order.

**Task Description** Dependency parsing is the task of creating the dependency tree for an input sentence, which is a directed graph and defines the grammatical relation between dependent words (e.g. *Mary*) and their heads (e.g. *prepared*), as shown in Figure 2 (a). The goal of cross-lingual dependency parsing is to train a parser on the source language and perform well on an unseen target language. Throughout this work, we systematically study the transferability of the proposed implicit word reordering between source and target languages with different word order distances.

**Backbone: Biaffine Dependency Parser** Following (Ahmad et al. 2019; Wu and Dredze 2019; Arviv et al. 2023), we adopt the graph-based bi-affine dependency parsing model (Dozat and Manning 2018) as the backbone of our parsers, which is composed of four layers, i.e., embedding layer, Bi-LSTM layer, MLP layer and scorer layer, as shown in the right side of Figure 3. Following (Ahmad et al. 2019; Liu et al. 2021), all the parsers take words as well as their gold part-of-speech (POS) tags as input. Formally, given the input sequence  $s$  with  $L$  words  $\{w_1, w_2, \dots, w_L\}$  and its POS tags  $\{p_1, p_2, \dots, p_L\}$ , the embedding layer creates a sequence of input embeddings  $e_{1:L}$  in which each  $e_i$  is a concatenation of its word embedding ( $e_{w_i}$ ) and POS embedding ( $e_{p_i}$ ). The POS embedding is trained from scratch, while the word embedding is initialized with the pre-trained multilingual language model, such as mBERT (Devlin et al. 2019).

$$e_i = e_{w_i} \oplus e_{p_i} \quad (1)$$

To further introduce the contextual information, we then encode each input embedding by a multilayer bidirectional LSTM:

$$r_i = \text{BiLSTM}(e_{1:L}, i) \quad (2)$$

Two dimension-reducing MLPs are then used to specialise each recurrent representation  $r_i$  into head-word and dependent-word representations for both the edge and the label prediction.

$$h_i^{\text{edge-head}} = \text{MLP}^{\text{edge-head}}(r_i) \quad (3)$$

$$h_i^{\text{edge-dep}} = \text{MLP}^{\text{edge-dep}}(r_i) \quad (4)$$

$$h_i^{\text{label-head}} = \text{MLP}^{\text{label-head}}(r_i) \quad (5)$$

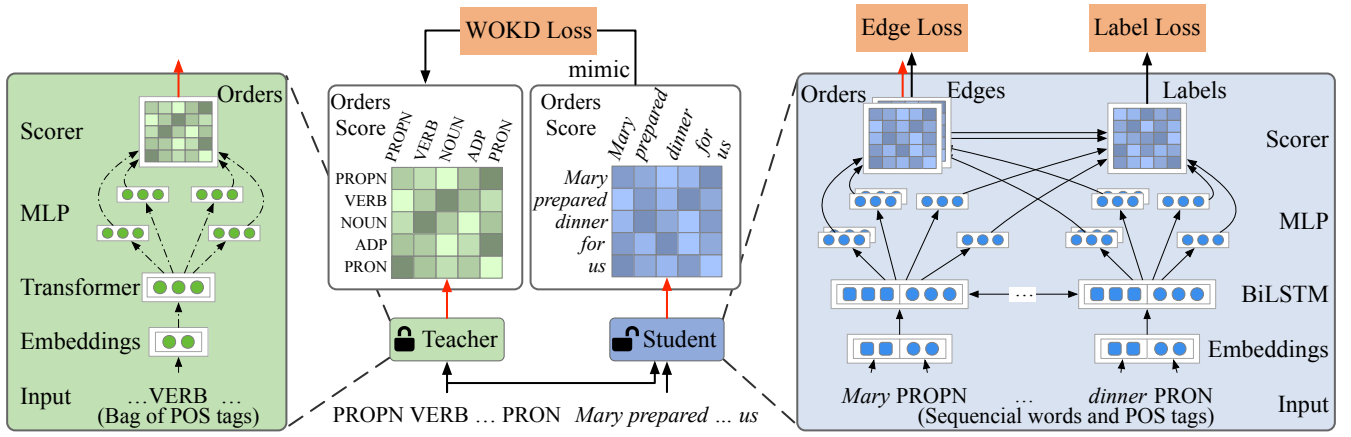


Figure 3: An overview of our IWR-KD: (i) Word Reordering Teacher decides the new direction between a dependent word and its head. (ii) Dependency Parsing Student is supervised by the teacher and the gold dependency parsing labels simultaneously.

$$h_i^{\text{label-dep}} = \text{MLP}^{\text{label-dep}}(r_i) \quad (6)$$

Following (Dozat and Manning 2018), we then use biaffine classifiers to compute the edge attention scores and label attention scores. These scores can be decoded into a graph by keeping only edges that received a positive score.

