

To be Continuous, or to be Discrete, Those are *Bits* of Questions

Yiran Wang Masao Utiyama

National Institute of Information and Communications Technology (NICT)
yiran.wang@nict.go.jp mutiyama@nict.go.jp

Abstract

Recently, binary representation has been proposed as a novel representation that lies between continuous and discrete representations. It exhibits considerable information-preserving capability when being used to replace continuous input vectors. In this paper, we investigate the feasibility of further introducing it to the output side, aiming to allow models to output binary labels instead. To preserve the structural information on the output side along with label information, we extend the previous contrastive hashing method as structured contrastive hashing. More specifically, we upgrade CKY from label-level to bit-level, define a new similarity function with span marginal probabilities, and introduce a novel contrastive loss function with a carefully designed instance selection strategy. Our model¹ achieves competitive performance on various structured prediction tasks, and demonstrates that binary representation can be considered a novel representation that further bridges the gap between the continuous nature of deep learning and the discrete intrinsic property of natural languages.

1 Introduction

Bridging the gap between the continuous nature of deep learning and the discrete intrinsic property of natural languages has been one of the most fundamental and essential questions since the very beginning. Continuous representation makes the training of neural networks effective and efficient. Nowadays, representing discrete natural languages in continuous format is the first and foremost step to leveraging the capabilities of deep learning. One could even argue that the exhilarating advancements in natural language processing in the past decade can largely be attributed to the word embedding technique, as it is the first successful attempt.

Word embedding (Bengio et al., 2000; Mikolov et al., 2013a,b) technique replaces the vocabulary-

¹<https://github.com/speedcell14/parserker>

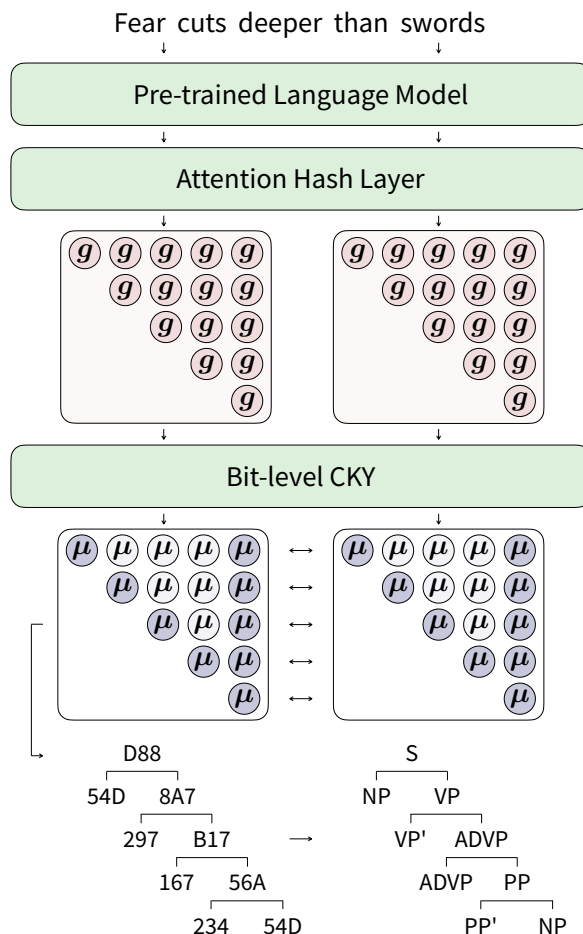


Figure 1: The model architecture. The attention hash layer produces span scores (pink circles), we only use the upper triangular part of these scores and feed them into the bit-level CKY to obtain the marginal probabilities of all valid spans (purple circles). During training, we only select the spans on the target trees for structured contrastive hashing and leave the other spans unused (transparent purple circles). During inference, as shown at the bottom, our model parses sentences by returning trees with label codes (hexadecimal numbers), which are then translated back to the original labels.

sized one-hot word representations with compact continuous vectors. Since then, input and output layers incorporating embedding matrices have be-

come standard components in neural models. Discrete tokens are mapped into continuous vectors by looking up the corresponding index in it, and continuous vectors are mapped back to discrete tokens by searching the most similar one from it. However, the essence of this operation is still one-hot encoding, even the following subword tokenization techniques (Sennrich et al., 2016; Kudo, 2018) attempt to mitigate this issue by decomposing words into subword units, these approaches still require the building vocabularies and embedding matrices that consists of tens of thousands of tokens. In the era of large language models (OpenAI et al., 2023; Touvron et al., 2023), these embedding matrices typically account for a considerable number of parameters, especially in cross-lingual models. Besides, parameter updates also solely rely on the sparse gradients backpropagated to the limited tokens present in sentences. Moreover, imposing structural constraints on continuous representations to model relations among tokens is considered difficult, whereas it is easy and common in discrete representations. Therefore, further bridging the gap has become increasingly important nowadays.

Recently, Wang et al. (2023) introduced a novel binary representation that lies between continuous and discrete representations. They proposed a contrastive hashing method to compress continuous hidden states into binary codes. These codes contain all the necessary task-relevant information, and using them as the only inputs can reproduce the performance of the original models. Unlike associating each token with only a single vector, their method allocates multiple bits to each token, and the token representation can be constructed by concatenating these bit vectors. In other words, their binary representation breaks tokens down into combinations of semantic subspaces. As a result, replacing the token embedding matrix in the input layer with a tiny bit embedding matrix without sacrificing performance becomes possible.

In this paper, we explore the possibility of further introducing this representation to output layers. In the input layer, structural information can only be implicitly obtained by introducing the task loss as an auxiliary. However, the output layers often involve complex intra-label constraints, especially for structured prediction tasks, structural information can and should be explicitly preserved along with plain label information. Therefore, we attempt to endow models with this capability by extending previous contrastive hashing to structured




Figure 2: An example of the geometric center issue. Orange circles are positive to the black circle instance, while the dotted orange circle is their geometric center. The difference between \mathcal{L}_{sup} and our \mathcal{L}_{max} is that we target the closest positive instead of their geometric center.

contrastive hashing.

