

# Error Diversity Matters: An Error-Resistant Ensemble Method for Unsupervised Dependency Parsing

Behzad Shayegh<sup>1</sup>, Hobie H.-B. Lee<sup>1</sup>, Xiaodan Zhu<sup>2</sup>, Jackie Chi Kit Cheung<sup>3,4</sup>, Lili Mou<sup>1,4</sup>

<sup>1</sup>Dept. Computing Science, Alberta Machine Intelligence Institute (Amii), University of Alberta

<sup>2</sup>Dept. Electrical and Computer Engineering & Ingenuity Labs Research Institute, Queen’s University

<sup>3</sup>Quebec Artificial Intelligence Institute (Mila), McGill University

<sup>4</sup>Canada CIFAR AI Chair

the.shayegh@gmail.com, hobie.lee@ualberta.ca,

xiaodan.zhu@queensu.ca, jcheung@cs.mcgill.ca, doublepower.mou@gmail.com

## Abstract

We address unsupervised dependency parsing by building an ensemble of diverse existing models through post hoc aggregation of their output dependency parse structures. We observe that these ensembles often suffer from low robustness against weak ensemble components due to error accumulation. To tackle this problem, we propose an efficient ensemble-selection approach that considers error diversity and avoids error accumulation. Results demonstrate that our approach outperforms each individual model as well as previous ensemble techniques. Additionally, our experiments show that the proposed ensemble-selection method significantly enhances the performance and robustness of our ensemble, surpassing previously proposed strategies, which have not accounted for error diversity.

**Code** — <https://github.com/MANGA-UOFA/ED4UDP>

**Extended version** — <https://arxiv.org/abs/2412.11543>

## 1 Introduction

Syntactic parsing, a fundamental task in natural language processing (NLP), refers to identifying grammatical structures in text (Zhang 2020), which can help downstream NLP tasks such as developing explainable models (Amara, Sevastjanova, and El-Assady 2024). Unsupervised syntactic parsing is particularly beneficial for processing languages or domains with limited resources, as it eliminates the need for human-annotated training data (Kann et al. 2019). In addition, autonomously discovering patterns and structures helps to provide evidence to test linguistic theories (Goldsmith 2001) and cognitive models (Bod 2009). Two common types of syntactic parsing include: (1) constituency parsing, organizing words and phrases in a sentence into a hierarchy (Chomsky 1967), and (2) dependency parsing, predicting dependence links between words in a sentence (Tesnière, Osborne, and Kahane 2015). The latter is particularly interesting given its possible applications to other fields, e.g., RNA structure prediction (Wang and Cohen 2025).

Unsupervised dependency parsing has seen diverse approaches over decades (Smith and Eisner 2005; Cai, Jiang, and Tu 2017; Han et al. 2020). Yang et al. (2020) show




Figure 1: (a) Performances of individuals. (b) Performances of ensembles by adding individuals in (a) from left to right. Numbers are UASs on the WSJ test set, reported by the mean and the standard deviation through five runs. The experiment setup is detailed in §4.1.

that the combination of two different models can outperform both individuals, suggesting the varied contributions of different models. We follow our previous work on constituency parsing (Shayegh et al. 2024; Shayegh, Wen, and Mou 2024) and build an ensemble of different unsupervised dependency parsers to leverage their diverse expertise.<sup>1</sup> We regard the ensemble output as the dependency structure that maximizes its overall similarity to all the *individual* models’ outputs. For the similarity metric, we employ the unlabeled attachment score (UAS; Nivre and Fang 2017), which is the most widely used evaluation metric of the task.

In our work, we observe that a naïve application of Shayegh et al. (2024)’s ensemble is sensitive to weak individuals in dependency parsing. As shown in Figure 1, a best-to-worst incremental ensemble experiment encounters a significant drop in performance when weaker individuals are added. This drop is due to the accumulation of individual errors, which arises from a low error diversity.

It is important to distinguish between two types of diversity: *expertise diversity* and *error diversity*. The former refers to the phenomenon that different models excel on different subsets of samples, while the latter indicates that models are wrong in different ways for each sample. Although both types of diversity are crucial for successful en-

<sup>1</sup>Expertise diversity may be due to different inductive biases, objective functions, data features (Rokach 2009), and/or different training subsets (Breiman 1996; Wen et al. 2023, 2025). In our work, we build an ensemble of different models exploring various structures, predictors, and training methods.

sembles (Zhang and Ma 2012), it is essential to emphasize that expertise diversity serves as a motivation to build an ensemble, whereas error diversity is a requirement for its success (outperforming ensemble individuals), because otherwise, the individuals’ errors cannot be eliminated through the ensemble process.

To this end, we propose a diversity-aware method for ensemble selection considering error diversity so that our approach is resistant to error accumulation. Having a diverse set of ensemble components has been a focus of researchers for years (Tang, Suganthan, and Yao 2006; Zhou 2012; Wu et al. 2021; Purucker et al. 2023). However, most previous studies focus on expertise diversity by considering only the individuals’ successes/failures, ignoring the differences in their mistakes (Yule and Pearson 1900; Fleiss 1981; Cunningham and Carney 2000). By contrast, we develop a metric, called *society entropy*, to capture both aspects of diversity. Our approach outperforms the existing expertise-diversity metrics in ensemble selection, demonstrating the importance of error diversity.

We conduct our experiments on the Wall Street Journal (WSJ) corpus in the Penn Treebank (Marcus, Santorini, and Marcinkiewicz 1993). Results show the effectiveness of our ensemble approach, outperforming all the individual parsers and ensemble baselines. We further boost the performance by employing our diversity-aware selection of individuals, achieving state-of-the-art performance in the task.

In conclusion, our main contributions include:

1. proposing an ensemble-based approach to unsupervised dependency parsing,
2. specifying error diversity as an important aspect in building an ensemble, and
3. utilizing society entropy as a diversity metric to perform selection of ensemble individuals.

