



Full length article

Dependency parsing with bottom-up Hierarchical Pointer Networks

Daniel Fernández-González*, Carlos Gómez-Rodríguez

Universidade da Coruña, CITIC, FASTPARSE Lab, LyS Group, Depto. de Ciencias de la Computación y Tecnologías de la Información, Campus de Elviña, s/n, A Coruña, 15071, Spain



ARTICLE INFO

Keywords:

Natural language processing
Computational linguistics
Parsing
Dependency parsing
Neural network
Deep learning

ABSTRACT

Dependency parsing is a crucial step towards deep language understanding and, therefore, widely demanded by numerous Natural Language Processing applications. In particular, left-to-right and top-down transition-based algorithms that rely on Pointer Networks are among the most accurate approaches for performing dependency parsing. Additionally, it has been observed for the top-down algorithm that Pointer Networks' sequential decoding can be improved by implementing a hierarchical variant, more adequate to model dependency structures. Considering all this, we develop a bottom-up oriented Hierarchical Pointer Network for the left-to-right parser and propose two novel transition-based alternatives: an approach that parses a sentence in right-to-left order and a variant that does so from the outside in. We empirically test the proposed neural architecture with the different algorithms on a wide variety of languages, outperforming the original approach in practically all of them and setting new state-of-the-art results on the English and Chinese Penn Treebanks for non-contextualized and BERT-based embeddings.

1. Introduction

Dependency parsing consists in representing the grammatical structure of a given sentence by attaching each word to another (which will be considered its *head* or *parent*), to finally gather all these directed links into a dependency tree as the one included in Fig. 1. Additionally, these directed links or dependencies are enhanced with *labels* that describe syntactic functions.

This syntactic information accurately provided by parsers in the form of *dependency trees* has proven highly useful for a huge variety of Natural Language Processing (NLP) tasks. In particular, dependency parsing has been recently used for machine translation [1], opinion mining [2,3], event extraction [4], question answering [5], sentiment analysis [6–8], coreference resolution [9], summarization [10] or semantic role labeling and named entity recognition [11], among others.

Pointer Networks [12] have notably succeeded in implementing highly accurate versions of one of the most widely used dependency parsing paradigms: *transition-based* dependency parsers. In particular, two different algorithms have been proposed: a *top-down* approach [13] that, at each step and starting from the root node, connects each word to one of its children; and a *left-to-right* variant [14] that, starting

from the left, attaches each word of the sentence to its parent. Apart from being twice as fast as the former, the latter outperforms the top-down variant in terms of accuracy in practically all datasets tested so far.

As transition-based algorithms [15], both perform the parsing process as a sequential decoding where, at each step or parsing state, all permissible actions (which do not violate the single-head or acyclicity constraints) are evaluated and the highest-scoring one is greedily applied, generating a new state. This sequential decoding differs from the approach followed by their main competitors: *graph-based* algorithms [16]. These previously score all possible arcs (or sets of arcs) and then, during decoding, search for the highest-scoring valid dependency tree. While recent graph-based models use a head-selection strategy [17,18] similar to that followed by the left-to-right transition-based parser [14], the training and decoding processes are completely different: at each decoding step, only permissible arc transitions are evaluated in the transition-based approach, and scores are computed taking into account a sequence of previous decisions encoded through the decoder. On the contrary, these simplified graph-based models independently score all possible parents for each word and then, during decoding, only the highest-scoring ones are kept regardless of the parents chosen for the other words. Then, at the end of the process, a maximum spanning tree algorithm is applied (if necessary) to output

* Corresponding author.

E-mail addresses: d.fgonzalez@udc.es (D. Fernández-González), carlos.gomez@udc.es (C. Gómez-Rodríguez).

URLs: <https://danifg.github.io> (D. Fernández-González), <http://www.grupolys.org/~cgomezr/> (C. Gómez-Rodríguez).




Fig. 1. Dependency tree for a given sentence.

a well-formed tree.¹ This difference can be also seen empirically, as the left-to-right transition-based parser achieves higher accuracies than the approach by Dozat and Manning [18] in several widely-known benchmarks.

While this sequential decoding (typically implemented by a recurrent neural network) seems to be beneficial for dependency parsing, it might lead to accuracy losses due to error propagation: mistaken past decisions will affect future actions, especially harming performance on long-range dependencies and attachments created in final steps. This limitation present in classic transition-based parsers also affects recent algorithms based on Pointer Networks, as decoder states located at the end of the sequence tend to forget relevant information from the past and are more prone to suffer from error propagation.

In fact, Liu et al. [19] propose a *hierarchical* decoding for the top-down algorithm [13] by having access, at each step, to information about the focus word’s parent and siblings created in the past, introducing not only knowledge about distant decoder states relevant for future decisions, but also an underlying tree structure to the decoding process (more appropriate for modeling dependency graphs).

In this paper, we initially develop a general Hierarchical Pointer Network architecture with a bottom-up structured decoding for the left-to-right transition-based algorithm [14]. This will allow the decoder to have access to information about partial structures created in the past and will help to make better future decisions. Alternatively to this algorithm, we design two novel bottom-up oriented transition systems that can be easily implemented on the proposed neural model.

Finally, we empirically show that the presented architecture with any transition-based algorithm provides improvements in accuracy on ten different languages from Universal Dependencies [20] and obtains state-of-the-art scores on the commonly-used English and Chinese Penn Treebanks [21,22].

The remainder of this article is organized as follows: Section 2 introduces the baseline left-to-right parser by Fernández-González and Gómez-Rodríguez [14] and briefly presents the top-down Hierarchical Pointer Network by Liu et al. [19]. In Section 3, we describe in detail the proposed bottom-up Hierarchical Pointer Network and how it is adapted to the left-to-right algorithm. Section 4 presents new transition systems and how they are implemented on the novel neural architecture. In Section 5, we extensively evaluate the proposed neural model with each parsing strategy on numerous datasets, as well as include a thorough analysis of their performance. Lastly, Section 6 contains a final discussion.

2. Preliminaries

2.1. Left-to-right transition system

Fernández-González and Gómez-Rodríguez [14] propose an efficient left-to-right transition system that is defined by a *focus word pointer* i , which is used to point to the word currently being processed w_i , and a single SHIFT-ATTACH- p transition, which assigns a parent word w_p (in

¹ Additionally, while in transition-based parsers like [13] and [14], the sequential decoding makes it possible to easily add high-order features and projectivity constraints without increasing runtime complexity, these enhancements in graph-based models lead to a performance penalty.




Fig. 2. Attachment order and leftmost (lm), rightmost (rm), last-attached left (la) and right (ra) dependents that are available at each decoding step when parsing the sentence in Fig. 1 with the three available transition-based algorithms. Note that words w_i are represented by the positional index i .

position p) to w_i (producing the dependency arc $w_p \rightarrow w_i$) and then moves the pointer making it point to the next token w_{i+1} . Starting at the beginning of a sentence (of length n), this algorithm sequentially attaches each token w_i to its parent in just n steps.² In Fig. 2(a), we depict the sequential prediction of SHIFT-ATTACH- p transitions that produces all arcs in the dependency tree in Fig. 1 following a left-to-right transition system.