$$s_{i \rightarrow j}^{\text{edge}} = h_j^{\text{edge-dep}} U^{\text{edge}} h_i^{\text{edge-head(T)}} + b^{\text{edge}} \quad (7)$$

$$s_{i \rightarrow j}^{\text{label}} = h_j^{\text{label-dep}} U^{\text{label}} h_i^{\text{label-head(T)}} + b^{\text{label}} \quad (8)$$

$$\hat{y}_{i \rightarrow j}^{\text{edge}} = \{s_{i \rightarrow j}^{\text{edge}} \geq 0\} \quad (9)$$

$$\hat{y}_{i \rightarrow j}^{\text{label}} = \arg \max s_{i \rightarrow j}^{\text{label}} \quad (10)$$

where  $U$  and  $b$  are linear transformation and bias term, respectively. For each position pair  $i, j$ , a binary cross-entropy loss is used for the existence of edge  $i \rightarrow j$ , and a cross-entropy loss is used for the labels of gold edges.

$$\mathcal{L}^{\text{edge}} = \text{BCELoss}(y_{i \rightarrow j}^{\text{edge}}, \hat{y}_{i \rightarrow j}^{\text{edge}}) \quad (11)$$

$$\mathcal{L}^{\text{label}} = \text{CELoss}(y_{i \rightarrow j}^{\text{label}}, \hat{y}_{i \rightarrow j}^{\text{label}}) \quad (12)$$

The training object for parsing is the summing of above two losses.

$$\mathcal{L} = \mathcal{L}^{\text{edge}} + \lambda_1 \mathcal{L}^{\text{label}} \quad (13)$$

**Prior Explicit Word Reordering** Given a sentence  $s = \{w_1, w_2, \dots, w_L\}$  in the source language, prior explicit word reordering (EWR) works (Liu et al. 2021; Arviv et al. 2023; Rasooli and Collins 2019) aim to permute the words in it to syntactically more similar to the order of the target language. Then the reordered sentences, denoted as  $s' = \{w'_1, w'_2, \dots, w'_L\}$ , as shown in Figure 2 (b), used to train the cross-lingual dependency parser.

## Method: Implicit Word Reordering

To mitigate the shortcomings of the existing EWR approach, we propose an **Implicit Word Reordering** algorithm via **Knowledge Distillation (IWR-KD)**. Unlike conventional EWR where words really are rearranged, our IWR-KD

adapts word order relations in the feature space and integrates this procedure into the training of cross-lingual dependency parser.

The overall structure of the proposed IWR-KD is shown in Figure 3, which consists of two important components: a *word reordering teacher* aims to decide the new word order given the source input, and a *dependency parsing student* not only learns dependency parsing from the labelled source input but also mimics target word order predictions of the teacher.

## Word Reordering Teacher Training

The goal of the word reordering teacher model is to decide the word order of each dependency according to the POS tags of the dependent words and their heads. Given the target language sentence  $\{w_1, w_2, \dots, w_L\}$  and their POS tags  $\{p_1, p_2, \dots, p_L\}$ , the following steps are applied to train the word reordering teacher:

1. **Head-Words Finding.** For each word  $w_i$  in the sentence, we predict its head word  $w_j$  using a pre-trained parser in the source language.
2. **Training Instances Extraction.** We create a training instance for each pairwise POS tags  $\langle p_i, p_j \rangle$  of dependent-head words  $\langle w_i, w_j \rangle$ . The word order label  $y_{i \rightarrow j}^{\text{order}} \in \{0, 1\}$  is made based on their positions in the original sentence, where  $y_{i \rightarrow j}^{\text{order}} = 0$  indicates the POS tag  $p_i$  is on the left of its head  $p_j$ , while  $y_{i \rightarrow j}^{\text{order}} = 1$  indicates  $p_i$  is on the right of  $p_j$  in the original sentence.
3. **Teacher Training.** As shown in the left-side in Figure 3, the teacher network consists of four layers. Given a sequence of POS tag  $\{p_1, p_2, \dots, p_L\}$ , we first obtain its POS embeddings  $\{e_{p_1}, e_{p_2}, \dots, e_{p_L}\}$ . We only use the universal POS tags to discard the word form discrepancy. We then introduce contextual information by an order-free Transformer layer, which could improve cross-lingual generalization compared with the BiLSTM layer.

$$z_{p_i} = \text{Transformer}(e_{p_i}) \quad (14)$$

We obtain the head and dependent representation for each  $z_{p_i}$  through two dimension-reducing MLPs.

$$h_i^{\text{order-head}} = \text{MLP}^{\text{order-head}}(z_{p_i}) \quad (15)$$

$$h_i^{\text{order-dep}} = \text{MLP}^{\text{order-dep}}(z_{p_i}) \quad (16)$$

We compute word order for the given pairwise POS tags having head-dependent edges. A binary cross-entropy loss is used for the word order learning.

$$s_{i \rightarrow j}^{\text{order-tea}} = h_j^{\text{order-dep}} U^{\text{order}} h_i^{\text{order-head}(T)} + b^{\text{order}} \quad (17)$$

$$\hat{y}_{i \rightarrow j}^{\text{order-tea}} = \arg \max s_{i \rightarrow j}^{\text{order-tea}} \quad (18)$$

$$\mathcal{L}^{\text{order}} = \text{BCELoss}(y_{i \rightarrow j}^{\text{order}}, \hat{y}_{i \rightarrow j}^{\text{order-tea}}) \quad (19)$$

After the teacher has well-trained in the target language, we take the teacher as guidance to train the dependency parsing student.