## 2 Related Work

### 2.1 Unsupervised Dependency parsing

Over the years, researchers have proposed various models to tackle unsupervised dependency parsing under different setups (Le and Zuidema 2015; He, Neubig, and Berg-Kirkpatrick 2018; Shen et al. 2021). The longest line of research began with Klein and Manning (2004) introducing the dependency model with valence (DMV), a probabilistic generative model of a sentence and its dependency structure, which is extended by Headden III, Johnson, and McClosky (2009) to include categories of valence and lexical information. In the deep learning era, Jiang, Han, and Tu (2016) utilize neural networks to learn probability distributions in DMV. Han, Jiang, and Tu (2017, 2019a,b) further equip neural DMVs with lexical and contextual information. More recently, Yang et al. (2020) introduce two second-order variants of DMV, which incorporate grandparent–child or sibling information, respectively. In this work, we build ensembles of different models to leverage their diverse knowledge.

All of the above work focuses on projective dependency parse structures, meaning that all the dependency arcs can be drawn on one side of the sentence without crossing each

other. The projectivity constraint comes from the context-free nature of human languages (Chomsky 1956; Gaifman 1965). Given that most English sentences are characterized by projective structures (Wu et al. 2020), we also follow previous studies and focus on this type of parses.

Unsupervised dependency parsing typically considers unlabeled structures, i.e., dependency arcs are not categorized by their linguistic types (e.g., subject or object of a verb). We also follow the previous line of research and focus on unlabeled structures.

### 2.2 Ensemble-Based Dependency Parsing

Che et al. (2018) and Lim et al. (2018) propose to build ensembles of dependency parsers to smooth out their noise and reduce the sensitivity to the initialization of the neural networks. They use the average of networks’ softmax outputs for prediction. This approach is restricted as it requires the models to have the same output space, which does not hold in our scenario as we aim to leverage the knowledge of diverse models with different architectures.

Kulkarni et al. (2022) and Shayegh et al. (2024) show the effectiveness of post hoc aggregation methods for ensemble-based supervised and unsupervised constituency parsing. Kulkarni, Eulenstein, and Li (2024) borrow different post hoc aggregation methods from graph studies, including the maximum spanning tree (Gavril 1987), conflict resolution on heterogeneous data (Li et al. 2014), and a customized Ising model (CIM; Ravikumar, Wainwright, and Lafferty 2010), and employ them for ensemble-based dependency parsing. They compare the performance of these aggregation methods and show the superiority of CIM. However, this method does not ensure the validity of the ensemble output as a projective dependency parse structure. Nevertheless, we include CIM as a baseline which underperforms when applied to our unsupervised setting, performing worse than the best individual.

In this work, we develop a dependency-structure aggregation method, based on which we further build an ensemble of unsupervised dependency parsers with different designs.

### 2.3 Ensemble Selection

Ensemble selection refers to selecting a set of models to form an ensemble, which has been extensively discussed in the machine learning literature (Kuncheva 2004; Caruana, Munson, and Niculescu-Mizil 2006). Caruana et al. (2004) introduce a forward-stepwise-selection method, inspired by the field of feature selection (Liu and Motoda 2007); in their approach, individuals are incrementally added to the ensemble based on the performance boost on a validation set. Such an approach is prone to overfitting to the validation set (Kohavi and John 1997); in addition, it is a time-consuming process to build the ensemble, as we need to build many ensembles at each increment.

Another line of research predicts the success of an ensemble by looking at the individuals’ properties (Ganaie et al. 2022; Mienye and Sun 2022). One commonly adopted criterion is to keep diversity among the ensemble individuals (Minku, White, and Yao 2010; Wood et al. 2023). This

brings multiple benefits, such as bringing different expertise (expertise diversity; Zhang and Ma 2012) and smoothing out individual errors (error diversity; Zhou 2012). Two branches of related studies are: (1) how to make a balance between the diversity and the quality of the individuals (Chandra, Chen, and Yao 2006; Wood et al. 2023), and (2) how to measure the diversity (Kuncheva and Whitaker 2003). In our work, we address the latter.

Most ensemble-selection studies focus on binary classification tasks. Yule and Pearson (1900) propose a  $Q$  statistic to measure diversity based on the association between random variables. Fleiss (1981) introduces the measure of inter-rater reliability, quantifying the non-coincidental agreements among classifiers. Kohavi and Wolpert (1996) measure the diversity based on the variability of predictions across all classifiers. Skalak (1996) and Ho (1998) measure the disagreement between a pair of classifiers by the ratio of observations on which only one classifier is correct. To address multi-class classification, Kuncheva and Whitaker (2003) and Kadkhodaei and Moghadam (2016) reduce the problem to binary classification by considering the success/failure of classifiers, which is a binary variable. Although this approach measures expertise diversity, it loses information about error diversity, which we show is important and should not be ignored.

### 3 Approach

In this work, we propose an ensemble approach to unsupervised dependency parsing. In general, ensemble methods in machine learning consist of two stages: obtaining individual models and aggregating the individuals’ predictions. We will address both in the rest of this section.

#### 3.1 Aggregating Dependency Parses

We aim to aggregate parses of different individual models, potentially with different architectures and output formats. We propose to have post hoc aggregation methods applied to the dependency parses obtained by previous parsers. Inspired by our previous work (Shayegh et al. 2024), we formulate our ensemble under the minimum Bayes risk (MBR) framework. Consider a set of individuals’ outputs  $A_1, \dots, A_K$  given a sentence. The ensemble output  $A^*$  is defined as the dependency parse structure maximizing its similarity to all the individuals’ outputs:

$$A^* = \operatorname{argmax}_{A \in \mathcal{A}} \sum_{k=1}^K \operatorname{similarity}(A, A_k) \quad (1)$$

where  $\mathcal{A}$  is the set of all the possible dependency parse structures and similarity is a customizable similarity measure for dependency parses.