In order to incrementally build a well-formed dependency tree during decoding, only SHIFT-ATTACH- p transitions that do not generate cycles in the already-built dependency graph are allowed. Additionally, the left-to-right transition system can efficiently produce unrestricted dependency graphs, including non-projective structures.³ However, the projectivity constraint can optionally be enforced by discarding transitions that produce crossing dependencies, which can be useful in treebanks with a negligible presence of non-projectivity [23].

2.2. Pointer networks for left-to-right dependency parsing

The left-to-right algorithm can be easily implemented by a Pointer Network [12] to undertake non-projective dependency parsing. Pointer Networks are one among the various recent neural models that seek to improve performance on various Artificial Intelligence tasks [24,25]. These sequence-to-sequence neural networks are trained to output a sequence of positions (discrete numbers) from an input sequence of items. In our case, the output is the sequence $\mathbf{p} = p_1, p_2, \dots, p_n$ of values necessary to parameterize the SHIFT-ATTACH- p transition and build a dependency graph for the input sentence $\mathbf{w} = w_1, \dots, w_n$. These positions (or values of p in the left-to-right transition system) are

² The top-down transition system by Ma et al. [13] needs $2n-1$ steps to parse a sentence of length n .

³ Non-projective dependency graphs are harder to produce since they are able to represent more complex syntactic phenomena by allowing crossing arcs.

selected by the attention mechanism [26] over the input sentence. This avoids having to keep a fixed output dictionary size, which will vary for each input sentence with respect to its length.

More specifically, the encoder–decoder architecture of the original left-to-right parser [14] is designed as follows:

Encoder Given an input sentence $\mathbf{w} = w_1, \dots, w_n$, each word w_i is initially represented as x_i : i.e., a concatenation of character-level representations (e_i^c) (extracted by max-pooling-based Convolutional Neural Networks [27]), POS tag embeddings (e_i^p) and word embeddings (e_i^w):

$$x_i = e_i^c \oplus e_i^w \oplus e_i^p$$

Then, each vector representation x_i is passed through a BiLSTM (*Bidirectional Long Short-Term Memory*) to generate the encoder hidden state \mathbf{h}_i .

$$\mathbf{h}_i = \mathbf{h}_{li} \oplus \mathbf{h}_{ri} = \text{BiLSTM}(x_i)$$

where \mathbf{h}_{li} and \mathbf{h}_{ri} are the outputs of a leftward and a rightward unidirectional LSTM (*Long Short-Term Memory*), respectively. This results in the sequence $\mathbf{h}_0 \dots \mathbf{h}_n$ that encodes the input sentence (where \mathbf{h}_0 is a special representation for the artificial ROOT node).

We also report accuracies with our encoder enhanced with contextualized word embeddings extracted from the widely-used language model BERT [28]. In those cases, BERT-based word embeddings are directly concatenated to the basic word representation x_i before passing it through the BiLSTM:

$$x'_i = x_i \oplus e_i^{BERT}$$

$$\mathbf{h}_i = \text{BiLSTM}(x'_i)$$

Decoder The sequential decoding that models the transition-based behavior is implemented by a unidirectional LSTM, which is in charge of generating a new *decoder hidden state* s_t at each time step t . As input it receives the encoder hidden state \mathbf{h}_i (representation of the current focus word w_i) and, as a standard recurrent neural network, it is also conditioned by the previous decoder state s_{t-1} :

$$s_t = \text{LSTM}(\mathbf{h}_i) = f(s_{t-1}, \mathbf{h}_i)$$

In the original work, authors added to \mathbf{h}_i the previous and next encoder hidden states as extra features; however, in a followup paper, they removed those additions since the bidirectional LSTM is already encoding such information and, therefore, they do not lead to significant improvements [23].

The resulting decoder state s_t (which represents the current focus word plus the past decisions made so far) is used for computing scores of all possible words w_j from the input sentence (encoded as vectors \mathbf{h}_j with $j \in [0, n]$ and $j \neq i$) as parent of w_i . These scores are obtained by a biaffine scoring function [18] to finally compute the attention distribution in a vector \mathbf{a}_t :

$$v_{ij} = \text{score}(s_t, \mathbf{h}_j) = f_1(s_t)^T \mathbf{W} f_2(\mathbf{h}_j) + \mathbf{U}^T f_1(s_t) + \mathbf{V}^T f_2(\mathbf{h}_j) + \mathbf{b};$$

$$\mathbf{a}_t = \text{softmax}(v_t)$$

where \mathbf{W} , \mathbf{U} and \mathbf{V} are the weights and, as observed by Dozat and Manning [18], $f_1(\cdot)$ and $f_2(\cdot)$ are multilayer perceptrons (MLP) for reducing dimensionality and avoiding model overfitting.

Then, the attention vector \mathbf{a}_t is employed for implementing a pointer over the input sentence, selecting, at each step t , the highest-scoring position (where the parent word w_p is located) and, therefore, providing a value p necessary to apply a $\text{SHIFT-ATTACH-}p$ transition and build the arc $w_p \rightarrow w_i$. Finally, as proposed by [13], a multi-class classifier (based on the approach developed by Dozat and Manning [18]) is separately trained to predict the arc label of each dependency created by the pointer.

Since the decoder performs n attachments to process an input sentence with length n and, additionally, attention is computed at each step over the whole sentence, the time complexity of the parsing process is $O(n^2)$.

2.3. Top-down Hierarchical Pointer Networks

Liu et al. [19] introduced Hierarchical Pointer Networks for the top-down transition-based algorithm developed by Ma et al. [13]. The top-down transition system consists of two actions and a stack: one transition for connecting the word on top of the stack to one of its children and push it into the stack, and another action for popping the current focus word on top. Liu et al. [19] design a structured decoding for the original Pointer Network, where not only the immediately-previous decoder state s_{t-1} of the LSTM is considered for choosing the next action to be applied on the current word on top, but also decoder hidden states generated when the current focus word's parent and last sibling were assigned following the top-down transition system. More graphically, in Fig. 1, when, for instance, the current focus word to be processed is *John*, decoder hidden states of its parent *play* and last-assigned sibling *together*⁴ are taken into consideration when choosing *John*'s children. With this strategy, the authors manage to introduce an explicit structural inductive bias into the original linear decoder, allowing a more adequate modeling for the generation of dependency graphs. Liu et al. [19] also empirically show that their top-down Hierarchical Pointer Networks are beneficial for the accuracy of arcs created in final steps, which is especially crucial on long sentences.

3. Bottom-up Hierarchical Pointer Networks

Based on the research work by Liu et al. [19], we propose a Hierarchical Pointer Network for the left-to-right transition-based algorithm [14] and other related parsing strategies where dependents of the focus word may already have been attached (as we will design more such algorithms later). In particular, instead of parent or sibling information, we design a bottom-up structured tree decoding by considering decoder hidden states of the focus word's relevant dependents. More concretely, to alleviate the gradual loss of relevant information during sequential decoding, we propose to consider decoder hidden states of the leftmost and rightmost dependents already attached to the current focus word, as well as the most recently attached left and right dependents. Note that not all this information is available in all strategies, e.g., the left-to-right algorithm has no access to right dependencies of the focus word.