We begin by upgrading the CKY, which parses sentences, returns spans with discrete labels, to support binary format labels (§3.1). Subsequently, we define a new similarity function by using span marginal probabilities obtained from this bit-level CKY (§3.2) to jointly learn label and structural information. Furthermore, we conduct a detailed analysis of several widely used contrastive learning losses, identifying the geometric center issue, and introduce a novel contrastive learning loss to remedy it (§3.3) through carefully selecting instances. By doing so, we show that it is feasible to introduce binary representation to output layers and have them output binary labels on trees. Moreover, since our model is based on contrastive learning, it also benefits from its remarkable representation learning capability, resulting in better performance than existing models. We conduct experiments on constituency parsing and nested named entity recognition. Experimental results (§4.2) demonstrate that our models achieve competitive performance with only around 12 and 8 bits, respectively.

2 Background

2.1 Constituency Parsing

For a given sentence w_1, \dots, w_n , constituency parsing aims at detecting its hierarchical syntactic structures. Previous work (Stern et al., 2017; Gaddy et al., 2018; Kitaev and Klein, 2018; Kitaev et al., 2019; Zhang et al., 2020) decompose tree score $g(\mathbf{t})$ as the sum of its constituents scores,

$$g(\mathbf{t}) = \sum_{\langle l_i, r_i, y_i \rangle \in \mathbf{t}} g(l_i, r_i, y_i) \quad (1)$$

where l_i and r_i indicate the left and right boundary of the i -th span, and $y_i \in \mathcal{Y}$ stands for the label. Constituent score $g(l_i, r_i, y_i)$ reflects the joint score

of selecting the specified span and assigning it the specified label. Previous work (Kitaev and Klein, 2018; Zhang et al., 2020) commonly compute this score using a linear or bilinear component. Under the framework of graphical probabilistic models, they can efficiently compute the conditional probability by applying the CKY algorithm.

$$p(\mathbf{t}) = \frac{\exp g(\mathbf{t})}{Z \equiv \sum_{\mathbf{t}' \in \mathcal{T}} \exp g(\mathbf{t}')} \quad (2)$$

Z is commonly known as the partition function which enumerates all valid constituency trees.

Besides, marginal probability $\mu(l_i, r_i, y_i)$ is also frequently mentioned. It stands for the proportion of scores for all trees that include the specified span with the specified label. As noted by Eisner (2016), computing the partial derivative of the log partition with respect to the span score is an efficient approach to obtain the marginal probability.

$$\mu(l_i, r_i, y_i) = \frac{\partial \log Z}{\partial g(l_i, r_i, y_i)} \quad (3)$$

Intuitively, marginal probability indicates the joint probability of selecting a specified span with a specified label. Therefore, it is easy to notice that merely summing the marginal probabilities for all labels of a given span does not always yield 1, i.e., $\sum_{y' \in \mathcal{Y}} \mu(l_i, r_i, y') \neq 1$, as there is no guarantee that this span will be selected. In other words, marginal probabilities contain not only label information but also structural information. If a span is unlikely to be selected, its marginal probability will not be high regardless of the label.

2.2 Contrastive Hashing

Contrastive learning (He et al., 2020; Gao et al., 2021; Khosla et al., 2020) is an effective yet simple representation learning method, which involves pulling together positive pairs and pushing apart negative pairs in a metric space. Recently, Wang et al. (2023) extended this approach as contrastive hashing. They append an untrained transformer to the end of a pre-trained language model and use its attention scores for both task learning and hashing. Specifically, its entire attention probabilities $a_{i,j}^k$ are used to compute hidden states for downstream tasks as usual, and its diagonal entries $s_{i,i}^k$ of the attention scores are employed for hashing.

$$s_{i,j}^k = \frac{(\mathbf{W}_k^Q \mathbf{h}_i)^\top (\mathbf{W}_k^K \mathbf{h}_j)}{\sqrt{d_k}} \quad (4)$$

$$a_{i,j}^k = \text{softmax}_j(s_{i,j}^k) \quad (5)$$

Where \mathbf{W}_k^Q and \mathbf{W}_k^K are parameters, \mathbf{h} are hidden states, d_k is the head dimension. These two learning objectives share the same attention matrix, therefore, task-relevant information is implicitly ensured to be preserved in these binary codes.

More specifically, to leverage the multi-head mechanism, they allow each head to represent one and only one bit. By increasing the number of heads to K , they obtain attention scores from K different semantic aspects. During the inference stage, codes are generated by binarizing these scores,

$$\mathbf{c}_i = [c_i^1, \dots, c_i^K] \in \{-1, +1\}^K \quad (6)$$

$$c_i^k = \text{sign}(s_{i,i}^k)$$

During the training stage, they approximate Hamming similarity by computing the cosine similarity, with one of its inputs binarized first.

$$s(i, j) = \cos(\mathbf{s}_{i,i}, \mathbf{c}_j) \quad (7)$$

Apart from this similarity function, they also propose a novel loss by carefully selecting instances and eliminating potential positives and negatives. They fine-tune the entire model using both the downstream task loss and the contrastive hashing loss, i.e., $\mathcal{L} = \mathcal{L}_{\text{task}} + \beta \cdot \mathcal{L}_{\text{contrastive}}$. Experiments show that they can reproduce the original performance on an extremely tiny model using only these 24-bit codes as inputs. Therefore, they claim that these codes preserve all the necessary task-relevant information.

3 Proposed Methods

Our model attempts to learn parsing and hashing simultaneously with a single structured contrastive hashing loss. In other words, we try to introduce the binary representation to output layers and eliminate the need for the $\mathcal{L}_{\text{task}}$ above. To achieve this, we first extend the CKY module to support binary labels (§3.1). Then, we replace the cosine similarity with a newly defined similarity function (§3.2) based on span marginal probabilities, because it contains not only label information but also structural information. After that, we analyze several commonly used contrastive losses, and propose a new one (§3.3) to mitigate the geometric center issue as shown in Figure 2. After training, we build code vocabulary by mapping binary codes back to their most frequently coinciding labels.

3.1 Constituency Parsing with Bits

We decompose tree scores as the sum of constituent scores as well, but with discrete labels y_i replaced with binary codes $\mathbf{c}_i \in \{-1, +1\}^K$.

$$g(\mathbf{t}) = \sum_{\langle l_i, r_i, \mathbf{c}_i \rangle \in \mathbf{t}} g(l_i, r_i, \mathbf{c}_i) \quad (8)$$

$$g(l_i, r_i, \mathbf{c}_i) = \sum_{k=1}^K g_k(l_i, r_i, c_i^k) \quad (9)$$

Where $g_k(l_i, r_i, c_i^k)$ represents the span score with the k -th bit position assigned as value c_i^k . We additionally assume that the bits are independent of each other, so we simply add their scores together to obtain the span score $g(l_i, r_i, \mathbf{c}_i)$.