In particular, we propose to use the sentence-level unlabeled attachment score (UAS; Nivre and Fang 2017) as the similarity measure because it is the main evaluation metric for unlabeled, projective dependency parsing.<sup>2</sup> Formally,

<sup>2</sup>As an alternative to UAS, we also follow our previous work (Shayegh et al. 2024) using  $F_1$  as the similarity measure. We show results in Appendix A. We exclude it from the main body as it

UAS is the accuracy of head detection for every word (Le and Zuidema 2015; Nivre and Fang 2017; Yang et al. 2020). In other words,  $\operatorname{UAS}(A_p, A_r) = \frac{1}{n} \sum_{j=1}^n \mathbb{1}[A_p^{(j)} = A_r^{(j)}]$ , where  $A_p$  and  $A_r$  are two parallel lists of attachment heads for  $n$  given words in a sentence or a corpus. Putting UAS in Eqn. (1), we have

$$A^* = \operatorname{argmax}_{A \in \mathcal{A}} \sum_{k=1}^K \operatorname{UAS}(A, A_k) \quad (2)$$

$$= \operatorname{argmax}_{A \in \mathcal{A}} \sum_{j=1}^n \underbrace{\sum_{k=1}^K \mathbb{1}[A^{(j)} = A_k^{(j)}]}_{m(A^{(j)}; j)} \quad (3)$$

Here,  $m(a; j)$  is the number of individuals assigning word  $a$  as the  $j$ th word’s head. Specially,  $a = 0$  indicates the ROOT. We solve this problem by reducing it to the dependency parsing model of Eisner (1996) and using their efficient dynamic programming algorithm.

$$A^* = \operatorname{argmax}_{A \in \mathcal{A}} \sum_{j=1}^n m(A^{(j)}; j) \quad (4)$$

$$= \operatorname{argmax}_{A \in \mathcal{A}} \exp \left\{ \sum_{j=1}^n m(A^{(j)}; j) \right\} \quad (5)$$

$$= \operatorname{argmax}_{A \in \mathcal{A}} \prod_{j=1}^n \exp \{ \underbrace{m(A^{(j)}; j)}_{\hat{m}(A^{(j)}; j)} \} \quad (6)$$

$$= \operatorname{argmax}_{A \in \mathcal{A}} \prod_{j=1}^n \frac{\hat{m}(A^{(j)}; j)}{\underbrace{\sum_{a=0}^n \hat{m}(a; j)}_{p(A^{(j)}; j)}} \quad (7)$$

In Eqn. (7), we normalize values of  $\hat{m}$  to be in  $[0, 1]$ , which can be viewed as estimated probabilities. Consequently, the argmax objective is the joint probability of all attachments in a sentence, i.e., the probability of the dependency parse  $A$ . Hereby, we have reduced our problem to the dependency parsing problem in Eisner (1996) and can use their efficient algorithm to obtain  $A^*$  in  $\mathcal{O}(n^3)$  time complexity.

#### 3.2 Ensemble Selection

In this work, we build an ensemble of various unsupervised dependency parsers. We further propose to perform ensemble selection, as we notice in Figure 1 that an ensemble including all individuals underperforms a moderately sized ensemble.

We hypothesize that, if ensemble individuals lack diversity in their errors, they introduce systematic bias and may hurt the ensemble performance. Therefore, it is important to balance the quality of individuals and their diversity. Our ensemble selection objective is to find  $\mathcal{K} \subseteq \{1, \dots, K\}$  being

underperforms individuals and is considered a failure, demonstrating the importance of using the task-specific evaluation metric as the ensemble objective.

the set of selected ensemble individuals maximizing

$$\text{objective}(\mathcal{K}) = \sum_{\kappa \in \mathcal{K}} \text{UAS}(A_{\kappa}, A_{\text{gold}}) + \alpha \cdot \text{diversity}(\{A_{\kappa}\}_{\kappa \in \mathcal{K}}) \quad (8)$$

where  $A_{\text{gold}}$  is the ground-truth parse<sup>3</sup> and  $\alpha$  is a balancing hyperparameter.

For diversity, we introduce *society entropy*. Essentially, we first define *society distribution* (SD) for the  $j$ th word as

$$\text{SD}(a; \{A_{\kappa}^{(j)}\}_{\kappa \in \mathcal{K}}) = \frac{\sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)}=a]}{|\mathcal{K}|} \quad (9)$$

where the probability of the head being  $a$  is the fraction of the individuals agreeing on that. Then, we define the *society entropy* (SE) as the entropy  $E(p) = -\sum_x p(x) \log p(x)$  of society distribution:

$$\text{SE}(\{A_{\kappa}\}_{\kappa \in \mathcal{K}}) = \frac{1}{n} \sum_{j=1}^n E\left(\text{SD}\left(\cdot; \{A_{\kappa}^{(j)}\}_{\kappa \in \mathcal{K}}\right)\right) \quad (10)$$

where  $n$  is the number of words. This metric not only considers expertise diversity, but also measures the error diversity across individuals.

We maximize our objective in a forward-stepwise manner (Caruana et al. 2004). In other words, we begin by selecting

$$\kappa_1 = \underset{\kappa' \in \{1, \dots, K\}}{\text{argmax}} \text{UAS}(A_{\kappa'}, A_{\text{gold}}) \quad (11)$$

Then, we add individuals to our collection one at a time by maximizing our objective, i.e., the  $t$ th selected individual is chosen by

$$\kappa_t = \underset{\kappa' \in \{1, \dots, K\} \setminus \mathcal{K}_{t-1}}{\text{argmax}} \text{objective}(\mathcal{K}_{t-1} \cup \{\kappa'\}) \quad (12)$$

where  $\mathcal{K}_t = \{\kappa_1, \dots, \kappa_t\}$  is the first  $t$  selected individuals. Overall, the selected set of individuals is  $\mathcal{K}_T$ , where  $T$  is a hyperparameter indicating the number of ensemble individuals.