All this will not only provide valuable knowledge about words processed in the past (especially those attached by long-range dependencies), but will also keep an underlying tree structure that will help the decoder to make good decisions throughout the parsing process. For instance, when the sentence from Fig. 1 is parsed following a purely bottom-up strategy⁵ and the current focus word *play* must be attached to its parent, it would be helpful to have access to information about its leftmost and rightmost dependents (*John* and *together*, respectively), as well as the most recently attached right dependent (*tennis*).

3.1. Structured tree decoder

To implement this novel bottom-up Hierarchical Pointer Network variant, we keep the same encoder as the original approach [14], but use a bottom-up hierarchical decoder instead.

More in detail, at each step t with current focus word w_i , the generated decoder hidden state s_t is directly conditioned, apart from by the previous decoder state s_{t-1} and the encoder hidden state \mathbf{h}_i , by the decoder states of the already-attached leftmost ($s_{l_{m(t)}}$) and rightmost

⁴ Assuming that in the top-down transition system, right children are attached first in an inside-out order.

⁵ Please note that the left-to-right transition system is not purely bottom-up since, when a word is attached to its parent, not all of its children have necessarily been assigned yet.

($s_{rm(t)}$) dependents of w_i , as well as the decoder states of the left ($s_{la(t)}$) and right ($s_{ra(t)}$) last-attached dependents.⁶

$$s_t = f(s_{lm(t)}, s_{rm(t)}, s_{la(t)}, s_{ra(t)}, s_{t-1}, \mathbf{h}_t)$$

where $f(\cdot)$ is a *fusion function* that combines all components into a single decoder state.

Following Liu et al. [19], we do not directly feed these six components to the decoder, but utilize a *gating mechanism* to allow our model to adequately extract the most useful information at each decoding step. In particular, we implement two different gating functions:

$$\mathbf{g}_t = \text{sigmoid}(\mathbf{W}_{gp}s_{t-1} + \mathbf{W}_{glm}s_{lm(t)} + \mathbf{W}_{grm}s_{rm(t)} + \mathbf{W}_{gla}s_{la(t)} + \mathbf{W}_{gra}s_{ra(t)} + \mathbf{b}_g) \quad (\text{GATE1})$$

$$\mathbf{g}_t = \text{sigmoid}(\mathbf{W}_{glm}(s_{t-1} \odot s_{lm(t)}) + \mathbf{W}_{grm}(s_{t-1} \odot s_{rm(t)}) + \mathbf{W}_{gla}(s_{t-1} \odot s_{la(t)}) + \mathbf{W}_{gra}(s_{t-1} \odot s_{ra(t)}) + \mathbf{b}_g) \quad (\text{GATE2})$$

where \mathbf{W}_{gp} , \mathbf{W}_{glm} , \mathbf{W}_{grm} , \mathbf{W}_{gla} , \mathbf{W}_{gra} and \mathbf{b}_g are gating weights. While all decoder states are equally weighted in GATE1, the element-wise product used in GATE2 has the effect of similarity comparison among the previous decoder state s_{t-1} and the other components.

These gating functions are then used to define the fusion function f as follows:

$$\mathbf{h}'_t = \tanh(\mathbf{W}_p s_{t-1} + \mathbf{W}_{lm}s_{lm(t)} + \mathbf{W}_{rm}s_{rm(t)} + \mathbf{W}_{la}s_{la(t)} + \mathbf{W}_{ra}s_{ra(t)});$$

$$\mathbf{h}''_t = \mathbf{g}_t \odot \mathbf{h}'_t;$$

$$s_t = \text{LSTM}(\mathbf{h}''_t, \mathbf{h}_t)$$

where \mathbf{W}_{lm} , \mathbf{W}_{rm} , \mathbf{W}_{la} and \mathbf{W}_{ra} are the weights for combining the five different decoder states into one intermediate hidden state \mathbf{h}'_t . After that, the gating mechanism \mathbf{g}_t is applied to control the information flow and generate the hidden state \mathbf{h}''_t , which will be fed into the decoder together with the encoder hidden state \mathbf{h}_t .

As shown in Section 2.2, the resulting decoder state s_t is then used for computing the attention vector \mathbf{a}_t that will work as a pointer over the input sentence.

3.2. Model specifics for the left-to-right transition system

While the left-to-right transition system can be directly implemented on the described Hierarchical Pointer Network, it cannot use its full potential. This transition-based approach does not follow a fully bottom-up strategy during the left-to-right decoding⁷ and, therefore, right dependents are not available when they are needed. In the example, when the word *play* is under processing, only the leftmost dependent *John* is available since no right dependents have been created yet and, when the words *tennis* and *together* are finally attached, they are no longer needed since the word *play* will have already been processed. This led us to simplify the definition of the gating-based fusion function f for the left-to-right transition system by removing decoder states of right dependents:

$$s_t = f(s_{lm(t)}, s_{la(t)}, s_{t-1}, \mathbf{h}_t) \quad (\text{l-adapted})$$

Since, in some languages, it might be the case that words have either only one left dependent ($s_{lm(t)}$ and $s_{la(t)}$ being the same) or the left last-attached dependent was created in the previous time step (s_{t-1} and $s_{la(t)}$

⁶ Note that $s_{lm(t)}$ and $s_{rm(t)}$ might have the same values as $s_{la(t)}$ and $s_{ra(t)}$, respectively, when just one left or right dependent were assigned for the current focus word.

⁷ Left dependents are added bottom-up and right dependents top-down.

being the same), we also experiment with a simpler variant of f that just considers decoder states s_{t-1} and $s_{lm(t)}$:

$$s_t = f(s_{lm(t)}, s_{t-1}, \mathbf{h}_t) \quad (\text{l-simple})$$

Finally, it is worth mentioning that the hierarchical decoder does not penalize the quadratic runtime complexity of the left-to-right parser and, as the original work [14], it is trained by minimizing the total log loss (cross entropy) for choosing the correct sequence of SHIFT-ATTACH- p transitions to build a gold dependency tree for the input sentence \mathbf{w} (i.e. predicting the correct sequence of indices p , with each decision (p_t) being conditioned by previous ones ($p_{<t}$):

$$\mathcal{L}(\theta) = - \sum_{t=1}^T \log P_{\theta}(p_t | p_{<t}, \mathbf{w})$$

By optimizing the sum of their objectives, we simultaneously train the pointer and the labeler.

4. Alternative transition-based algorithms

Apart from the existing left-to-right algorithm, other transition systems can be implemented on the proposed bottom-up Hierarchical Pointer Network. In particular, we develop two alternative approaches to sequentially parse a sentence by attaching each word to its parent.