Following Wang et al. (2023), we maintain the one-head-one-bit design and also utilize attention scores for hashing. Furthermore, since we attempt to eliminate $\mathcal{L}_{\text{task}}$, we do not need to compute the final outputs of the transformer layer. Therefore, we only retain the query \mathbf{W}_k^Q and key \mathbf{W}_k^K to calculate the span score for getting $+1$ in the k -th bit position by using the token hidden states of the left and right span boundary \mathbf{h}_{l_i} and \mathbf{h}_{r_i} . For the -1 case, we simply leave its score as 0.

$$g_k(l_i, r_i, +1) = \frac{(\mathbf{W}_k^Q \mathbf{h}_{l_i})^\top (\mathbf{W}_k^K \mathbf{h}_{r_i})}{\sqrt{d_k}} \quad (10)$$

$$g_k(l_i, r_i, -1) = 0 \quad (11)$$

With these definitions, we can extend the CKY module to the bit-level and calculate the conditional probability and partition function using Equation 2 as usual. Additionally, the bit-level marginal probability is defined as below.

$$\mu_k(l_i, r_i, c_i^k) = \frac{\partial \log Z}{\partial g_k(l_i, r_i, c_i^k)} \quad (12)$$

3.2 Contrastive Hashing with Structures

Wang et al. (2023) emphasize that the key to their loss function is first hashing one of its inputs as codes and then calculating the similarity between continuous scores and the discrete codes. We define our similarity function in a similar way. Since we can straightforwardly obtain the span marginal probabilities with Equation 12, we then binarize scores into codes towards the sides with the higher span marginal probabilities.

$$\mathbf{c}_i = [c_i^1, \dots, c_i^K] \in \{-1, +1\}^K \quad (13)$$

$$c_i^k = \begin{cases} +1 & \mu_k(l_i, r_i, +1) > \mu_k(l_i, r_i, -1) \\ -1 & \text{otherwise} \end{cases}$$

Naturally, we define the similarity function between two spans as the marginal probability of selecting i -th span while assigning \mathbf{c}_j as its code.

$$s(i, j) = \frac{1}{K} \sum_{k=1}^K \mu_k(l_i, r_i, c_j^k) \quad (14)$$

As we mentioned above (§2.1), we use marginal probabilities to define the similarity function because they reflect the joint probability of both selecting the specified span and assigning the specified label to it. If a span is unlikely to be selected as a phrase, then both $\mu_k(l_i, r_i, +1)$ and $\mu_k(l_i, r_i, -1)$ will be close to zero. Thus, the model learns structural and label information simultaneously. By leveraging this similarity function, we extend contrastive hashing to structured contrastive hashing. This approach eliminates the label embedding from the output layer, as the hashing layer returns labels in binary format now.

Moreover, for a sentence with n tokens, the total number of spans is $(n^2 + n)/2$, and bit-level CKY returns marginal probabilities for them all. However, using them for contrastive hashing leads to an intractable time complexity of $\mathcal{O}(n^4)$. In practice, we select only spans from the target trees, reducing the number of spans to $2n - 1$ to maintain the time complexity at $\mathcal{O}(n^2)$. This is another reason why we prefer marginal probability.

3.3 Instance Selection

Following the contrastive learning framework (Gao et al., 2021; Wang et al., 2023), we feed sentences into the neural network twice to obtain two semantically identical but slightly augmented representations. In this way, we get two different marginal probabilities for each span. We calculate contrastive losses by comparing each span across views and average them as the batch loss (§3.5). Note that each batch contains spans from multiple sentences, so our contrastive hashing also compares spans across different sentences. For clarity, we omit subscripts in the following equations when there is no ambiguity.

As we mentioned above (§2.2), the fundamental concept of contrastive learning is pulling together positives and pushing apart negatives. The most commonly used objective function is defined as,

$$\mathcal{L}_{\text{self}} = -\log \frac{\exp s(i, i)}{\sum_{j \in \mathcal{N} \cup \mathcal{P}} \exp s(i, j)} \quad (15)$$

$\mathcal{N} = \{j \mid y_j \neq y_i\}$ and $\mathcal{P} = \{j \mid y_j = y_i\}$ stands for the negative and positive sets, respectively. We

additionally define $\mathcal{S} = \{i\}$ as the set that contains span i as its only entry. It is obvious that $\mathcal{S} \subseteq \mathcal{P}$ always holds.

In addition, the $\log \sum \exp$ operator is commonly considered a differentiable approximation of the max operator. Therefore, by slightly tweaking the equation, we reinterpret it as below.

$$\begin{aligned} \mathcal{L}_{\text{self}} &= \log \sum_j \exp s(i, j) - s(i, i) \\ &\approx \max_{j \in \mathcal{N} \cup \mathcal{P}} s(i, j) - s(i, i) \end{aligned} \quad (16)$$

Moreover, [Khosla et al. \(2020\)](#) also proposed a loss function for supervised settings where multiple positives are present. By applying the same tricks, we can rewrite the loss function as follows.

$$\begin{aligned} \mathcal{L}_{\text{sup}} &= -\frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \log \frac{\exp s(i, p)}{\sum_{j \in \mathcal{N} \cup \mathcal{P}} \exp s(i, j)} \\ &= \log \sum_j \exp s(i, j) - \frac{1}{|\mathcal{P}|} \sum_p s(i, p) \\ &\approx \max_{j \in \mathcal{N} \cup \mathcal{P}} s(i, j) - \text{mean}_{j \in \mathcal{P}} s(i, j) \end{aligned} \quad (17)$$

[Wang et al. \(2023\)](#) have also proposed a contrastive loss function for hashing. They claim that identical tokens may even not contain identical information due to different contexts. Therefore, they treat \mathcal{P} as potential false positives and negatives, and replace it with \mathcal{S} in both terms.

$$\begin{aligned} \mathcal{L}_{\text{hash}} &= -\log \frac{\exp s(i, i)}{\sum_{j \in \mathcal{N} \cup \mathcal{S}} \exp s(i, j)} \\ &\approx \max_{j \in \mathcal{N} \cup \mathcal{S}} s(i, j) - s(i, i) \end{aligned} \quad (18)$$