## 4 Experiments

### 4.1 Settings

**Dataset.** We performed experiments on the Wall Street Journal (WSJ) corpus in the Penn Treebank (Marcus, Santorini, and Marcinkiewicz 1993). We adopted the standard split: Sections 2–21 for training, Section 22 for validation, and Section 23 for testing. For training and validation, we followed previous work and only used sentences with at most 10 words after being stripped of punctuation and terminals. We used the entire test set for evaluation, regardless of sentence lengths. As all our individuals are unsupervised parsers, we did not use linguistic annotations in the training set. We use an annotated validation set<sup>3</sup> with only 250 sentences for validation during training and ensemble selection.

<sup>3</sup>We assume a small labeled validation set is available, which is a common assumption in both unsupervised syntactic structure detection (Jiang, Han, and Tu 2016; Deshmukh et al. 2021; Wu et al. 2024) and ensemble selection (Mienye and Sun 2022).

**Ensemble Individuals.** We consider the following unsupervised dependency parsers as our ensemble components<sup>4</sup>:

- CRFAE (Cai, Jiang, and Tu 2017), which is a discriminative and globally normalized conditional random field autoencoder model that predicts the conditional distribution of the structure given a sentence.
- NDMV (Jiang, Han, and Tu 2016), a dependency model with valence (DMV; Klein and Manning 2004) model that learns probability distributions using neural networks. We used Viterbi expectation maximization to train the model.
- NE-DMV (Jiang, Han, and Tu 2016), which applies the neural approach to the extended DMV model (Headen III, Johnson, and McClosky 2009; Gillenwater et al. 2010). We employed Tu and Honavar (2012) for initialization.
- L-NDMV (Han, Jiang, and Tu 2017), an extension of NDMV that utilizes lexical information of tokens.
- Sib-NDMV (Yang et al. 2020), which applies the neural approach to Sib-DMV, an extension of DMV that utilizes information of sibling tokens.
- Sib&L-NDMV (Yang et al. 2020), a joint model incorporating Sib-NDMV and L-NDMV.

For L-NDMV, Sib-NDMV, and Sib&L-NDMV, we use Naseem et al. (2010) for initialization. For the hyperparameters and configurations of individuals, we use the default values specified in the respective papers or codebases. We train five instances of each model using different random seeds to assess the stability of our approach. We observe that CRFAE, NDMV, and NE-DMV exhibit instability and lower performance on average than what the authors reported. To achieve comparable performance, we run the training for these models 20 times and select the top 25% based on their performance on the validation set. We report a comparison between the quoted results and our replication results in Appendix B, showing the success of our replication.

### 4.2 Aggregation Results

Table 1 shows the main results on the WSJ dataset, where we performed five runs of each individual and built ensembles of corresponding runs. We may or may not exclude Sib&L-NDMV as it is already a fusion model, and one may argue about including it in an ensemble where Sib-NDMV and L-NDMV are also present.

We compare our aggregation approach against the customized Ising model (CIM; Row 6; Ravikumar, Wainwright, and Lafferty 2010), which calculates per-sample vote weights based on individuals’ correlation over samples in classification tasks. Indeed, Kulkarni, Eulenstein, and Li (2024), as the only previous work on post hoc ensemble-based dependency parsing, shows that CIM outperforms alternatives, including the maximum spanning tree (Gavril 1987) and conflict resolution on heterogeneous data (Li et al.

<sup>4</sup>Our choice of individuals was highly based on the replicability of the baselines.

Method	UAS
Ensemble individuals	
1 CRFAE	53.0 $\pm$ 4.5
2 NDMV	48.1 $\pm$ 3.6
3 NE-DMV	51.0 $\pm$ 3.4
4 L-NDMV	62.4 $\pm$ 0.6
5 Sib-NDMV	64.3 $\pm$ 0.2
Ensembles	
6 CIM aggregation	58.8 $\pm$ 0.7
7 Our unweighted aggregation	65.7 $\pm$ 0.9
8 Our weighted aggregation	66.6 $\pm$ 0.4
Selected ensembles (our aggregation)	
9 w/o diversity	63.8
10 w/ Kuncheva’s diversity	65.8
11 w/ society entropy (ours)	<u>67.3</u>
12 Ensemble validation	<b>67.8</b>
Additional individual	
13 Sib&L-NDMV	67.9 $\pm$ 0.1
Ensembles (our aggregation)	
14 Unweighted	67.9 $\pm$ 0.3
15 Weighted	<u>68.4</u> $\pm$ 0.4
16 Selected w/ society entropy (ours)	<b>68.8</b>
17 Selected by ensemble validation	<u>68.4</u>

Table 1: Mean and standard deviation across five runs, evaluated on the WSJ test set. \*, \*\*, and \*\*\* denote statistically significant improvements over the best individual, as determined by a paired-sample T-test at confidence levels of 92%, 98%, and 99.99%, respectively. Selected ensembles do not include variance statistics, and thus, statistical tests can not be conducted for them.

2014).<sup>5</sup> The CIM is considered an unweighted approach, as it does not use labeled validation data. We see that this approach does not retain the performance of the best individual (Row 5) and is thus unsuccessful.

By contrast, our approach (Row 7) outperforms all the individuals and baselines, demonstrating the effectiveness of our MBR formulation. We also observe high stability in the performance of our ensemble despite the unstable individuals (Rows 1-3). However, we observe that building an ensemble including Sib&L-NDMV (Row 14) does not bring a performance boost over the best individual (Row 13) due to the huge performance gap between individuals, showing the limited robustness of the ensemble against weak individuals.

We further improve the performance of our ensemble by weighting individuals based on their performance on the validation set (Rows 8 and 15). We can build weighted ensembles with rational numbers as weights by virtually duplicating individuals (Shayegh, Wen, and Mou 2024). The weighted variant overcomes the challenge of high variance in individuals’ performance, exhibiting higher performance

<sup>5</sup>We strictly follow Kulkarni, Eulenstein, and Li (2024) to employ CIM in dependency parsing.

than the best individual and the unweighted ensemble. These results highlight the importance of controlling the negative effect of weak individuals. In the next section, we explore ensemble selection as a superior alternative technique, further enhancing robustness against weak individuals.