### Dependency Parsing Student Learning

The target of the student model is to parse a sentence based on the transformed features corresponding to the reordered source sentence. As illustrated in the right side of Figure 3, we implement the student model by adding a word order learning network based on the backbone. The student model is trained to mimic the prediction probability distribution of word order generated by the teacher model on the POS tags in the target language. This process aims to transfer word order knowledge from the teacher model to the student model while allowing the student model to leverage language-specific word ordering knowledge available in the target-language POS tags.

Here, the student model consists of four layers. Specifically, given a source input sentence  $s$  with  $L$  words  $\{w_1, w_2, \dots, w_L\}$  and its POS (part-of-speech) tags  $\{p_1, p_2, \dots, p_L\}$ . We first obtain its representation sequence  $\{r_1, r_2, \dots, r_L\}$  through the embedding layer and BiLSTM layer, as shown in Equation 1-2.

We then use Equation 3-6 and 15-16 to specialize each recurrent representation  $r_i$  into head-word  $h_{1:L}^{\text{head}}$  and dependent-word  $h_{1:L}^{\text{dep}}$  representations using two dimension-reduction MLPs for edge, label and word order prediction, respectively. In addition to calculating the edge prediction and label prediction according to the Equation 7-10, we also calculate the word order prediction as follows:

$$s_{i \rightarrow j}^{\text{order-stu}} = h_j^{\text{order-dep}} U^{\text{order}} h_i^{\text{order-head}(T)} + b^{\text{order}} \quad (20)$$

$$\hat{y}_{i \rightarrow j}^{\text{order-stu}} = \arg \max s_{i \rightarrow j}^{\text{order-stu}} \quad (21)$$

The student model learns dependency parsing by the following two losses:

$$\mathcal{L}^{\text{edge}} = \text{BCELoss}(y_{i \rightarrow j}^{\text{edge}}, \hat{y}_{i \rightarrow j}^{\text{edge-stu}}) \quad (22)$$

$$\mathcal{L}^{\text{label}} = \text{CELoss}(y_{i \rightarrow j}^{\text{label}}, \hat{y}_{i \rightarrow j}^{\text{label-stu}}) \quad (23)$$

The word order knowledge distillation learning loss is formulated as the mean squared error loss:

$$\mathcal{L}^{\text{order}} = \text{MSELoss}(\hat{y}_{i \rightarrow j}^{\text{order-tea}}, \hat{y}_{i \rightarrow j}^{\text{order-stu}}) \quad (24)$$

And the whole student training loss is the summation of three losses:

$$\mathcal{L} = \mathcal{L}^{\text{edge}} + \lambda_1 \mathcal{L}^{\text{label}} + \lambda_2 \mathcal{L}^{\text{order}} \quad (25)$$

## Experiments

In this section, we conduct extensive experiments on 31 languages across a broad spectrum of language families to validate the effectiveness and reasonableness of our proposed IWR-KD method for cross-lingual dependency parsing.

### Setup

**Datasets** Following the setup of (Ahmad et al. 2019), we conduct experiments on Universal Dependencies (UD) Treebanks (v2.14) (Zeman et al. 2024), in which 31 different languages are selected for evaluation. In our main experiments, we take English as the source language and 30 other languages as the target ones. We only use the source language for both training and hyper-parameter tuning.

**Performance Metric** The evaluation metrics are unlabeled attachment score (UAS) and labeled attachment score (LAS). Each experiment is conducted in three runs with different random seeds and the average scores are reported.

**Implement Details** We employ mBERT (Devlin et al. 2019) to derive the cross-lingual word embeddings. Since the mBERT embeddings are subword-level, we follow previous work (Wu and Dredze 2019; Liu et al. 2021) in taking the first subword as word-level embeddings. The maximum subword sequence length is 512. We train the POS embeddings from scratch and set the dimension size of POS embeddings as 50. The batch size is set to 32. We use Adam (Kingma and Ba 2015) to train models with  $\beta_1$  of 0.9,  $\beta_2$  of 0.9 and  $L_2$  weight decay of  $1e-5$ . The model is trained for 50 epochs with the learning rate of  $1e-5$  for mBERT and  $3e-5$  for other network layers. We choose the best hyper-parameters according to the development set in the source language. We empirically set  $\lambda_1 = 1, \lambda_2 = 0.001$ . We implement our method using PyTorch 1.8.0 based on *Hugging Face* transformer library<sup>2</sup>.