Right-to-left While a left-to-right hierarchical decoding might be more adequate for modeling left-branching languages where long-range dependencies tend to be leftward arcs (such as Turkish and Korean), for right-branching languages with a high predominance of long rightward arcs (such as Arabic and Hebrew), a right-to-left transition system could be a perfect fit.⁸ This parses a sentence starting from the last word w_n and it provides a SHIFT-ATTACH- p transition that, at each step, assigns a parent to the current focus word w_i and moves i to point to word w_{i-1} . In Fig. 2(b), we describe the order in which arcs are created with the right-to-left transition system to generate the dependency tree in Fig. 1 and the available dependents at each decoding step. Symmetrically to the left-to-right algorithm, we adjust the fusion function f to process only decoder states of right dependents:

$$s_t = f(s_{rm(t)}, s_{ra(t)}, s_{t-1}, \mathbf{h}_t) \quad (\text{r-adapted})$$

Additionally, following the same reasoning as for the left-to-right algorithm, we also implement a simplified alternative that removes $s_{ra(t)}$ from the equation:

$$s_t = f(s_{rm(t)}, s_{t-1}, \mathbf{h}_t) \quad (\text{r-simple})$$

Outside-in In order to fully use the proposed neural architecture and, therefore, have access to left and right dependents during the decoding process, we design an *outside-in* transition system: it parses a sentence starting from the leftmost word (w_1) and continuing with the rightmost word (w_n) towards the middle of the sentence, alternating words from the left and the right. To that end, a canonical order $w_1, w_n, w_2, w_{n-1}, \dots, w_{\lfloor \frac{n+1}{2} \rfloor}$ is determined at the beginning. At each decoding step the SHIFT-ATTACH- p transition, apart from attaching w_i to its parent, moves i according to that order. In Fig. 2(c), we can see an example of how this outside-in algorithm works and how left and

⁸ Please note that, following common usage in the parsing literature, naming conventions in this paper use the terms “left” and “right” to refer to temporal order (under the convention that a word is to the “left” of another if it comes first), not to spatial order when the words are written in a given script. This means that in languages with a right-to-left writing system, such as Arabic and Hebrew, the original left-to-right parser will begin reading each line of text from the right (i.e., the beginning), and the right-to-left algorithm will begin from the left (i.e., the end). Similarly, the graphical representation of what we call a “leftward arc” would be an arrow pointing to the right in those languages.

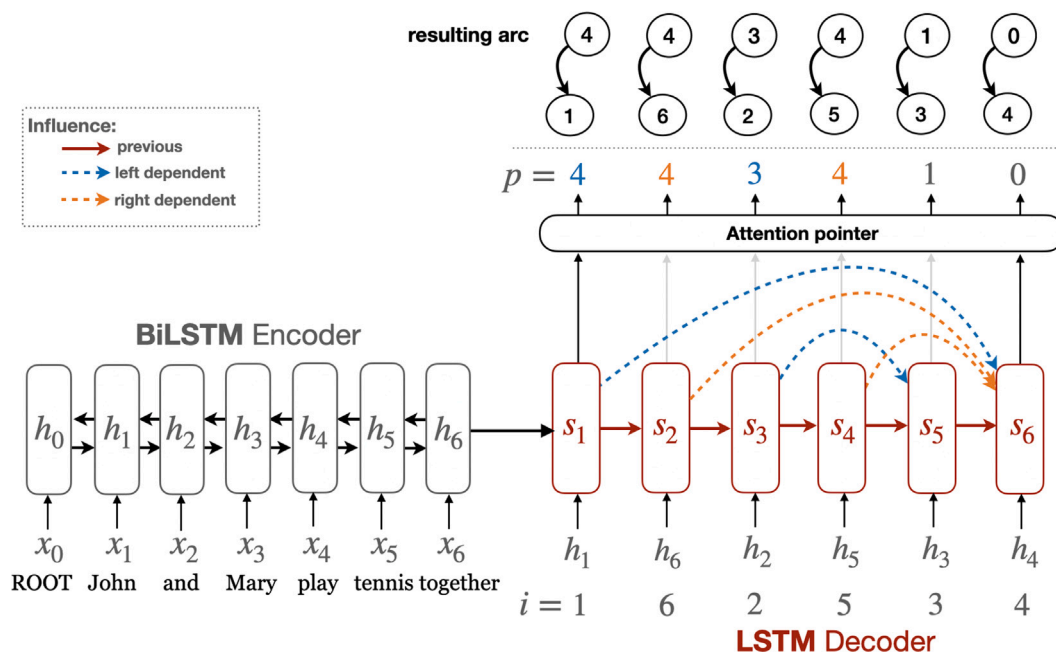


Fig. 3. Neural architecture for the Hierarchical Pointer Network and decoding steps necessary to produce the dependency tree in Fig. 1 with the outside-in transition system. Note that some dependent influences, such as the one generated by arc 1 → 3 in decoder state s_5 , will not be used as the word in position 1 was already processed.

Table 1

Details of the treebanks used in our experiments. Family = Afro-Asiatic (AA), Indo-European (IE), Koreanic (KO), Language isolate (LI), Sino-Tibetan (ST), Turkic (TU) or Uralic (UR). Order = dominant word order according to WALS [29]. Size = number of training sentences. %long = percentage of long arcs (length > 4) in the dev split. %left = percentage of long arcs that are leftward.

Language	Treebank	Family	Order	Size	%long	%left
Arabic	PADT	AA	VSO	6.1k	16.62	10.30
Basque	BDT	LI	SOV	5.4k	20.17	44.46
Chinese	GSD	ST	SVO	4.0k	25.31	59.44
English	EWT	IE	SVO	12.5k	16.43	25.49
Finnish	TDT	UR	SVO	12.2k	17.45	26.83
Hebrew	HTB	AA	SVO	5.2k	17.82	20.60
Italian	ISDT	IE	SVO	13.1k	16.20	22.69
Korean	GSD	KO	SOV	4.4k	14.76	78.40
Swedish	Talbanken	IE	SVO	4.3k	19.22	29.51
Turkish	IMST	TU	SOV	3.7k	14.84	70.69

right dependents are available during the decoding process. Since this strategy allows us to use decoder states of left and right dependents, the fusion function f can be fully used. For the same reasons as stated for the other transition systems, we also propose a simplification that discards last-attached dependents as follows:

$$s_t = f(s_{lm(t)}, s_{rm(t)}, s_{t-1}, \mathbf{h}_t) \tag{simple}$$

Note that both alternative transition systems are guaranteed to parse a sentence of length n in just n steps, keeping the same runtime complexity as the left-to-right algorithm, and are likewise trained, except for the order in which transitions (values of p) are predicted.

Fig. 3 depicts the proposed Hierarchical Pointer Network and presents the decoding procedure for producing the dependency tree described in Fig. 1 with the outside-in transition system. In this sketch, it can be graphically seen how decoder hidden states of dependents influence the decoding process.

Finally, it is worth mentioning that it is not possible to design a purely bottom-up transition system on the Pointer Network framework using n transitions (one per word) as, regardless of the order in which words are considered, it is not possible to guarantee that dependents

are always processed before heads (e.g., the first node considered will not necessarily be a leaf).

5. Experiments

5.1. Data

We conduct experiments on a wide variety of languages from Universal Dependencies v2.6 [20]. Following [30], we choose ten treebanks from different language families, with different morphological complexity and with different predominances of long-range dependencies. These are detailed in Table 1.

Additionally, we evaluate the proposed neural architecture on two widely-known benchmarks: the English Penn Treebank (PTB) [21] version with Stanford Dependencies [31] and, following [23], without any PoS tags⁹; and the dependency conversion [32] of the Chinese Penn Treebank 5.1 (CTB) [22] with gold POS tags.