### 4.3 Ensemble Selection Results

We compare our diversity-based selection method with the classic forward-stepwise-selection approach (Caruana et al. 2004), referred to as *ensemble validation*. This baseline validates each possible ensemble on a validation set, and thus is slow and computationally expensive. Moreover, we compare our proposed diversity metric “society entropy” versus several previously proposed measures, listed below:

- Disagreement (Skalak 1996; Ho 1998), which is a pairwise metric, indicating the fraction of times only one classifier is true. We extend this metric to more than two classifiers by averaging disagreements through all pairs:

$$\sum_{\kappa_1, \kappa_2 \in \mathcal{K}} \sum_{j=1}^n \frac{\mathbb{1}[A_{\kappa_1}^{(j)} = A_{\text{gold}}^{(j)}] \oplus \mathbb{1}[A_{\kappa_2}^{(j)} = A_{\text{gold}}^{(j)}]}{n|\mathcal{K}|^2} \quad (13)$$

where  $n$  is the number of samples based on which we compute the diversity,  $A_{\text{gold}}$  is the ground truth,  $\mathcal{K}$  is the selected individuals, and  $\oplus$  is the “exclusive or” operator.

- KW variance (Kohavi and Wolpert 1996), which differs from the averaged disagreement metric above by the constant coefficient  $\frac{K-1}{2K}$  (Kuncheva and Whitaker 2003):

$$\sum_{j=1}^n \frac{\left( \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} = A_{\text{gold}}^{(j)}] \right) \left( \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} \neq A_{\text{gold}}^{(j)}] \right)}{n|\mathcal{K}|^2} \quad (14)$$

- Fleiss’ Kappa (Fleiss 1981), which differs from KW variance by a coefficient:

$$\sum_{j=1}^n \frac{\left( \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} = A_{\text{gold}}^{(j)}] \right) \left( \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} \neq A_{\text{gold}}^{(j)}] \right)}{n|\mathcal{K}|(1 - |\mathcal{K}|\bar{p})(1 - \bar{p})} \quad (15)$$

where  $\bar{p} = \frac{1}{n|\mathcal{K}|} \sum_{j=1}^n \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} = A_{\text{gold}}^{(j)}]$  is the fraction of true hits by all the individuals.

- Kuncheva’s diversity<sup>6</sup> (Kuncheva and Whitaker 2003), which is a non-pairwise metric measuring the smoothness of the oracle distribution:

$$\frac{1}{n(|\mathcal{K}| - \lceil |\mathcal{K}|/2 \rceil)} \sum_{j=1}^n \min \left\{ \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} = A_{\text{gold}}^{(j)}], \sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} \neq A_{\text{gold}}^{(j)}] \right\} \quad (16)$$

<sup>6</sup>Kuncheva and Whitaker (2003) name this measure as *entropy*. We avoid following this naming due to the different formulation of the standard entropy which we use for our society entropy.

Selection method	Selection time (ms)	Ensemble UAS			
		$k = 3$	$k = 4$	$k = 5$	Overall
w/o diversity	$0.03_{\pm 0}$	$64.0_{\pm 0.9}$	$65.6_{\pm 1.3}$	$65.9_{\pm 0.7}$	$65.2_{\pm 1.3}$
w/ Kuncheva's diversity	$12_{\pm 004}$	$65.0_{\pm 0.4}$	$67.7_{\pm 0.5}$	$66.1_{\pm 0.6}$	$66.3_{\pm 1.2}$
w/ society entropy (ours)	$22_{\pm 006}$	$67.1_{\pm 0.3}$	$66.9_{\pm 0.3}$	$67.2_{\pm 0.3}$	<b><math>67.1_{\pm 0.4}</math></b>
Ensemble validation	$11,066_{\pm 166}$	$66.6_{\pm 0.8}$	$66.6_{\pm 0.4}$	$67.7_{\pm 0.5}$	$67.0_{\pm 0.8}$

Table 2: UASs of selected ensembles with  $k$  ensemble components, selected from 15 individuals using different strategies. Numbers are the means and standard deviations across 30 runs. The Overall column represents all the 90 results for  $k = 3, 4, 5$ . Time is measured using 28 Intel(R) Core(TM) i9-9940X (@3.30GHz) CPUs.

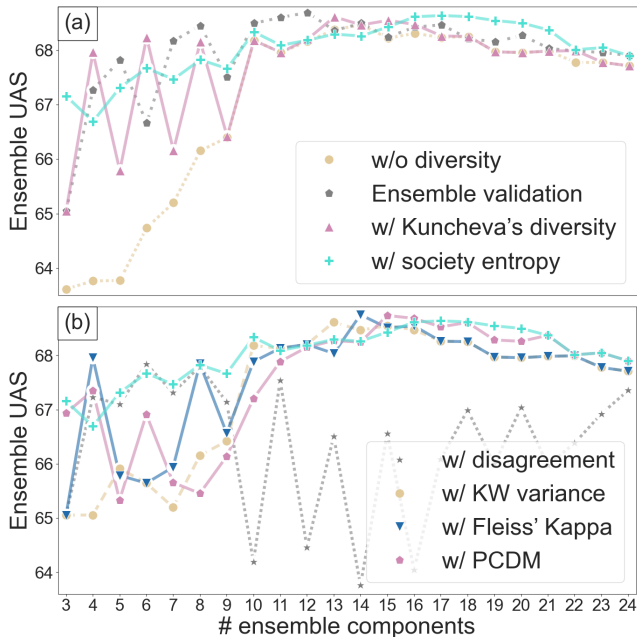


Figure 2: Ensemble performance by the number of selected ensemble components using different selection strategies. Ensemble selection happens over 25 individuals (five runs for each of CRFAE, NDMV, NE-DMV, L-NDMV, and Sib-NDMV). Results are split into two figures for easier reading.