To quantify the *word order distance* between two languages, we following (Ahmad et al. 2019) select the 52 most frequently occurring dependency triples (Dep  $\curvearrowright$  Head, Label) across 31 languages, then concatenate relative frequency of the left-direction (dependent word before its head) for all triples in each language as the word order feature. We use Manhattan distance as our word order distance. Due to differences in the versions of the datasets used, our word order distance does not strictly align with previous word order distance (Ahmad et al. 2019)<sup>3</sup>.

**Baseline Parsers** We compare our IWR-KD with several competing baselines involving word order learning as follows. *SelfAtt-Direct* (Ahmad et al. 2019) adopts the self-attention based order-free encoder for cross-lingual parsing. *mBERT-Direct* (Wu and Dredze 2019) is fine-tuned by adding the graph-based biaffine dependency parser on top of it. *Frozen PE* (Liu et al. 2021) freezes the positional embedding in mBERT during the fine-tuning. *Subtree-EWR* (Arviv et al. 2023) explicitly reorders words through subtree-aware constraints for cross-lingual parsing. *WOL* is a variant of ours, which learns the word order knowledge in the

<sup>2</sup><https://huggingface.co>

<sup>3</sup>Using order typology distance to distinguish it from ours

Lang.	Dist.	SelfAttn	mBERT	Frozen PE	Subtree-EWR	WOL	IWR-KD (ours)
en	0.00	90.4/88.4	92.4/90.3	92.2/90.3 ↓	92.4/90.3	<u>93.0/91.5</u>	92.5/90.5
no	0.06	80.8/72.8	86.7/77.4	87.0/78.0	84.5/76.6 ↓	88.3/78.8	87.9/78.8
sv	0.07	81.0/73.2	84.3/76.7	84.2/76.4 ↓	81.8/74.8 ↓	84.9/77.3	<u>85.3/78.1</u>
fr	0.09	77.9/72.8	83.4/71.7	83.4/72.0	84.7/73.5	84.4/72.5	<u>85.2/73.4</u>
pt	0.09	76.6/67.8	81.2/73.0	80.4/72.4 ↓	81.4/73.0	82.6/74.6	<u>83.1/75.4</u>
da	0.10	76.6/67.9	81.7/72.6	81.3/72.8 ↓	80.0/71.6 ↓	82.7/73.2	<u>83.0/73.8</u>
es	0.12	74.5/66.4	80.5/71.1	79.1/70.1 ↓	79.8/70.3 ↓	81.6/72.1	<u>82.3/72.7</u>
it	0.12	80.8/75.8	84.3/77.7	83.6/77.2 ↓	84.9/78.9	85.6/79.5	<u>85.9/79.5</u>
ca	0.13	73.8/65.1	80.5/70.4	79.1/69.4 ↓	81.6/72.1	81.5/71.5	<u>82.3/71.9</u>
hr	0.13	61.9/52.9	78.0/64.4	77.7/64.2 ↓	77.7/65.3	77.0/63.1 ↓	<u>78.2/65.0</u>
pl	0.13	74.6/62.2	88.4/75.7	87.2/75.4 ↓	88.3/76.8	88.8/75.9	88.8/76.6
sl	0.13	68.2/56.5	79.2/66.2	77.5/64.9 ↓	77.1/64.4 ↓	79.2/65.0 ↓	<u>79.7/67.2</u>
uk	0.13	60.1/52.3	73.9/63.0	75.8/64.5	73.8/62.3 ↓	75.8/63.8	<u>77.1/65.4</u>
bg	0.14	79.4/68.2	85.3/73.9	85.6/74.4	76.8/66.7 ↓	87.0/74.8	86.3/74.4
cs	0.14	63.1/53.8	78.3/64.2	77.7/64.0 ↓	78.4/65.0	77.5/63.0 ↓	<u>78.5/64.5</u>
de	0.14	71.3/61.6	77.9/69.3	78.7/70.3	80.8/73.0	78.7/68.6	79.9/70.4
he	0.14	55.3/48.0	68.0/53.8	68.3/54.2	69.7/55.1	68.3/54.5	<u>70.6/55.0</u>
nl	0.14	68.6/60.3	79.7/71.3	79.4/71.2 ↓	79.8/71.9	79.3/70.6 ↓	<u>80.7/72.2</u>
ru	0.14	60.6/51.6	77.3/65.8	76.0/64.9 ↓	75.6/65.6 ↓	77.3/65.9	<u>77.8/66.1</u>
ro	0.15	65.1/54.1	78.2/64.6	78.5/66.0	79.1/66.0	79.2/65.2	76.1/63.2 ↓
id	0.17	49.2/43.5	60.4/48.9	62.1/50.3	70.5/58.4	60.3/49.1	62.2/50.7
sk	0.17	66.7/58.2	82.9/69.4	83.4/70.9	82.4/69.8 ↓	81.8/67.8 ↓	<u>84.5/70.6</u>
lv	0.18	70.8/49.3	78.2/56.9	78.0/56.9 ↓	72.3/54.1 ↓	78.8/56.9	<u>79.2/57.8</u>
et	0.20	65.7/44.9	73.7/52.2	73.8/52.4	71.0/51.7 ↓	74.3/52.2	<u>74.7/53.0</u>
fi	0.20	66.3/48.7	74.5/53.7	75.4/54.9	71.1/53.5 ↓	76.1/55.5	75.9/55.4
zh	0.23	42.5/25.1	55.7/37.5	56.0/37.6	60.0/40.7	57.2/37.8	55.9/37.8
ar	0.26	38.1/28.0	44.8/33.1	47.1/35.5	54.1/41.1	46.5/34.9	48.5/36.2
la	0.28	48.0/35.2	57.8/41.4	57.0/40.6 ↓	57.9/42.5	57.0/40.9 ↓	56.6/40.8 ↓
ko	0.33	34.5/16.4	44.8/25.1	45.0/25.1	47.9/29.7	43.8/23.1 ↓	47.7/26.6
hi	0.40	35.5/26.5	41.4/26.8	42.5/28.0	41.0/27.0 ↓	40.5/27.1 ↓	<u>43.3/28.2</u>
ja	0.49	28.2/20.9	25.2/16.6	27.1/17.8	26.4/18.4	25.5/16.6	<u>28.8/20.6</u>
AVG ±STD	0.17	64.1/53.8 /	72.9/60.5 ± 0.4/0.1	72.9/60.7 ± 0.3/0.2	73.0/61.3 ± 0.2/0.2	73.4/60.8 ± 0.3/0.4	<u>74.1/61.7</u> ± 0.3/0.2