By unifying them all in a common format, we can observe that their main differences lie in the instance selection strategies. Both $\mathcal{L}_{\text{self}}$ and \mathcal{L}_{sup} pull instances towards the geometric center of their positive instances in **the second term**. The only difference is that $\mathcal{L}_{\text{self}}$ assumes there is only one positive instance, making the geometric center merely i -th span itself, whereas \mathcal{L}_{sup} has access to the ground-truth labels and thus can obtain a more specific center. On the other hand, $\mathcal{L}_{\text{hash}}$ differs from $\mathcal{L}_{\text{self}}$ in **the first term**, as it suggests that dividing instances solely based on ground-truth labels may also introduce false positives and negatives. Therefore, $\mathcal{L}_{\text{hash}}$ excludes potential false positives from this term. Additionally, it is noteworthy that among these three losses, all **first terms** employ max, while

all **second terms** use mean. Intuitively speaking, the max operator pulls towards the most likely true positive instance, while mean operator pulls towards the geometric center of all positive instances. Although it is hard to determine whether instances in \mathcal{P} are false positives or not, what we can be certain of is that there is at least one true positive instance, since $\mathcal{S} \subseteq \mathcal{P}$ always holds. Therefore, using the max operator to pull towards the most probable one is a more effective approach.

$$\begin{aligned} \mathcal{L}_{\text{max}} &\approx \max_{j \in \mathcal{N} \cup \mathcal{S}} s(i, j) - \max_{j \in \mathcal{P}} s(i, j) \\ &\approx \log \sum_{j \in \mathcal{N} \cup \mathcal{S}} \exp s(i, j) \\ &\quad - \log \sum_{p \in \mathcal{P}} \exp s(i, p) \\ &= -\log \frac{\sum_{p \in \mathcal{P}} \exp s(i, p)}{\sum_{j \in \mathcal{N} \cup \mathcal{S}} \exp s(i, j)} \end{aligned} \quad (19)$$

3.4 Architecture

As shown in [Figure 1](#), our model consists of a pre-trained language model, an attention hash layer, and a bit-level CKY. The attention hash layer derives from the transformer layer, but it preserves only the query and key components for calculating span score $g(l_i, r_i, c_i)$. All the other components, such as layer normalization ([Ba et al., 2016](#)) and feed-forward layers, are removed, as there is no need to calculate hidden states.

Compared with existing constituency parsers ([Kitaev and Klein, 2018](#); [Zhang et al., 2020](#)), the output layer of our parser does not include label embedding matrices. Instead, it utilizes the attention hash layer and a bit-level CKY to predict the binary representation of labels. Consequently, the number of parameters of this part changes from $|\mathcal{Y}| \times d$ to two $K \times \lceil \frac{d}{K} \rceil \times d$.

3.5 Training and Inference

During the training stage, we feed sentences into the model twice to obtain marginal probabilities μ^1 and μ^2 of the two views, and binarize spans on the target trees $\hat{\mathbf{t}}^1 = \{\langle l_i, r_i, c_i^1 \rangle\}_{i=1}^{2n-1}$ and $\hat{\mathbf{t}}^2 = \{\langle l_i, r_i, c_i^2 \rangle\}_{i=1}^{2n-1}$ using [Equation 13](#), respectively. Besides, we use the ground-truth labels y_i to divide \mathcal{N} and \mathcal{P} . After that, we calculate the contrastive hashing loss for each span with [Equation 19](#) and average them as the batch loss.

$$\mathcal{L} = \text{mean}_{1 \leq i < 2n} (\mathcal{L}(i, \mu^1, \hat{\mathbf{t}}^2) + \mathcal{L}(i, \mu^2, \hat{\mathbf{t}}^1)) \quad (20)$$

MODEL	PTB		CTB
	BERT	XLNET	BERT
Kitaev et al. (2019) ^b	95.59	-	91.75
Zhou and Zhao (2019) ^b	95.84	96.33	92.18
Mrini et al. (2020) ^p	-	96.38	92.64
Zhang et al. (2020) ^b	95.69	-	92.27
Yang and Deng (2020) [#]	95.79	96.34	93.59
Tian et al. (2020) ^b	95.86	<u>96.40</u>	<u>92.66</u>
Nguyen et al. (2021) [◇]	95.70	-	-
Xin et al. (2021) ^b	95.92	-	92.50
Cui et al. (2022) ^b	95.92	-	92.31
Yang and Tu (2022) [◇]	<u>96.01</u>	-	-
Yang and Tu (2023) [◇]	96.04	96.48	92.41
Ours (6 bits) ^b	94.81	95.70	91.45
Ours (8 bits) ^b	95.95	96.34	91.99
Ours (10 bits) ^b	96.03	96.37	<u>92.25</u>
Ours (12 bits) ^b	96.00	96.36	92.33
Ours (14 bits) ^b	<u>96.02</u>	96.43	92.06
Ours (16 bits) ^b	95.98	<u>96.40</u>	92.18

Table 1: The constituency parsing results. The **bold numbers** and the underlined numbers indicate the best and the second-best performance of each column. ^b[#][◇] stands for the graph-based, transition-based, and sequence-to-sequence models, respectively.

After training, we switch the model to evaluation mode, i.e., turning off dropouts (Srivastava et al., 2014), and then feed all the training sentences into the model again. During this pass, we count the frequency of each pair of binary code and its corresponding ground-truth label, i.e., $f(\mathbf{c}, y)$. Then, we can reconstruct a code vocabulary to map codes back to their most frequently coinciding labels.

$$y_{\mathbf{c}} \leftarrow \arg \max_{y \in \mathcal{Y}} f(\mathbf{c}, y) \quad (21)$$

During the inference stage, we search the most probable tree from all valid trees using the Cocke-Kasami-Younger (CKY) algorithm (Kasami, 1966). Our bit-level parsers return span boundaries and their binary codes, we translate them back to labels by using the Equation above.

4 Experiments

4.1 Settings

We validate the effectiveness of our model on various structured prediction tasks, i.e., constituency parsing and nested named entity recognition tasks. Dataset statistics can be found in Appendix A.