- PCDM (Banfield et al. 2005), which is based on the proportion of classifiers getting each sample correct:

$$\frac{1}{n} \sum_{j=1}^n \mathbb{1} \left[ 0.1 \leq \frac{\sum_{\kappa \in \mathcal{K}} \mathbb{1}[A_{\kappa}^{(j)} = A_{\text{gold}}^{(j)}]}{|\mathcal{K}|} \leq 0.9 \right] \quad (17)$$

To obtain the above metrics, we consider our task as a classification of each word's head, which is consistent with the computation of our society entropy and is a simplification of the task yet incorporating the diversity notion.

For each diversity metric, we finetune the balancing hyperparameter  $\alpha$  in Eqn. (8) with one decimal-place precision based on the performance of the corresponding selected five-individual ensembles.

Figure 2 illustrates the performance of ensembles by different numbers of selected individuals. The selection meth-

ods include using only individuals' performance (w/o diversity), the ensemble objective given in Eqn. (8) employing different diversity metrics, and the classic forward-stepwise-selection approach (ensemble validation).

Our results show that ignoring diversity performs worse than most other methods, underscoring the importance of diversity in ensemble selection. On the other hand, society entropy outperforms the other approaches when a larger number of ensemble components are selected. This is expected, given the higher likelihood of error accumulation with a greater number of weak individuals and that society entropy is the only metric that accounts for error diversity. Moreover, this metric demonstrates significantly more stability than others when we have a small number of individuals. Overall, society entropy performs consistently well, making it the most reliable approach for selecting any number of ensemble components.

Additionally, we present a stability and efficiency analysis in Table 2 for selecting a small number of individuals, as this not only represents a more realistic use case for ensemble selection but also addresses the instability observed in the performance of Kuncheva's diversity and ensemble validation in Figure 2. To this end, we randomly selected 15 individuals from a group of 25. This random sampling was repeated 30 times to ensure robustness. For each of these 30 sample sets, we applied different selection strategies to choose subsets of 3, 4, and 5 individuals from the 15. Results show that society entropy excels and is more stable than the diversity-free and Kuncheva's diversity approaches, achieving performance comparable to ensemble validation but around 500x faster, making it a practical method.

Finally, we report in Table 1 the performance of selected ensembles. Here, we selected comparable numbers of individuals (five models for Rows 9-12 and six models for Rows 16-17) from all individuals across all runs. Our results are consistent with our analysis, demonstrating that our society entropy is comparable with ensemble validation, outperforming other ensemble-selection methods. Moreover, our society entropy-based selected ensemble outperforms all unselected ensembles and individuals, exhibiting the highest unsupervised dependency parsing performance among all competitors.

#### 4.4 Additional Analysis

**Inference Efficiency.** In Table 3, we report the inference runtime for every individual, along with the execution time




Figure 3: UAS by dependents’ POS tags on the WSJ test set. The seven represented POS tags cover around 67% of the cases.

Method	Time <sub>(ms)</sub>
Individual inference	
CRFAE <sup>†</sup>	14.9
NDMV	22.3
NE-DMV	22.3
L-NDMV	56.4
Sib-NDMV	114.2
Sib&L-NDMV	174.5
Our aggregation	19.6

Table 3: Inference time in milliseconds, measured by a single Intel(R) Core(TM) i9-9940X (@3.30GHz) CPU, without GPU, averaged through samples in the WSJ test. <sup>†</sup>Implemented in Java. Others are in Python.

of our aggregation algorithm. While the inference of an ensemble is generally slow as it requires the inference of all the individuals, results show that our proposed aggregation step does not further slow down this process significantly, demonstrating the efficiency of our proposed approach that makes use of Eisner (1996)’s algorithm.

### Performance by Dependents’ Part-of-Speech (POS).

We would like to investigate the effect of our ensemble approach from a linguistic perspective. To this end, we report in Figure 3 the breakdown performance of our weighted ensemble and its individuals by the POS tags of the dependents. We note that different individuals have their expertise in different types of dependents. In fact, NE-DMV surpasses the others in NNP and VBD cases, L-NDMV is outstanding in IN, RB, and VB, while Sib-NDMV outperforms other individuals in all other types. In most cases, our ensemble roughly matches the performance of the best individual, outperforming all of them in terms of overall performance.

**Appendices.** We include more results in the appendix: (A)aggregation based on  $F_1$  scores, (B)aseline replication, and (C)ase studies.

## 5 Conclusion

In this work, we propose a post hoc ensemble approach to unsupervised dependency parsing, equipped with a diversity-aware ensemble-selection method. We emphasize the importance of error diversity alongside expertise diversity and introduce society entropy as a measure that accounts for both. Our experiments demonstrate the effectiveness of our ensemble approach and the critical role of error diversity in ensemble selection.

**Future Work.** We identify future directions from algorithmic, linguistic, and machine learning perspectives. On the algorithmic side, it is intriguing to develop aggregation methods for other structures, including non-projective dependency parses or other graph structures in tasks such as drug discovery. For the linguistic aspect, it is promising to investigate cross-task ensembles by aggregating diverse knowledge of different constituency and dependency parsers. From the machine learning perspective, our ensemble selection is limited to classification tasks or those that can be framed as classification. We aim to explore ensemble-selection approaches, incorporating error diversity, for a broader range of problems.

In parallel with our work, Charakorn, Manoonpong, and Dilokthanakul (2024) identify a similar differentiation between expertise diversity (called *specialization*) and general diversity (which includes error diversity) in multi-agent reinforcement learning. It further supports our claim that both aspects of diversity should be taken into account, and suggests a great promise for our approach in multi-agent reinforcement learning.

## Acknowledgments

We thank all reviewers and chairs for their valuable and constructive comments. The research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Amii Fellow Program, the Canada CIFAR AI Chair Program, an Alberta Innovates Program, a donation from DeepMind, and the Digital Research Alliance of Canada (alliancecan.ca). We also thank Hans-Werner Eroms and Timothy J. Osborne for providing advice on dependency parse representations.