Since random initializations are performed, we report the average Labeled and Unlabeled Attachment Scores (LAS and UAS) over 3 repetitions for each experiment; and, during evaluation, we just exclude punctuation on PTB and CTB following standard practice.

5.2. Settings

We use the code of the original parser by Fernández-González and Gómez-Rodríguez [14] (with the modifications proposed for the single-task parser in [23]) and implement the bottom-up Hierarchical Pointer Network plus the two novel transition systems, allowing a homogeneous comparison against the baseline. Similarly to [19], we experiment with the two available gate mechanisms in combination with the proposed transition systems and fusion function implementations. All models were executed on an Intel(R) Core(TM) i7-8700K CPU @ 3.70 GHz with two 12 GB GeForce GTX 1080 Ti GPUs.

⁹ It has been shown that using predicted PoS tags does not lead to accuracy gains in parsers built on Pointer Networks [13].

Table 2
Hyper-parameter selection.

Adam optimizer	
Batch size	32
β_1	0.9
β_2	0.9
Initial learning rate	0.001
Gradient clipping	5.0
Decay rate	0.75
Architecture	
BiLSTM size	512
BiLSTM number of layers	3
LSTM size	512
LSTM number of layers	1
LSTM layers dropout	0.33
CNN window size	3
CNN number of filters	50
MLP activation function	ELU
MLP number of layers	1
Arc MLP size	512
Label MLP size	128
UNK replacement probability	0.5
Character embedding dimension	100
POS tag embedding dimension	100
Word embedding dimension	100
Embeddings dropout	0.33
English BERT embedding dimension	1024
Chinese BERT embedding dimension	768

While character and PoS tag embeddings are randomly initialized, we use pre-trained embeddings for initializing word vectors: structured-skipgram embeddings [33] for English and Chinese, and Polyglot embeddings [34] for other languages. All embeddings are fine-tuned during training.

For PTB and CTB, we additionally concatenate contextualized word embeddings obtained from the pre-trained language model BERT [28]. We follow the greener and less resource-consuming approach undertaken by [23] and directly feed BERT fixed weights as described in Section 2.2, without any fine-tuning to our specific task. More specifically, we extract a combination of weights from layers 17–20 of BERT_{LARGE} for PTB and weights from the 11th layer of BERT_{BASE} for CTB, averaging BERT-based embeddings when a word is tokenized into more than one subword.

Finally, we use beam size 10 for PTB and CTB, and 1 for Universal Dependencies treebanks; we employ the Adam optimizer [35] for parameter optimization and further details about hyper-parameter selection are reported in Table 2.

5.3. Results

Table 3 reports LAS on dev splits of Universal Dependencies treebanks obtained by the proposed bottom-up Hierarchical Pointer Networks with different transition systems and parser configurations (gate and fusion function) and choose the best configuration on average to be evaluated on test splits. As shown in the reported results, the proposed architecture with the three algorithms improves over the baseline parser in practically all datasets, obtaining higher gains on those languages with a larger amount of long arcs (such as Basque and Chinese). We can also observe that the simplified fusion functions improve over the adapted and full versions on average, meaning that long-distance dependents are more valuable than last-assigned ones and, in some cases, the usage of the latter harms parsing accuracy. Regarding the novel transition systems, the right-to-left parser outperforms other algorithms on languages with a significant predominance of long rightward arcs (such as Arabic), and the outside-in algorithm clearly underperforms the other transition systems on average in spite of having access to left and right relevant dependents, which probably means that a sequential human-like strategy is more suitable for parsing

natural languages, or that the outside-in order is too complex to learn effectively. Finally, GATE1 obtains higher accuracy in general.

We also compare the proposed approach with the three implemented transition systems against the most recent dependency parsers on PTB and CTB. In order to choose the best parser configuration for the Penn Treebanks, we use scores on PTB development splits reported in Table 4. Since GATE1 outperforms GATE2 in general (as shown for Universal Dependencies treebanks in Table 3), we simply use GATE1 in all experiments and vary the fusion function implementation. Lastly, for BERT augmentations, we directly use simplified fusion functions, significantly reducing training time. As we can see in Table 5, the left-to-right approach achieves the highest accuracy obtained so far on PTB and CTB test splits with neither contextualized word embeddings nor extra constituent information. Regarding BERT augmentations, all transition systems not only outperform the baseline [23], but also the right-to-left model achieves the best LAS to date among approaches that use (or fine-tune) BERT (even improving over those enhanced with constituent information on PTB). These results provide some evidence that our proposed neural architecture has access to some structural information not learnt by the pre-trained language model BERT. Finally, all these improvements in accuracy provided by the proposed neural architecture come without harming the $O(n^2)$ runtime complexity of the original left-to-right transition-based parser, overcoming, therefore, their main competitors that have a higher $O(n^3)$ runtime complexity: the second-order graph-based parsers by Zhang et al. [40] and Wang and Tu [41].

Finally, while the runtime complexity of the proposed transition-based parsers is quadratic as that of the original sequential approach, the leverage of structural information during decoding comes at a cost in practice. In particular, the left-to-right, right-to-left and outside-in algorithms implemented on the Hierarchical Pointer Network respectively deliver 9.62 sent./s., 9.82 sent./s. and 8.81 sent./s. during decoding time on the test split of the PTB. This means that the original 22.57-sentences-per-second speed of the sequential Pointer Network is approximately halved. Although the current version of our neural architecture is less efficient than the baseline, it is worth mentioning that it was not optimized for speed and, therefore, further enhancements in handling structural information would improve final decoding speed.

5.4. Dependent information availability

To further understand the availability of dependent information for each transition system and its impact on parsing performance, we report in Table 6 the number of dependents per sentence that are accessible for each transition system on Universal Dependencies treebanks. While, at a first glance, we might think that the outside-in algorithm is the strategy that leverages dependent information the most since it receives both left and right dependents, it seems that, overall and according to gold trees, the left-to-right parser is the strategy that has access to a larger amount of (in this case, left) dependents, and the right-to-left approach is the one that uses more long-range (right) dependents (typically present in right-branching languages, where rightward dependency arcs tend to be longer). This might explain the good results of the right-to-left approach on Universal Dependencies datasets on average (since the majority of tested languages have a higher percentage of long rightward arcs and long-range dependents are considered more valuable for reducing error propagation) and the fact that the l-adapted fusion function has a better performance on the left-to-right parser (since this function gathers information about closer dependents and this algorithm has access to a notable amount of dependent information). Finally, we can also observe that, as expected, the right-to-left algorithm is more adequate to model languages with a high predominance of rightward arcs such as Hebrew and especially Arabic, where it has access to a notable amount of right dependents not seen by the other transition systems throughout the parsing process.