For the constituency parsing task, we conduct experiments on the datasets PTB (Marcus et al.,

MODEL	ACE'4	ACE'5	GENIA
	BERT	BERT	BERT
Wang et al. (2020) [◇]	86.28	84.66	79.19
Wang et al. (2021) [◇]	86.06	84.71	78.67
Xu et al. (2021) [◇]	86.30	85.40	79.60
Fu et al. (2021) ^b	86.60	85.40	78.20
Yu et al. (2020) [◇]	86.70	85.40	<u>80.50</u>
Shen et al. (2021) [◇]	87.41	86.67	80.54
Tan et al. (2021) [◇]	87.26	<u>87.05</u>	80.44
Lou et al. (2022) ^b	<u>87.90</u>	86.91	78.44
Zhu and Li (2022) [◇]	87.98	87.15	-
Yang and Tu (2022) [◇]	86.94	85.53	78.16
Ours (4 bits) ^b	85.81	83.37	73.54
Ours (6 bits) ^b	87.39	85.23	78.57
Ours (8 bits) ^b	87.93	85.90	78.79
Ours (10 bits) ^b	<u>87.87</u>	<u>85.75</u>	<u>78.72</u>
Ours (12 bits) ^b	87.52	85.26	78.40

Table 2: The nested named entity recognition results. ^b[◇][◇] stands for graph-based, span-based, sequential labeling, and sequence-to-sequence models, respectively.

1993) and CTB5.1 (Xue et al., 2005). We transform the original trees into those of Chomsky normal form and adopt left binarization with NLTK (Bird and Loper, 2004). We study model performance by employing pre-trained language models with checkpoints bert-large-cased (Devlin et al., 2019) and xlnet-large-cased (Yang et al., 2019) for PTB, and bert-base-chinese for CTB.

For nested named entity recognition task, we use datasets ACE2004 (Doddington et al., 2004), ACE2005 (Walker et al., 2006), and GENIA (Kim et al., 2003). We follow the data splitting of Shibuya and Hovy (2020). Nested named entity recognition, as Fu et al. (2021) claimed, can be considered as a partially observed constituency parsing task. Therefore, we add a dummy span TOP as the top span to each sentence to ensure all the observed spans form a valid parsing tree, and apply the same transformation and binarization on it. We use the checkpoint bert-large-cased on ACE2004 and ACE2005, and for the GENIA dataset we use dmis-lab/biobert-large-cased-v1.1 (Lee et al., 2019) as the pre-trained language model.

We utilize the deep learning framework PyTorch (Paszke et al., 2019) to implement our models and download pre-trained language checkpoints from huggingface/transformers (Wolf et al., 2020).

To keep the training of contrastive hashing stable, we collect sentences until the total number of tokens in each batch reaches 1024. We employ Adam

NEG	POS	LOSS	PTB
	\mathcal{S}	$\mathcal{L}_{\text{self}}$	81.08
$\max_{\mathcal{N} \cup \mathcal{P}}$	mean \mathcal{P}	\mathcal{L}_{sup}	95.58
	max \mathcal{P}	-	95.75
	\mathcal{S}	$\mathcal{L}_{\text{hash}}$	94.26
$\max_{\mathcal{N} \cup \mathcal{S}}$	mean \mathcal{P}	-	95.88
	max \mathcal{P}	\mathcal{L}_{max}	96.03

Table 3: Ablation study of instance selection strategies in constituency parsing experiments. The columns NEG and POS display the selection strategies for negatives and positives, respectively. LOSS shows this combination corresponds to which loss definition.

optimizer (Kingma and Ba, 2015; Loshchilov and Hutter, 2019), and the total number of training steps of constituency parsing and nested named entity recognition are 50,000 and 20,000, and the warm-up are 4000 and 2000 steps, respectively. To provide harder negatives by augmenting inputs, we also randomly mask a portion of tokens as [MASK].

Experiments are all conducted on a single NVIDIA Tesla V100 graphics card, the total training wall time is around 3 hours and 1 hour, respectively. All experiments are run with two different random seeds and the reported numbers in the following tables are their averages.

4.2 Main Results

Our model consistently achieves competitive performance on various structured prediction tasks and datasets, as presented in Table 1 and Table 2.

For constituency parsing, our models reach the peak with around 12 bits. Continuously increasing the number of bits does not further improve performance, on the contrary, it leads to a slight decline. We attribute this to the disproportionately large hashing space, as the amount of information carried by each task and dataset is fixed. For example, assigning K bits to a task with only K labels leads to an extreme case. Models in this case tend to produce the most trivial bit-level one-hot representation, making them nothing different from traditional static embedding models. On the contrary, decreasing the number of bits to fewer than 8 bits is also harmful, due to the insufficient representation capability. Besides, our models outperform almost all previous graph-based methods that rely on maximizing the log-likelihoods of target trees (Kitaev et al., 2019; Mrini et al., 2020; Zhang et al., 2020; Xin et al., 2021; Cui et al., 2022). There-

LABEL	CODE	COVERAGE (%)
S	110110001000	56.48
	110100001101	42.80
S'	101100001000	59.69
	110110101001	38.36
NP	010101001101	99.38
NP'	010001010100	98.70
VP	100010100111	99.52
VP'	001010010111	98.23
ADJP	100100000111	66.22
	000011010110	29.08
ADJP'	000000010110	93.34
ADVP	101100010111	84.53
	000101100111	11.87
ADVP'	001100010110	52.03
	000101110110	37.00

Table 4: Example of the hashing results on the constituency parsing task. The LABEL column shows the labels and their corresponding incomplete labels, which are introduced during the Chomsky normal form transformation. The CODE and COVERAGE columns display binary codes and their frequency proportions among all possible codes under each label. For instance, label S' is supposed to be assigned to an incomplete span within a larger S span, and 59.69% of S' labels are translated from the code 101100001000.

fore, we claim that leveraging contrastive learning is beneficial to representation learning.

For nested named entity recognition, all datasets show the best performance at the 8-bit settings, and decreasing to fewer than 6 bits also results in insufficient representing capability. Similarly, our methods outperform the previous sequential labeling methods (Shibuya and Hovy, 2020; Wang et al., 2021; Xin et al., 2021) and graph-based methods (Fu et al., 2021; Lou et al., 2022). In addition, some other papers (Yu et al., 2020; Shen et al., 2021; Tan et al., 2021; Zhu and Li, 2022) propose to straightforwardly enumerate all spans and directly train the model to classify them. These methods currently show the best performance, our method can also achieve comparable results to them.