## References

- Amara, K.; Sevasťjanova, R.; and El-Assady, M. 2024. SyntaxShap: Syntax-aware explainability method for text generation. In *Findings of ACL*, 4551–4566.
- Banfield, R. E.; Hall, L. O.; Bowyer, K. W.; and Kegelmeyer, W. 2005. Ensemble diversity measures and their application to thinning. *Information Fusion*, 6(1): 49–62.
- Bod, R. 2009. From exemplar to grammar: A probabilistic analogy-based model of language learning. *Cognitive Science*, 33(5): 752–793.
- Breiman, L. 1996. Bagging predictors. *Machine Learning*, 24: 123–140.
- Cai, J.; Jiang, Y.; and Tu, K. 2017. CRF autoencoder for unsupervised dependency parsing. In *EMNLP*, 1638–1643.
- Caruana, R.; Munson, A.; and Niculescu-Mizil, A. 2006. Getting the most out of ensemble selection. In *ICDM*, 828–833.
- Caruana, R.; Niculescu-Mizil, A.; Crew, G.; and Ksikes, A. 2004. Ensemble selection from libraries of models. In *ICML*.
- Chandra, A.; Chen, H.; and Yao, X. 2006. Trade-off between diversity and accuracy in ensemble generation. In *Multi-Objective Machine Learning*, 429–464.
- Charakorn, R.; Manoonpong, P.; and Dilokthanakul, N. 2024. Diversity is not all you need: Training a robust cooperative agent needs specialist partners. In *NeurIPS*.
- Che, W.; Liu, Y.; Wang, Y.; Zheng, B.; and Liu, T. 2018. Towards better UD parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation. In *CoNLL Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, 55–64.
- Chomsky, N. 1956. Three models for the description of language. *IRE Transactions on Information Theory*, 2(3): 113–124.
- Chomsky, N. 1967. *Syntactic Structures*. The Hague: Mouton.
- Cunningham, P.; and Carney, J. 2000. Diversity versus quality in classification ensembles based on feature selection. In *ECML*, 109–116.
- Deshmukh, A. A.; Zhang, Q.; Li, M.; Lin, J.; and Mou, L. 2021. Unsupervised chunking as syntactic structure induction with a knowledge-transfer approach. In *Findings of EMNLP*, 3626–3634.
- Eisner, J. M. 1996. Three new probabilistic models for dependency parsing: An exploration. In *COLING*, 340–345.
- Fleiss, J. L. 1981. *Statistical Methods for Rates and Proportions*. Wiley.
- Gaifman, H. 1965. Dependency systems and phrase-structure systems. *Information and Control*, 8(3): 304–337.
- Ganaie, M.; Hu, M.; Malik, A.; Tanveer, M.; and Suganthan, P. 2022. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115: 105151.
- Gavril, F. 1987. Generating the maximum spanning trees of a weighted graph. *Journal of Algorithms*, 8(4): 592–597.
- Gillenwater, J.; Ganchev, K.; Graça, J.; Pereira, F.; and Taskar, B. 2010. Sparsity in dependency grammar induction. In *ACL Short Papers*, 194–199.
- Goldsmith, J. 2001. Unsupervised learning of the morphology of a natural language. *Computational Linguistics*, 27(2): 153–198.
- Han, W.; Jiang, Y.; Ng, H. T.; and Tu, K. 2020. A survey of unsupervised dependency parsing. In *COLING*, 2522–2533.
- Han, W.; Jiang, Y.; and Tu, K. 2017. Dependency grammar induction with neural lexicalization and big training data. In *EMNLP*, 1683–1688.
- Han, W.; Jiang, Y.; and Tu, K. 2019a. Enhancing unsupervised generative dependency parser with contextual information. In *ACL*, 5315–5325.
- Han, W.; Jiang, Y.; and Tu, K. 2019b. Lexicalized neural unsupervised dependency parsing. *Neurocomputing*, 349: 105–115.
- He, J.; Neubig, G.; and Berg-Kirkpatrick, T. 2018. Unsupervised learning of syntactic structure with invertible neural projections. In *EMNLP*, 1292–1302.
- Headden III, W. P.; Johnson, M.; and McClosky, D. 2009. Improving unsupervised dependency parsing with richer contexts and smoothing. In *HLT-NAACL*, 101–109.
- Ho, T. K. 1998. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8): 832–844.
- Jiang, Y.; Han, W.; and Tu, K. 2016. Unsupervised neural dependency parsing. In *EMNLP*, 763–771.
- Kadkhodaei, H.; and Moghadam, A. M. E. 2016. An entropy based approach to find the best combination of the base classifiers in ensemble classifiers based on stack generalization. In *International Conference on Control, Instrumentation, and Automation*, 425–429.
- Kann, K.; Mohanane, A.; Bowman, S. R.; and Cho, K. 2019. Neural unsupervised parsing beyond English. In *Proceedings of the Workshop on Deep Learning Approaches for Low-Resource NLP*, 209–218.
- Klein, D.; and Manning, C. 2004. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *ACL*, 478–485.
- Kohavi, R.; and John, G. H. 1997. Wrappers for feature subset selection. *Artificial Intelligence*, 97(1): 273–324.
- Kohavi, R.; and Wolpert, D. 1996. Bias plus variance decomposition for zero-one loss functions. In *ICML*, 275–283.
- Kulkarni, A.; Eulenstein, O.; and Li, Q. 2024. Empirical analysis for unsupervised universal dependency parse tree aggregation. *arXiv preprint arXiv:2403.19183*.
- Kulkarni, A.; Sabetpour, N.; Markin, A.; Eulenstein, O.; and Li, Q. 2022. CPTAM: Constituency parse tree aggregation method. In *Proceedings of the SIAM International Conference on Data Mining*, 630–638.
- Kuncheva, L. I. 2004. *Combining Pattern Classifiers: Methods and Algorithms*. Wiley.
- Kuncheva, L. I.; and Whitaker, C. J. 2003. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine Learning*, 51(2): 181–207.