Table 1: Cross-lingual dependency parsing results by language (UAS%/LAS%). We order the languages by order typology distances (Ahmad et al. 2019) to English. We use "↓" to indicate results below mBERT and use underlined text to highlight the best performance.

source language with an auxiliary loss. For fair comparison, all methods are implemented with the same datasets and training configurations. We report the results of *SelfAtt-Direct* in the original paper and re-implement other methods. Note that we do not compare ours with methods that use extra resources, e.g., *SubDP* (Shi, Gimpel, and Livescu 2022) utilizes translation and word alignment to augment data, *SFDP* (Sun, Li, and Zhao 2023) employs large amounts of unlabelled data to improve the performance.

## Main Results

Table 2 presents the results on the test sets. The languages are ordered by their order typology distance to English. From the experimental results, we can make the following observations. (1) Our proposed IWR-KD method achieves state-of-the-art performance on most languages and average performance across all languages, demonstrating the ef-

Model	UAS%/LAS%
<b>IWR-KD (Ours)</b>	<b>74.1/61.7</b>
w/ Pseudo-Labelling	73.5 <sub>-0.6</sub> /61.3 <sub>-0.4</sub>
w/ Silver POS Tags	73.4 <sub>-0.7</sub> /60.9 <sub>-0.8</sub>
w/o FT	72.6 <sub>-1.5</sub> /60.3 <sub>-1.4</sub>

Table 2: Ablation study (averaged results over all languages).

fectiveness and generalization of IWR-KD. (2) The order-sensitive direct transfer method (mBERT) significantly surpasses the traditional order-agnostic self-attention method (SelfAttn) by effectively capturing richer contextual information, including word order. (3) The performance of the mBERT-based word order-agnostic model (Frozen PE) de-

(Dep $\curvearrowright$ Head, Label)	Word Order Frequency	Source/Target Examples
(PRON $\curvearrowright$ VERB, obj)	EN: 0.1351 EN+IWR-KD: 0.9348 FR: 0.9109	EN: Recently I 'm having trouble <span style="background-color: #e0ffe0;">training<sup>head</sup><sub>VERB</sub></span> <span style="background-color: #ffe0e0;">him<sup>dep</sup><sub>PRON</sub></span> . FR: ... se sont permis de <span style="background-color: #ffe0e0;">les<sup>dep</sup><sub>PRON</sub></span> <span style="background-color: #e0ffe0;">traiter<sup>head</sup><sub>VERB</sub></span> en tant ...
(NOUN $\curvearrowright$ VERB, obl)	EN: 0.0996 EN+IWR-KD: 0.5055 DE: 0.6031	EN: ... be <span style="background-color: #e0ffe0;">causing<sup>head</sup><sub>VERB</sub></span> us trouble for <span style="background-color: #ffe0e0;">years<sup>dep</sup><sub>NOUN</sub></span> to come. DE: ... französischen <span style="background-color: #ffe0e0;">Radsportverband<sup>dep</sup><sub>NOUN</sub></span> <span style="background-color: #e0ffe0;">gezahlt<sup>head</sup><sub>VERB</sub></span> ...
(NOUN $\curvearrowright$ NOUN, nmod)	EN: 0.0056 EN+IWR-KD: 0.4867 ET: 0.8296	EN: Hope there is useful <span style="background-color: #e0ffe0;">grist<sup>head</sup><sub>NOUN</sub></span> for the <span style="background-color: #ffe0e0;">mill<sup>dep</sup><sub>NOUN</sub></span> here . ET: ... <span style="background-color: #ffe0e0;">tänapäeva<sup>dep</sup><sub>NOUN</sub></span> ühe olulisema ameerika <span style="background-color: #e0ffe0;">luuletaja<sup>head</sup><sub>NOUN</sub></span>