Table 3

LAS comparison of the original left-to-right parser with sequential decoding against the three available transition systems on Hierarchical Pointer Networks combined with different fusion functions implementations and gating mechanisms on ten treebanks from Universal Dependencies. Only best models on average on dev splits are evaluated on test sets. In [Appendix](#), we report the standard deviations over 3 runs on test splits. We use ISO 639-1 codes to represent languages.

tran.	f	gate	ar	en	eu	fi	he	it	ko	sv	tr	zh	Avg.
dev													
L2R	–	–	83.82	90.29	83.80	88.46	89.11	92.52	83.29	86.78	64.75	82.54	84.54
L2R	l-adapted	GATE1	84.15	90.51	84.67	88.97	89.27	92.63	83.81	87.41	65.35	83.62	85.04
	l-adapted	GATE2	84.05	90.44	84.53	88.95	89.22	92.60	83.79	87.32	65.24	83.48	84.96
	l-simple	GATE1	84.06	90.49	84.84	88.90	89.30	92.59	83.84	87.44	65.40	83.65	85.05
	l-simple	GATE2	84.05	90.29	84.77	88.89	89.20	92.62	83.86	87.43	65.12	83.57	84.98
R2L	r-adapted	GATE1	84.21	90.41	84.83	88.95	89.21	92.47	83.66	87.28	65.25	83.50	84.98
	r-adapted	GATE2	84.20	90.26	84.93	88.81	89.21	92.52	83.68	87.14	65.09	83.53	84.94
	r-simple	GATE1	84.28	90.52	84.75	88.86	89.28	92.56	83.71	87.41	65.33	83.69	85.04
	r-simple	GATE2	84.26	90.46	84.88	88.85	89.20	92.40	83.69	87.26	65.36	83.52	84.99
O-I	full	GATE1	84.07	90.43	84.59	88.74	89.28	92.51	83.46	86.91	65.14	83.54	84.87
	full	GATE2	84.00	90.33	84.42	88.80	89.09	92.38	83.45	86.96	65.07	83.47	84.80
	simple	GATE1	84.13	90.48	84.54	88.77	89.36	92.55	83.56	87.06	65.12	83.49	84.91
	simple	GATE2	84.04	90.35	84.53	88.77	89.10	92.33	83.53	87.07	65.14	83.28	84.81
test													
L2R	–	–	84.20	89.32	84.53	88.45	87.75	92.46	86.01	88.92	66.12	83.68	85.14
L2R	l-simple	GATE1	84.45	89.43	85.24	89.00	88.23	92.50	86.21	89.16	66.63	84.99	85.58
R2L	r-simple	GATE1	84.51	89.43	85.46	88.77	88.26	92.62	86.16	89.17	67.05	84.75	85.62
O-I	simple	GATE1	84.31	89.33	85.33	88.88	88.27	92.54	86.14	88.99	66.60	84.43	85.48

Table 4

Accuracy performance on PTB dev splits. We mark in bold the chosen configurations.

tran.	f	UAS	LAS
L2R	l-adapted	96.03	94.15
	l-simple	96.01	94.15
R2L	r-adapted	96.00	94.14
	r-simple	96.04	94.16
O-I	full	95.96	94.10
	simple	96.00	94.14

5.5. Error analysis

We study the mitigation of error propagation by the proposed neural architecture and characterize errors relative to sentence length and word position¹⁰ for a concatenation of all ten Universal Dependencies datasets¹¹ and for the PTB with BERT-based embeddings.

In particular, [Figs. 4\(a\)](#) and [\(b\)](#) show the accuracy relative to sentence lengths of the three transition systems plus the baseline parser on Universal Dependencies treebanks and PTB, respectively. On Universal Dependencies datasets, we see accuracy gains regardless of sentence length; however, while the improvements tend to be higher as sentences are longer, these narrow when the length is larger than 40 (especially affecting the left-to-right and outside-in variants). On PTB, the novel approaches not only outperform the baseline on long sentences (as expected), but also surprisingly on the shortest ones. We also note that, while the right-to-left parser obtains the highest performance on the longest sentences (better dealing with error-propagation), the outside-in algorithm is suffering a drop on sentences with length greater than 40 on Universal Dependencies treebanks, and on those with length between 31 and 40 on PTB (possibly because the information about left

¹⁰ Errors relative to dependency length (as described in [\[44\]](#)) are no longer an effective method to measure the impact of error propagation in transition systems designed for Pointer Networks (except for the top-down algorithm [\[13\]](#)), since these require a single $\text{SHIFT-ATTACH-}p$ transition to directly build any arc (regardless of its length), not involving several decoder states and, therefore, not directly reflecting the presence of error propagation.

¹¹ In order to prevent data sparsity, we randomly select a number of sentences that approximately gather 10,000 tokens per treebank.

Table 5

Performance comparison of dependency parsers on PTB and CTB. The second block gathers approaches that are enhanced with constituent information, and the last block includes the performance of the top-down transition-based model with Hierarchical Pointer Networks, since only scores with gold PoS tags are reported. * denotes those models that fine-tune BERT and those parsers marked with † report scores based on the best single run on the development set, instead of reporting the average score on the test set over several runs, i.e., instead of averaging to mitigate the effect of random seeds in reported accuracy, they use model selection to choose the most promising seed using the dev set (following this method, our best model *Hier.Ptr.Net L2R* w/o BERT obtains UAS 96.19 LAS 94.61 on PTB).

Parser	PTB		CTB	
	UAS	LAS	UAS	LAS
Zhang et al. [17]	94.10	91.90	87.84	86.15
Ma and Hovy [36]	94.88	92.96	89.05	87.74
Dozat and Manning [18]	95.74	94.08	89.30	88.23
Li et al. [37]	94.11	92.08	88.78	86.23
Ma et al. [13]	95.87	94.19	90.59	89.29
Ji et al. [38] [†]	95.97	94.31	–	–
Fernández-González and Gómez-Rodríguez [14]	96.04	94.43	–	–
Li et al. [39]	95.83	94.54	90.47	89.44
Fernández-González and Gómez-Rodríguez [23]	96.06	94.50	90.61	89.51
Zhang et al. [40] [†]	96.14	94.49	–	–
Wang and Tu [41]	95.98	94.34	90.81	89.57
Hier. Ptr. Net. L2R	96.18	94.59	90.76	89.67
Hier. Ptr. Net. R2L	96.14	94.53	90.72	89.62
Hier. Ptr. Net. O-I	96.07	94.48	90.64	89.50
+BERT				
Li et al. [39]	96.44	94.63	90.89	89.73
Li et al. [39] [*]	96.57	95.05	–	–
Mohammadshahi and Henderson [42] [*]	96.66	95.01	92.86	91.11
Wang and Tu [41] [*]	96.91	95.34	92.55	91.38
Fernández-González and Gómez-Rodríguez [23]	96.91	95.35	92.58	91.42
Hier. Ptr. Net. L2R	97.05	95.47	92.70	91.50
Hier. Ptr. Net. R2L	97.01	95.48	92.75	91.62
Hier. Ptr. Net. O-I	96.95	95.36	92.65	91.47
Zhou and Zhao [43]	96.09	94.68	–	–
Fernández-González and Gómez-Rodríguez [23]	96.25	94.64	90.79	89.69
+BERT				
Zhou and Zhao [43] [*]	97.00	95.43	91.21	89.15
Fernández-González and Gómez-Rodríguez [23]	96.97	95.46	92.78	91.65
Liu et al. [19]	96.09	95.03	–	–

and right dependents is exclusively used by this transition system and it tends to be available at final steps – as shown in [Fig. 2\(c\)](#) –, where