- Le, P.; and Zuidema, W. 2015. Unsupervised dependency parsing: Let's use supervised parsers. In *NAACL-HLT*, 651–661.
- Li, Q.; Li, Y.; Gao, J.; Zhao, B.; Fan, W.; and Han, J. 2014. Resolving conflicts in heterogeneous data by truth discovery and source reliability estimation. In *SIGMOD*, 1187–1198.
- Lim, K.; Park, C.; Lee, C.; and Poibeau, T. 2018. SEx BiST: A multi-source trainable parser with deep contextualized lexical representations. In *CoNLL Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, 143–152.
- Liu, H.; and Motoda, H. 2007. *Computational Methods of Feature Selection*. CRC Press.
- Marcus, M. P.; Santorini, B.; and Marcinkiewicz, M. A. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2): 313–330.
- Mienye, I. D.; and Sun, Y. 2022. A survey of ensemble learning: Concepts, algorithms, applications, and prospects. *IEEE Access*, 10: 99129–99149.
- Minku, L. L.; White, A. P.; and Yao, X. 2010. The impact of diversity on online ensemble learning in the presence of concept drift. *IEEE Transactions on Knowledge and Data Engineering*, 22(5): 730–742.
- Naseem, T.; Chen, H.; Barzilay, R.; and Johnson, M. 2010. Using universal linguistic knowledge to guide grammar induction. In *EMNLP*, 1234–1244.
- Nivre, J.; and Fang, C.-T. 2017. Universal dependency evaluation. In *Proceedings of the NoDaLiDa Workshop on Universal Dependencies*, 86–95.
- Purucker, L. O.; Schneider, L.; Anastacio, M.; Beel, J.; Bischl, B.; and Hoos, H. 2023. Q(D)O-ES: Population-based quality (diversity) optimisation for post hoc ensemble selection in AutoML. In *AutoML*.
- Ravikumar, P.; Wainwright, M. J.; and Lafferty, J. D. 2010. High-dimensional Ising model selection using  $\ell_1$ -regularized logistic regression. *The Annals of Statistics*, 38(3): 1287 – 1319.
- Rokach, L. 2009. *Pattern Classification using Ensemble Methods*. World Scientific.
- Shayegh, B.; Cao, Y.; Zhu, X.; Cheung, J. C.; and Mou, L. 2024. Ensemble distillation for unsupervised constituency parsing. In *ICLR*.
- Shayegh, B.; Wen, Y.; and Mou, L. 2024. Tree-averaging algorithms for ensemble-based unsupervised discontinuous constituency parsing. In *ACL*, 15135–15156.
- Shen, Y.; Tay, Y.; Zheng, C.; Bahri, D.; Metzler, D.; and Courville, A. 2021. StructFormer: Joint unsupervised induction of dependency and constituency structure from masked language modeling. In *ACL-IJCNLP*, 7196–7209.
- Skalak, D. B. 1996. The sources of increased accuracy for two proposed boosting algorithms. In *Proceedings of the AAAI Workshop on Integrating Multiple Learned Models*, 120–125.
- Smith, N. A.; and Eisner, J. 2005. Guiding unsupervised grammar induction using contrastive estimation. In *Proceedings of the IJCAI Workshop on Grammatical Inference Applications*, 73–82.
- Tang, E. K.; Suganthan, P. N.; and Yao, X. 2006. An analysis of diversity measures. *Machine Learning*, 65: 247–271.
- Tesnière, L.; Osborne, T. J.; and Kahane, S. 2015. *Elements of Structural Syntax*. John Benjamins.
- Tu, K.; and Honavar, V. 2012. Unambiguity regularization for unsupervised learning of probabilistic grammars. In *EMNLP-CoNLL*, 1324–1334.
- Wang, K.; and Cohen, S. B. 2025. DEPFold: RNA secondary structure prediction as dependency parsing. In *ICLR*.
- Wen, Y.; Hao, Y.; Cao, Y.; and Mou, L. 2023. An equal-size hard EM algorithm for diverse dialogue generation. In *ICLR*.
- Wen, Y.; Shayegh, B.; Huang, C.; Cao, Y.; and Mou, L. 2025. EBBS: An ensemble with bi-Level beam search for zero-shot machine translation. In *AAAI*.
- Wood, D.; Mu, T.; Webb, A. M.; Reeve, H. W. J.; Luján, M.; and Brown, G. 2023. A unified theory of diversity in ensemble learning. *JMLR*, 24(359): 1–49.
- Wu, Y.; Liu, L.; Xie, Z.; Chow, K.-H.; and Wei, W. 2021. Boosting ensemble accuracy by revisiting ensemble diversity metrics. In *CVPR*, 16469–16477.
- Wu, Z.; Chen, Y.; Kao, B.; and Liu, Q. 2020. Perturbed masking: Parameter-free probing for analyzing and interpreting BERT. In *ACL*, 4166–4176.
- Wu, Z.; Deshmukh, A. A.; Wu, Y.; Lin, J.; and Mou, L. 2024. The emergence of chunking structures with hierarchical HRNN. *Computational Linguistics*.
- Yang, S.; Jiang, Y.; Han, W.; and Tu, K. 2020. Second-order unsupervised neural dependency parsing. In *COLING*, 3911–3924.
- Yule, G. U.; and Pearson, K. 1900. On the association of attributes in statistics: With illustrations from the material of the childhood society, &c. *Philosophical Transactions of the Royal Society A*, 194: 257–319.
- Zhang, C.; and Ma, Y. 2012. *Ensemble Machine Learning: Methods and Applications*. Springer.
- Zhang, M. 2020. A survey of syntactic-semantic parsing based on constituent and dependency structures. *Science China Technological Sciences*, 63(10): 1898–1920.
- Zhou, Z.-H. 2012. *Ensemble Methods: Foundations and Algorithms*. Chapman & Hall/CRC.