Table 3: Case study on Implicit Word Reordering. Word order frequency indicates the relative frequency of the left direction (dependent before its head). The GREEN (RED) highlight indicates a head (dependent) word.

clines relative to direct transfer (mBERT) in both the source language and similar languages, such as English (en) and Norwegian (no), due to underfitting caused by the frozen position representation. (4) The explicit word reordering method (Subtree-EWR) substantially declines in some languages, such as Slovenian (sl) and Finnish (fi), due to unnatural sentence rearrangements that introduce noise detrimental to model learning. Note that the EWR method indirectly accesses the target language dependency annotation information (Arviv et al. 2023), so it can achieve high performance in some languages. While IWR-KD relies solely on low-cost target language POS tags, it outperforms SubtreeWR in average performance. (5) WOL performs better in languages closer to English, while mBERT outperforms it in languages that are more "distant" from English. This indicates capturing word order is critical in cross-lingual dependency parsing.

### Ablation Study

To verify the effect of each part of our method, we do experiments with the following variants of our IWR-KD. (1) *w/ pseudo-labelling*, which converts the soft word order probability distribution into hard labels to guide the student model training. Hard labels result in performance degradation, highlighting that our word order distillation loss conveys richer word order knowledge than hard labels. (2) *w/ Silver UPOS*, which employs the Stanza<sup>4</sup> tool to annotate POS tags to simulate a more realistic application scenario. In this case, our method can improve the performance of the direct transfer method (mBERT in Table 2). (3) *w/o FT*, which eliminates the fine-tuning of mPLMs to make the model more lightweight. However, performance drops significantly without fine-tuning, indicating that fine-tuning enables the encoder to effectively adapt to the specific task.

### Case Study

In this section, we conduct a series of case studies to bring up insights into why the proposed IWR-KD works.

<sup>4</sup><https://stanfordnlp.github.io/stanza>

The IWR-KD method can help the parser correctly identify dependency labels by exploiting the learned order relationship between dependent words and their heads. Specifically, if the model can align the word order relationships of the source language to the target language, misidentified dependency labels in the target language may be corrected. As shown in Table ??, in the first example (PRON  $\curvearrowright$  VERB, obj), the relative frequency of PRON before VERB differs greatly between the source and target languages, which leads to migration failure. Our model's prediction of the relative frequency of PRON before VERB in the source language is close to that of the target language. At the same time, our model learned that their label is obj on the source language training set, and then the model can correctly predict PRON and PRON in the target language. The dependency tag of VERB is obj. Two additional examples present the same results using different languages.

### Performance versus Word Order Distance

As shown in the upper subfigure of Figure 4, transfer ability decreases as the word order distance increases from left to right, indicating that word order is a crucial factor for transfer performance (Perf.). The Pearson correlation coefficient between word order distance and transfer performance is -0.8803 (p-value = 1e-10), demonstrating a strong negative correlation. As shown in the lower subfigure of Figure 4, the reduction (Redn.) in word order distance positively influences performance improvement (Impr.). The Pearson correlation coefficient is 0.4308 (p-value = 0.0175), indicating a strong positive correlation.

### Effect of Different Teacher Models

This section investigates the effect of the different word reordering teacher models. We compare three learning methods: (1) *Rand*: Randomly learning the word order relationship between pairs of words. (2) *Heur*: Heuristically learning the ordering relationships of pairs belonging to 52 selected dependency triples. (3) *Our*: Based on a pre-trained parser in the source language, we predict the head of each




Figure 4: Word order distance and performance. Languages (x-axis) are sorted by their order typology distances (Ahmad et al. 2019) to English from left to right.

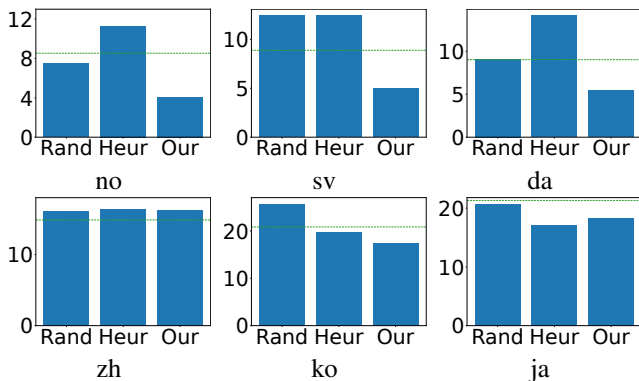


Figure 5: Word order distance predicted by different word reordering teachers. The green bar indicates original word order distance between English and the target language.

word and learn their ordering relationships.