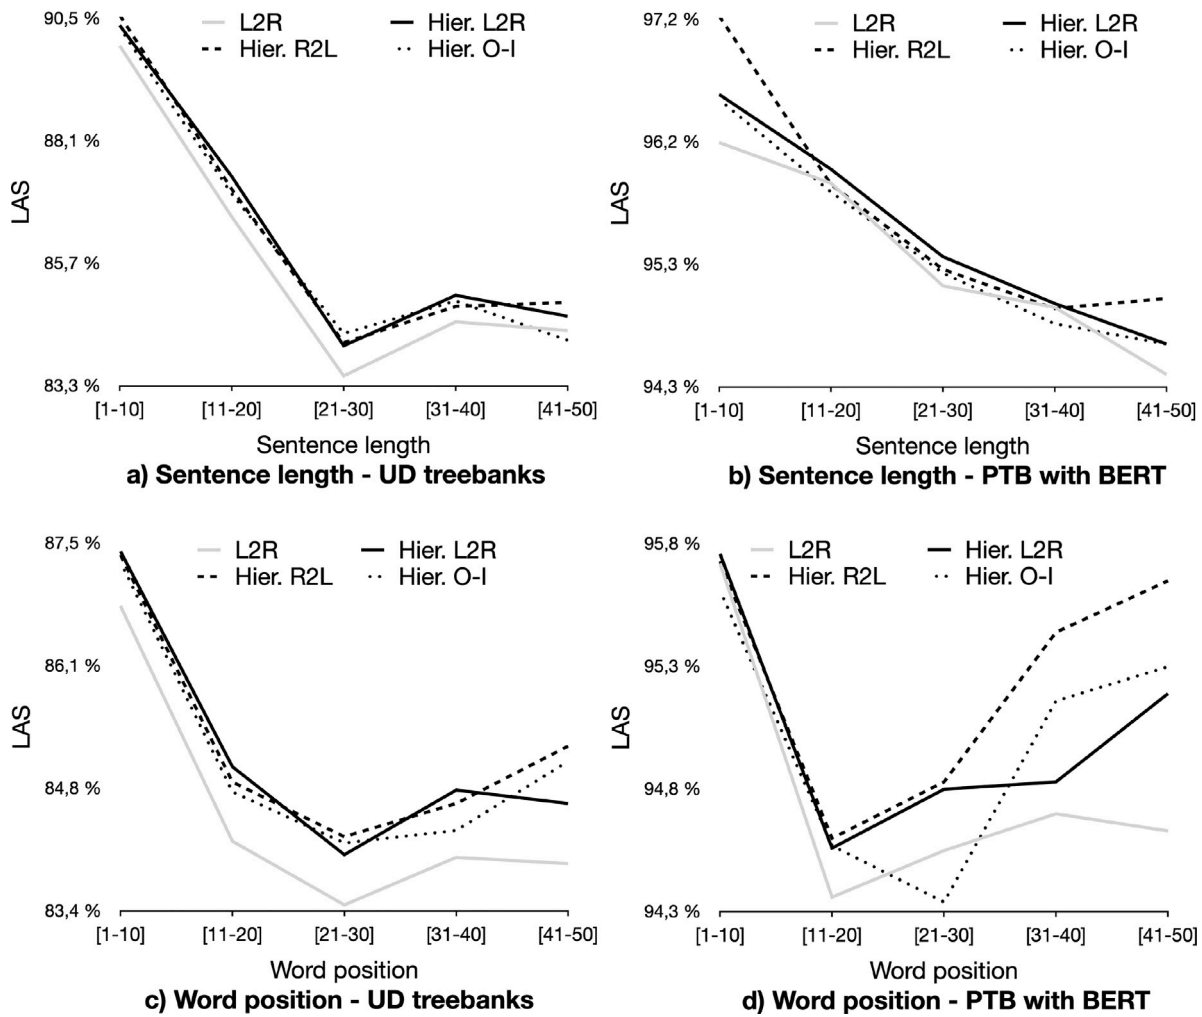


Fig. 4. Parsing performance of the original left-to-right parser and the proposed variants implemented on Hierarchical Pointer Networks relative to sentence length and word position.

Table 6

Number of dependents and long-range dependents per sentence that are available when processing gold trees in Universal Dependencies dev splits with each proposed transition system. Note that the number of long-range dependents (with arc lengths > 4) per sentence is notably low since short sentences are also considered for the computation.

Language	Left2Right		Right2Left		Outside-in	
	all	long	all	long	all	long
Arabic	10.07	1.20	22.20	8.00	14.59	3.75
Basque	6.65	2.29	5.75	2.17	7.10	2.81
Chinese	16.34	6.17	7.99	3.83	12.60	5.34
English	7.02	1.66	4.54	2.54	6.10	2.33
Finnish	7.06	1.69	5.36	2.70	6.61	2.45
Hebrew	10.80	1.69	11.77	5.47	11.69	3.95
Italian	11.15	1.74	8.97	4.88	10.64	3.73
Korean	6.94	2.45	4.64	0.79	5.77	1.61
Swedish	10.63	2.70	7.81	4.18	9.61	3.71
Turkish	5.71	1.85	3.46	0.74	5.07	1.27

it might be difficult to manage when sentences are substantially long and the amount of information is likely to be considerably large).

The reduction of error propagation can be seen more clearly in Figs. 4(c) and (d), where we report the LAS relative to word positions within the sentence. In comparison to the sequential variant, the left-to-right transition system with structured decoding obtains the highest gains in accuracy on attachments made on words at final positions of the sentence (the most affected by error propagation). We can also

Table 7

Standard deviations over 3 runs on test splits for scores reported in Table 5.

Parser	PTB		CTB	
	UAS	LAS	UAS	LAS
Hier. Ptr. Net. L2R	±0.02	±0.03	±0.05	±0.06
Hier. Ptr. Net. R2L	±0.04	±0.03	±0.06	±0.08
Hier. Ptr. Net. O-I	±0.03	±0.05	±0.04	±0.06
+BERT				
Hier. Ptr. Net. L2R	±0.01	±0.02	±0.04	±0.05
Hier. Ptr. Net. R2L	±0.03	±0.02	±0.06	±0.05
Hier. Ptr. Net. O-I	±0.02	±0.02	±0.03	±0.04

observe how the right-to-left and outside-in approaches significantly outperform the left-to-right algorithm on words located near the end of the sentence, since these transition systems attach those words in initial steps of the parsing process. Lastly, there is a significant drop in accuracy by the outside-in model in words at middle positions on PTB (from 21 to 30) and, less significant, on Universal Dependencies datasets (from 31 to 40). This is probably due to the fact that arcs in those positions are created in final steps following the outside-in strategy (thus being the most affected by error-propagation); and, apart from the tight Y axis scale in the plot of Fig. 4(d), the significant drop in PTB can be explained by the fact that the amount of words in absolute positions from 21 to 30 (not necessarily being at the middle of their respective sentences) is substantially higher than in other languages.