4.3 Ablation Studies

We enumerate all combinations of the contrastive learning objective functions in Table 3 to study the influence of instance selection strategies.

We note that selecting $\mathcal{N} \cup \mathcal{S}$ is always superior to

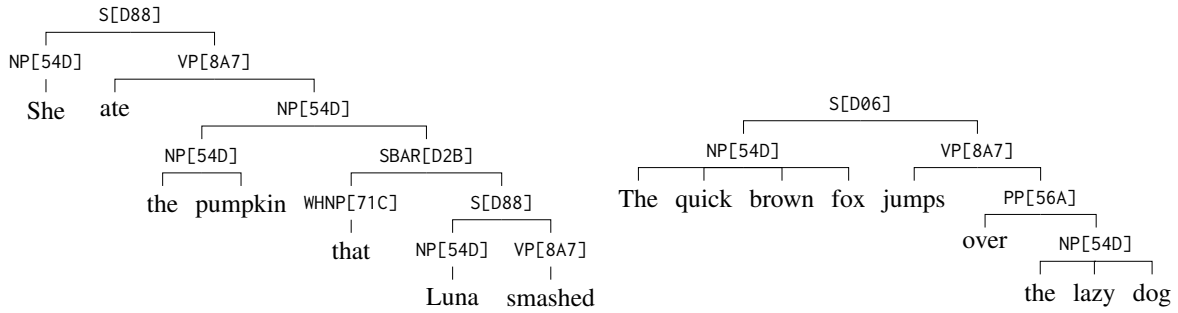


Figure 3: Examples of the hashing and constituency parsing results. The hexadecimal numbers in the brackets indicate the generated binary codes, and the span labels are translated from them.

selecting $\mathcal{N} \cup \mathcal{P}$ in the **negative term**, but using \mathcal{S} is inferior to using \mathcal{P} in the **positive term**. This leads to the conclusion that adding positive instances \mathcal{P} to the negative term is harmful, but it turns out to be beneficial when added to the positive term. This seems contradictory to the conclusion of Wang et al. (2023). However, this discrepancy arises because the training mechanisms are different. Their model relies on downstream task loss to implicitly learn structural information, whereas our model is designed to learn it explicitly. Therefore, we claim for the setting that aims at straightforwardly hashing both the label and structural information using a single loss, our \mathcal{L}_{\max} is the best choice.

Besides, Ouchi et al. (2020, 2021) appear to be the most relevant publications to our research. They also attempt to directly apply contrastive learning methods to structured prediction tasks. However, they pull instances towards the geometric centers, and their results show that this does not lead to noteworthy improvements than conventional graph-based models. This conclusion is consistent with ours, i.e., $\max_{\mathcal{P}}$ is better than $\text{mean}_{\mathcal{P}}$. We hypothesize that calculating the geometric center largely erases the contextual information of each instance. Pulling towards this center is not essentially different from pulling towards a static label embedding. Therefore, their approach loses the advantage of contrastive learning and trivially falls back to the conventional static embedding classification. On the contrary, our $\max_{\mathcal{P}}$ primarily targets the closest instance instead, therefore, it can fully leverage the capability of contrastive learning and distribute instances more uniformly, and allows each label to associate with multiple codes.

4.4 Case Studies

We show some frequent labels and their associated codes in Table 4. We only display the codes that

occupy the top 90% frequency of each label. It can be observed that some labels have two codes with roughly the same frequency, while others are predominantly represented by a single code. However, even in these cases, none reach 100% coverage. This result further demonstrates that the term $\max_{\mathcal{P}}$ imposes only a weak supervision signal and does not force all instances towards the geometric centers, thereby avoiding false positives and allowing a more uniform distribution of instances, rather than crowding at a few specific points.

In Figure 3, we display some parsing examples, using the same sentences as Wang et al. (2023). Our parser takes discrete tokens as inputs and returns binary labels, while theirs, conversely, receives binary codes of tokens and parses them into trees with discrete labels. Our codes preserve not only static but also contextual information of each label, for instance, the two S labels above are translated from two different codes. Therefore, our methods can also be considered as implicitly clustering methods, since these two codes refer to two different sub-labels of label S. This phenomenon is unlikely to be observed in models suffering from the geometric center issue.

5 Related Work

The ease of training has allowed contiguous representation (Hinton and Salakhutdinov, 2006) to achieve great success in applications such as masked language models (Devlin et al., 2019; Liu et al., 2020) and causal language models (Peters et al., 2018; Lewis et al., 2020; OpenAI et al., 2023; Touvron et al., 2023). Different from them, auto-encoder (Vincent et al., 2010; Kingma and Welling, 2014; Rezende et al., 2014) and contrastive learning (He et al., 2020; Gao et al., 2021) do not classify instances into specific classes, but rather reconstruct the input or determine whether a given pair

belongs to the same class. The advantage of these approaches is that they do not require introducing embedding matrices, making representation learning more effective. However, because embeddings are not introduced, continuous vectors are hard to be discretized to specific classes and need to train an additional classifier for such purpose.

Learning discrete representation has also been widely focused on since the early era of deep learning (Salakhutdinov and Hinton, 2009; Courville et al., 2011; Mnih and Gregor, 2014; Mnih and Rezende, 2016), as it is a necessary component of many downstream applications. For example, languages are inherently discrete, and mapping discrete tokens to continuous vectors and back is the foremost step for modern natural language processing. Additionally, many applications involving the control and interpretation of intermediate variables also require interleaving non-differentiable layers. However, the key challenge is that backpropagating gradients through discrete representation is generally intractable, making them hard to train.

The most straightforward solution is the straight-through estimator (Bengio et al., 2013; Yin et al., 2019), which simply copies gradients across non-differentiable layers, treating the Jacobian matrix as an identity matrix. SPIGOT (Peng et al., 2018; Mihaylova et al., 2020) further introduces a structured projection to reduce gradient errors. Building upon this previous work, VQ-VAE (van den Oord et al., 2017; Razavi et al., 2019) allows interleaving a discrete layer with continuous inputs mapped as discrete codes. These returned informative codes preserve necessary information for downstream tasks, and at the same time, provide an interface for controlling and interpreting the hidden states in intermediate layers of neural networks.