As shown in Figure 5, we calculate the word order distance predicted by three different teacher models on six different target languages. Three languages, i.e. no, sv, da, are closer to the source language, while the other three languages, i.e. zh, ko, ja, are far away from the source language. It can be seen that our transfer-based teacher model can effectively reduce the word order distance between the source language and target languages. We notice that our method fails to reduce the word order distance between the source language and zh (Chinese), but it improves the transfer performance as in Table 2. Through more thorough observation, we find that our method reduces the word order difference on some dependency triples. For example, for the word order frequency of (NOUN  $\curvearrowright$  NOUN, nmod), EN: 0.0056, ZH: 0.9980, EN+IWR-KD: 0.5413. We conjecture that different dependencies have different effects on transfer performance, which may be one of our future research directions.

## Related Work

Word order refers to the arrangement of words in a sentence or phrase to convey meaning in a particular language. In the cross-lingual community, word order is widely explored from two perspectives: order-agnostic methods (Ahmad et al. 2019; Liu et al. 2021; Ding, Wang, and Tao 2020; Chen, Zhang, and Fu 2023; Si et al. 2019; Hessel and Schofield 2021) and word reordering methods (Rasooli and Collins 2019; Ji et al. 2021; Liu et al. 2020; Chen et al. 2019; Goyal and Durrett 2020; Al-Negheimish, Madhyastha, and Russo 2023; Pham et al. 2021). The former argues that word order encoding is a risk for cross-lingual transfer, as the model often fits in the language-specific order. Reducing the word order information fitted into the models can improve the cross-lingual adaptation performance in position representation (Ding, Wang, and Tao 2020) and dependency parsing (Ahmad et al. 2019; Liu et al. 2021). The latter is dedicated to reordering the word from one language to another. For example, Arviv et al. (Arviv et al. 2023) achieves better cross-lingual transfer results by rearranging the words in the source language to meet the word order in the target language conditioned on the syntactic constraints.

## Conclusion

In this paper, we explore a new attempt, *implicit word reordering*, and propose an implicit word reordering method with knowledge distillation (IWR-KD), which uses the word reordering model as a teacher to guide the student model to simultaneously learn target-language word order knowledge and source-language dependency parsing knowledge on the labeled data of the source language. IWR-KD can effectively capture vital word order information and does not require the actual generation and selection of re-ordered sentences, which addresses the limitations of previous order-agnostic and explicit word reordering methods. We conduct extensive experiments on 31 different languages on Universal Dependency Treebanks, which showed that IWR-KD outperforms multiple competitive methods.

## Acknowledgements

We thank Beijing Advanced Innovation Center for Big Data and Brain Computing for providing computation resources. We also thank the anonymous reviewers and the area chair for their insightful comments. This work was supported by the National Natural Science Foundation of China (No. U2433212, No. 62306026), in part by the China Postdoctoral Science Foundation (No. 2023M740184), in part by the Fundamental Research Funds for the Central Universities, and in part by the State Key Laboratory of Complex & Critical Software Environment.

## References

- Ahmad, W.; Zhang, Z.; Ma, X.; Hovy, E.; Chang, K.-W.; and Peng, N. 2019. On Difficulties of Cross-Lingual Transfer with Order Differences: A Case Study on Dependency Parsing. In Burstein, J.; Doran, C.; and Solorio, T., eds., *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)*, 2440–2452. Minneapolis, Minnesota: Association for Computational Linguistics.
- Al-Negheimish, H.; Madhyastha, P.; and Russo, A. 2023. Towards preserving word order importance through Forced Invalidation. In Vlachos, A.; and Augenstein, I., eds., *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 2563–2570. Dubrovnik, Croatia: Association for Computational Linguistics.
- Arviv, O.; Nikolaev, D.; Karidi, T.; and Abend, O. 2023. Improving Cross-lingual Transfer through Subtree-aware Word Reordering. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*, 718–736. Singapore: Association for Computational Linguistics.
- Bengio, Y.; Mesnil, G.; Dauphin, Y.; and Rifai, S. 2013. Better Mixing via Deep Representations. In Dasgupta, S.; and McAllester, D., eds., *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, 552–560. Atlanta, Georgia, USA: PMLR.
- Chen, K.; Wang, R.; Utiyama, M.; and Sumita, E. 2019. Neural Machine Translation with Reordering Embeddings. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1787–1799. Florence, Italy: Association for Computational Linguistics.
- Chen, X.; Zhang, G.; and Fu, J. 2023. TPDM: Selectively Removing Positional Information for Zero-shot Translation via Token-Level Position Disentangle Module. *CoRR*, abs/2305.19857.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8440–8451. Online: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *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)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Ding, L.; Wang, L.; and Tao, D. 2020. Self-Attention with Cross-Lingual Position Representation. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 1679–1685. Online: Association for Computational Linguistics.
- Dozat, T.; and Manning, C. D. 2018. Simpler but More Accurate Semantic Dependency Parsing. In Gurevych, I.; and Miyao, Y., eds., *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 484–490. Melbourne, Australia: Association for Computational Linguistics.
- Goyal, T.; and Durrett, G. 2020. Neural Syntactic Preordering for Controlled Paraphrase Generation. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 238–252. Online: Association for Computational Linguistics.
- Hessel, J.; and Schofield, A. 2021. How effective is BERT without word ordering? Implications for language understanding and data privacy. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 204–211. Online: Association for Computational Linguistics.
- Ji, T.; Jiang, Y.; Wang, T.; Huang, Z.; Huang, F.; Wu, Y.; and Wang, X. 2021. Word Reordering for Zero-shot Cross-lingual Structured Prediction. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 4109–4120. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kiperwasser, E.; and Goldberg, Y. 2016. Simple and Accurate Dependency Parsing Using Bidirectional LSTM Feature Representations. *Transactions of the Association for Computational Linguistics*, 4: 313–327.
- Liu, L.; Zhou, Y.; Xu, J.; Zheng, X.; Chang, K.-W.; and Huang, X. 2020. Cross-Lingual Dependency Parsing by POS-Guided Word Reordering. In Cohn, T.; He, Y.; and

- Liu, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2938–2948. Online: Association for Computational Linguistics.
- Liu, Z.; Winata, G. I.; Cahyawijaya, S.; Madotto, A.; Lin, Z.; and Fung, P. 2021. On the Importance of Word Order Information in Cross-lingual Sequence Labeling. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15): 13461–13469.
- McDonald, R. 2006. *Discriminative learning and spanning tree algorithms for dependency parsing*. Ph.D. thesis, USA. AAI3225503.
- Nivre, J. 2008. Algorithms for Deterministic Incremental Dependency Parsing. *Computational Linguistics*, 34(4): 513–553.
- Pham, T.; Bui, T.; Mai, L.; and Nguyen, A. 2021. Out of Order: How important is the sequential order of words in a sentence in Natural Language Understanding tasks? In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 1145–1160. Online: Association for Computational Linguistics.
- Rasooli, M. S.; and Collins, M. 2019. Low-Resource Syntactic Transfer with Unsupervised Source Reordering. In Burstein, J.; Doran, C.; and Solorio, T., eds., *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)*, 3845–3856. Minneapolis, Minnesota: Association for Computational Linguistics.
- Shi, F.; Gimpel, K.; and Livescu, K. 2022. Substructure Distribution Projection for Zero-Shot Cross-Lingual Dependency Parsing. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6547–6563. Dublin, Ireland: Association for Computational Linguistics.
- Si, C.; Wang, S.; Kan, M.; and Jiang, J. 2019. What does BERT Learn from Multiple-Choice Reading Comprehension Datasets? *CoRR*, abs/1910.12391.
- Sun, K.; Li, Z.; and Zhao, H. 2023. Cross-Lingual Universal Dependency Parsing Only From One Monolingual Treebank. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11): 13393–13407.
- Upchurch, P.; Gardner, J. R.; Pleiss, G.; Pless, R.; Snively, N.; Bala, K.; and Weinberger, K. Q. 2017. Deep Feature Interpolation for Image Content Changes. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 6090–6099. IEEE Computer Society.
- Wang, Y.; Pan, X.; Song, S.; Zhang, H.; Huang, G.; and Wu, C. 2019. Implicit Semantic Data Augmentation for Deep Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 12635–12644.
- Wei, J.; Huang, C.; Vosoughi, S.; Cheng, Y.; and Xu, S. 2021. Few-Shot Text Classification with Triplet Networks, Data Augmentation, and Curriculum Learning. In Toutanova, K.; Rumshisky, A.; Zettlemoyer, L.; Hakkani-Tur, D.; Beltagy, I.; Bethard, S.; Cotterell, R.; Chakraborty, T.; and Zhou, Y., eds., *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 5493–5500. Online: Association for Computational Linguistics.
- Wu, S.; and Dredze, M. 2019. Beto, Bentz, Becas: The Surprising Cross-Lingual Effectiveness of BERT. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 833–844. Hong Kong, China: Association for Computational Linguistics.
- Zeman, D.; Nivre, J.; Abrams, M.; Ackermann, E.; Aepli, N.; Aghaei, H.; Agić, Ž.; Ahmadi, A.; Ahrenberg, L.; Ajede, C. K.; et al. 2024. Universal Dependencies 2.14. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.