Another line of research turns to relaxation techniques, as they approximate discrete representation with continuous relaxation to mitigate the intractability issue. The Gumbel-softmax estimator (Jang et al., 2017) was first proposed by using the Gumbel-max trick (Maddison et al., 2014, 2017) to provide a differentiable approximation of the $\arg \max$ operation. The marginal approaches (Friesen and Domingos, 2016; Kim et al., 2017; Liu and Lapata, 2018; Eisner, 2016) even further relax the discreteness restriction, not limiting themselves to selecting only a single discrete token, but instead aggregating information from all discrete tokens according to softmax distribution.

Different from all of them, our binary represen-

tion lies between continuous and discrete representations. On the one hand, the easy-to-train property of continuous representation allows us to avoid designing complex gradient estimators while also benefiting from state-of-the-art representation learning methods, i.e., contrastive learning. On the other hand, the easy-to-discretize property allows us to instantly transform these continuous representations into discrete representations, thus making the imposition of various task-relevant structures tractable and efficient. Besides, in other potential applications such as machine translation and language models, using binary representation removes the need for a large softmax, thus, reduces memory consumption and accelerates training and inference. Additionally, the final step of re-interpreting binary codes as labels or tokens allows the model to preserve much more contextual information within the codes. Consequently, this implicit clustering capability may potentially address issues like polysemy through discovering subclasses (§4.4).

6 Conclusions

In this paper, we demonstrated that introducing the binary representation to the output side is also feasible. Binary representation inherits the advantages of easy training from contiguous representation and the structure modeling capability from discrete representation. Therefore, binary representation can be considered as a novel representation that further bridges the gap between the continuous nature of deep learning and the discrete intrinsic property of natural languages. We achieved this by extending previous contrastive hashing into structured contrastive hashing. Specifically, we designed a bit-level CKY, defined a similarity function based on marginal probability, and proposed a novel contrastive hashing loss to mitigate the geometric center issue. Experiments show that our methods perform remarkably on various structured prediction tasks by using only a few bits. Our model also demonstrates a certain degree of implicit clustering capability in discovering subclasses. Most importantly, although we primarily focus on structured prediction tasks, our method can be easily applied to other natural language processing tasks by equipping the corresponding structures.

Limitations

Our contrastive hashing method yields informative binary representations. Nonetheless, information

theory suggests that labels with different frequencies carry varying amounts of information. Therefore, compressing all continuous label representation into codes with fixed number of bits is not the optimal solution, and it may even lead to the problem of high-frequency labels being overly concentrated in a narrow area. We remain to solve this issue by extending our method to a variant number of bits in future work.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *Preprint*, arXiv:1607.06450.
- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. 2000. [A neural probabilistic language model](#). In *Advances in Neural Information Processing Systems*, volume 13. MIT Press.
- Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. [Estimating or propagating gradients through stochastic neurons for conditional computation](#). *CoRR*, abs/1308.3432.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Aaron Courville, James Bergstra, and Yoshua Bengio. 2011. [A spike and slab restricted boltzmann machine](#). In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 233–241, Fort Lauderdale, FL, USA. PMLR.
- Leyang Cui, Sen Yang, and Yue Zhang. 2022. [Investigating non-local features for neural constituency parsing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2065–2075, Dublin, Ireland. Association for Computational Linguistics.
- 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. [The automatic content extraction \(ACE\) program – tasks, data, and evaluation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Jason Eisner. 2016. [Inside-outside and forward-backward algorithms are just backprop \(tutorial paper\)](#). In *Proceedings of the Workshop on Structured Prediction for NLP*, pages 1–17, Austin, TX. Association for Computational Linguistics.
- Abram L. Friesen and Pedro M. Domingos. 2016. [The sum-product theorem: A foundation for learning tractable models](#). *CoRR*, abs/1611.03553.
- Yao Fu, Chuanqi Tan, Mosha Chen, Songfang Huang, and Fei Huang. 2021. [Nested named entity recognition with partially-observed treecrfs](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12839–12847.
- David Gaddy, Mitchell Stern, and Dan Klein. 2018. [What’s going on in neural constituency parsers? an analysis](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 999–1010, New Orleans, Louisiana. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. [Momentum contrast for unsupervised visual representation learning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- G. E. Hinton and R. R. Salakhutdinov. 2006. [Reducing the dimensionality of data with neural networks](#). *Science*, 313(5786):504–507.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *International Conference on Learning Representations*.
- Tadao Kasami. 1966. [An efficient recognition and syntax-analysis algorithm for context-free languages](#). *Coordinated Science Laboratory Report no. R-257*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. [Supervised contrastive learning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.
- J-D Kim, Tomoko Ohta, Yuka Tateisi, and Jun’ichi Tsujii. 2003. [Genia corpus—a semantically annotated corpus for bio-textmining](#). *Bioinformatics*, 19(suppl_1):i180–i182.

- Yoon Kim, Carl Denton, Luong Hoang, and Alexander M. Rush. 2017. [Structured attention networks](#). In *International Conference on Learning Representations*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations*.
- Diederik P Kingma and Max Welling. 2014. [Auto-encoding variational bayes](#). In *International Conference on Learning Representations*.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. [Multilingual constituency parsing with self-attention and pre-training](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- 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*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2018. [Learning structured text representations](#). *Transactions of the Association for Computational Linguistics*, 6:63–75.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Ro{bert}a: A robustly optimized {bert} pretraining approach](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Chao Lou, Songlin Yang, and Kewei Tu. 2022. [Nested named entity recognition as latent lexicalized constituency parsing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6183–6198, Dublin, Ireland. Association for Computational Linguistics.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. [The concrete distribution: A continuous relaxation of discrete random variables](#). In *International Conference on Learning Representations*.
- Chris J Maddison, Daniel Tarlow, and Tom Minka. 2014. [A* sampling](#). In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Tsvetomila Mihaylova, Vlad Niculae, and André F. T. Martins. 2020. [Understanding the mechanics of SPIGOT: Surrogate gradients for latent structure learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2186–2202, Online. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient estimation of word representations in vector space](#). *Preprint*, arXiv:1301.3781.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. [Distributed representations of words and phrases and their compositionality](#). *Preprint*, arXiv:1310.4546.
- Andriy Mnih and Karol Gregor. 2014. [Neural variational inference and learning in belief networks](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1791–1799, Beijing, China. PMLR.
- Andriy Mnih and Danilo Rezende. 2016. [Variational inference for monte carlo objectives](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2188–2196, New York, New York, USA. PMLR.
- Khalil Mrini, Franck Dernoncourt, Quan Hung Tran, Trung Bui, Walter Chang, and Ndapa Nakashole. 2020. [Rethinking self-attention: Towards interpretability in neural parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 731–742, Online. Association for Computational Linguistics.
- Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty, and Xiaoli Li. 2021. [A conditional splitting framework for efficient constituency parsing](#). In *Proceedings of the 59th Annual Meeting of the Association for*

Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 5795–5807, Online. Association for Computational Linguistics.

- OpenAI, :, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madeleine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondrasiuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameez Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2023. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Hiroki Ouchi, Jun Suzuki, Sosuke Kobayashi, Sho Yokoi, Tatsuki Kuribayashi, Ryuto Konno, and Kentaro Inui. 2020. [Instance-based learning of span representations: A case study through named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6452–6459, Online. Association for Computational Linguistics.
- Hiroki Ouchi, Jun Suzuki, Sosuke Kobayashi, Sho Yokoi, Tatsuki Kuribayashi, Masashi Yoshikawa, and Kentaro Inui. 2021. [Instance-based neural dependency parsing](#). *Transactions of the Association for Computational Linguistics*, 9:1493–1507.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Hao Peng, Sam Thomson, and Noah A. Smith. 2018. [Backpropagating through structured argmax using a SPIGOT](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1863–1873, Melbourne, Australia. Association for Computational Linguistics.
- Matthew E. 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)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Ali Razavi, Aaron van den Oord, and Oriol Vinyals. 2019. [Generating diverse high-fidelity images with VQ-VAE-2](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. [Stochastic backpropagation and approximate inference in deep generative models](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1278–1286, Beijing, China. PMLR.
- Ruslan Salakhutdinov and Geoffrey Hinton. 2009. [Deep boltzmann machines](#). In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 448–455, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA. PMLR.
- 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)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Yongliang Shen, Xinyin Ma, Zeqi Tan, Shuai Zhang, Wen Wang, and Weiming Lu. 2021. [Locate and label: A two-stage identifier for nested named entity recognition](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2782–2794, Online. Association for Computational Linguistics.
- Takashi Shibuya and Eduard Hovy. 2020. [Nested named entity recognition via second-best sequence learning and decoding](#). *Transactions of the Association for Computational Linguistics*, 8:605–620.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15(56):1929–1958.
- Mitchell Stern, Jacob Andreas, and Dan Klein. 2017. [A minimal span-based neural constituency parser](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 818–827, Vancouver, Canada. Association for Computational Linguistics.
- Zeqi Tan, Yongliang Shen, Shuai Zhang, Weiming Lu, and Yueting Zhuang. 2021. [A sequence-to-set network for nested named entity recognition](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3936–3942. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Yuanhe Tian, Yan Song, Fei Xia, and Tong Zhang. 2020. [Improving constituency parsing with span attention](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1691–1703, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. 2017. [Neural discrete representation learning](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. 2010. [Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion](#). *Journal of Machine Learning Research*, 11(110):3371–3408.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. [Ace 2005 multilingual training corpus ldc2006t06](#), 2006. *URL* <https://catalog.ldc.upenn.edu/LDC2006T06>.
- Jue Wang, Lidan Shou, Ke Chen, and Gang Chen. 2020. [Pyramid: A layered model for nested named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5918–5928, Online. Association for Computational Linguistics.
- Yiran Wang, Hiroyuki Shindo, Yuji Matsumoto, and Taro Watanabe. 2021. [Nested named entity recognition via explicitly excluding the influence of the](#)

- best path. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3547–3557, Online. Association for Computational Linguistics.
- Yiran Wang, Taro Watanabe, Masao Utiyama, and Yuji Matsumoto. 2023. [24-bit languages](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 408–419, Nusa Dua, Bali. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Xin Xin, Jinlong Li, and Zeqi Tan. 2021. [N-ary constituent tree parsing with recursive semi-Markov model](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2631–2642, Online. Association for Computational Linguistics.
- Yongxiu Xu, Heyan Huang, Chong Feng, and Yue Hu. 2021. [A supervised multi-head self-attention network for nested named entity recognition](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14185–14193.
- Naiwen Xue, Fei Xia, Fu-Dong Chiou, and Marta Palmer. 2005. [The penn chinese treebank: Phrase structure annotation of a large corpus](#). *Natural Language Engineering*, 11(2):207–238.
- Kaiyu Yang and Jia Deng. 2020. [Strongly incremental constituency parsing with graph neural networks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 21687–21698. Curran Associates, Inc.
- Songlin Yang and Kewei Tu. 2022. [Bottom-up constituency parsing and nested named entity recognition with pointer networks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2403–2416, Dublin, Ireland. Association for Computational Linguistics.
- Songlin Yang and Kewei Tu. 2023. [Don’t parse, choose spans! continuous and discontinuous constituency parsing via autoregressive span selection](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8420–8433, Toronto, Canada. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Penghang Yin, Jiancheng Lyu, Shuai Zhang, Stanley J. Osher, Yingyong Qi, and Jack Xin. 2019. [Understanding straight-through estimator in training activation quantized neural nets](#). In *International Conference on Learning Representations*.
- Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. [Named entity recognition as dependency parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6470–6476, Online. Association for Computational Linguistics.
- Yu Zhang, Houquan Zhou, and Zhenghua Li. 2020. [Fast and accurate neural crf constituency parsing](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4046–4053. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Junru Zhou and Hai Zhao. 2019. [Head-Driven Phrase Structure Grammar parsing on Penn Treebank](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2396–2408, Florence, Italy. Association for Computational Linguistics.
- Enwei Zhu and Jinpeng Li. 2022. [Boundary smoothing for named entity recognition](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7096–7108, Dublin, Ireland. Association for Computational Linguistics.

A Dataset Statistics

DATASET	TRAIN	DEV	TEST	LABEL	CNF
PTB	39,832	1700	2416	26	130
CTB	18,104	352	348	26	133
ACE’4	6198	742	809	8	22
ACE’5	7285	968	1058	8	39
GENIA	15,022	1669	1855	5	21

Table 5: Statistics of these five datasets. Columns TRAIN, DEV, and TEST show the number of sentences, LABEL and CNF display the number of labels before and after Chomsky normal form transformation.