Table 8
Standard deviations over 3 runs on test splits for scores reported in Table 3.

tran.	<i>f</i>	gate	ar	en	eu	fi	he	it	ko	sv	tr	zh
L2R	–	–	±0.06	±0.01	±0.05	±0.06	±0.14	±0.04	±0.06	±0.03	±0.13	±0.04
L2R	l-simp.	GATE1	±0.02	±0.04	±0.06	±0.02	±0.08	±0.02	±0.04	±0.04	±0.08	±0.04
R2L	r-simp.	GATE1	±0.06	±0.04	±0.02	±0.06	±0.12	±0.02	±0.06	±0.05	±0.11	±0.08
O-I	simple	GATE1	±0.02	±0.08	±0.05	±0.04	±0.06	±0.07	±0.09	±0.02	±0.12	±0.08

6. Conclusions

We manage to introduce structural knowledge to the sequential decoding of the left-to-right dependency parser with Pointer Networks developed by Fernández-González and Gómez-Rodríguez [14]. The resulting neural architecture, named Bottom-up Hierarchical Pointer Network, is able to leverage relevant information about the focus word's dependents at each decoding step, instead of just receiving the previous decoder hidden state as input. This information is especially useful when it comes from long-range dependents attached in the past, notably helping the decoder to take better decisions in current and future steps. Additionally, we implement two novel transition systems that can be added to the proposed neural architecture: an algorithm that parses a sentence in right-to-left order and a transition system that processes it from the outside in.

We extensively test Bottom-up Hierarchical Pointer Networks on the English and Chinese Penn Treebanks as well as a wide range of languages from different families, with different degrees of morphological complexity and with different predominances of long-range dependencies. In our experiments, we prove that they lead to significant accuracy gains regardless of the fusion function implementation or gating mechanism; and show that the left-to-right approach achieves the highest accuracy obtained so far on the English and Chinese Penn Treebanks (without contextualized word embeddings) and the right-to-left model achieves the best LAS to date when BERT-based embeddings are available. Apart from the improvements in accuracy, the proposed neural architecture does not penalize the quadratic runtime complexity of the original left-to-right parser. Finally, we additionally undertake a thorough error analysis and provide evidence that error propagation (considered the main cause of accuracy loss in transition-based parsing) is mitigated with the presented structured decoding.

CRedit authorship contribution statement

Daniel Fernández-González: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing - review & editing, Visualization. **Carlos Gómez-Rodríguez:** Conceptualization, Validation, Formal analysis, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Source code available at <https://github.com/danifg/BottomUp-Hierarchical-PtrNet>.

Acknowledgments

We acknowledge the European Research Council (ERC), which has funded this research under the European Union's Horizon 2020 research and innovation programme (FASTPARSE, grant agreement No 714150) and the Horizon Europe research and innovation programme (SALSA, grant agreement No 101100615), ERDF/MICINN-AEI (SCANNER-UDC, PID2020-113230RB-C21), Xunta de Galicia (ED431C 2020/11), and Centro de Investigación de Galicia "CITIC", funded by Xunta de Galicia and the European Union (ERDF - Galicia 2014–2020 Program), by grant ED431G 2019/01. Funding for open access charge: Universidade da Coruña/CISUG.

Appendix. Standard deviations

See Tables 7 and 8

References

- [1] M. Zhang, Z. Li, G. Fu, M. Zhang, Dependency-based syntax-aware word representations, *Artif. Intell.* 292 (2021) 103427, <http://dx.doi.org/10.1016/j.artint.2020.103427>, URL <https://www.sciencedirect.com/science/article/pii/S0004370220301764>.
- [2] S. Sun, C. Luo, J. Chen, A review of natural language processing techniques for opinion mining systems, *Inf. Fusion* 36 (2017) 10–25, <http://dx.doi.org/10.1016/j.inffus.2016.10.004>, URL <https://www.sciencedirect.com/science/article/pii/S1566253516301117>.
- [3] B. Zhang, Y. Zhang, R. Wang, Z. Li, M. Zhang, Syntax-aware opinion role labeling with dependency graph convolutional networks, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, 2020a, pp. 3249–3258, <http://dx.doi.org/10.18653/v1/2020.acl-main.297>, Online, URL <https://www.aclweb.org/anthology/2020.acl-main.297>.
- [4] D.Q. Nguyen, K. Verspoor, From pos tagging to dependency parsing for biomedical event extraction, *BMC Bioinform.* 20 (2019) 72, <http://dx.doi.org/10.1186/s12859-019-2604-0>.
- [5] Q. Cao, X. Liang, B. Li, L. Lin, Interpretable visual question answering by reasoning on dependency trees, *IEEE Trans. Pattern Anal. Mach. Intell.* 43 (2021) 887–901, <http://dx.doi.org/10.1109/tpami.2019.2943456>.
- [6] J. Bai, Y. Wang, Y. Chen, Y. Yang, J. Bai, J. Yu, Y. Tong, Syntax-BERT: improving pre-trained transformers with syntax trees, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Association for Computational Linguistics, 2021, pp. 3011–3020, Online, <https://www.aclweb.org/anthology/2021.eacl-main.262>.
- [7] J. Barnes, R. Kurtz, S. Oepen, L. Øvrelid, E. Veldal, Structured sentiment analysis as dependency graph 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), Association for Computational Linguistics, 2021, pp. 3387–3402, <http://dx.doi.org/10.18653/v1/2021.acl-long.263>, Online, <https://aclanthology.org/2021.acl-long.263>.
- [8] E. Cambria, R. Mao, S. Han, Q. Liu, Sentic parser: a graph-based approach to concept extraction for sentiment analysis, in: Proceedings of ICDM Workshops, 2022, URL <https://sentic.net/sentic-parser.pdf>.
- [9] R. Sukthankar, S. Poria, E. Cambria, R. Thirunavukarasu, Anaphora and coreference resolution: a review, *Inf. Fusion* 59 (2020) 139–162, <http://dx.doi.org/10.1016/j.inffus.2020.01.010>, URL <https://www.sciencedirect.com/science/article/pii/S1566253519303677>.
- [10] V. Balachandran, A. Pagnoni, J.Y. Lee, D. Rajagopal, J. Carbonell, Y. Tsvetkov, StructSum: summarization via structured representations, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Association for Computational Linguistics, 2021, pp. 2575–2585, Online, <https://www.aclweb.org/anthology/2021.eacl-main.220>.
- [11] D. Sachan, Y. Zhang, P. Qi, W.L. Hamilton, Do syntax trees help pre-trained transformers extract information?, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Association for Computational Linguistics, 2021, pp. 2647–2661, Online, URL <https://www.aclweb.org/anthology/2021.eacl-main.228>.

- [12] O. Vinyals, M. Fortunato, N. Jaitly, Pointer networks, in: C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 28, Curran Associates, Inc., 2015, pp. 2692–2700, URL <http://papers.nips.cc/paper/5866-pointer-networks.pdf>.
- [13] X. Ma, Z. Hu, J. Liu, N. Peng, G. Neubig, E. Hovy, Stack-pointer networks for dependency parsing, in: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Melbourne, Australia, 2018, pp. 1403–1414, <http://dx.doi.org/10.18653/v1/P18-1130>, URL <https://www.aclweb.org/anthology/P18-1130>.
- [14] D. Fernández-González, C. Gómez-Rodríguez, Left-to-right dependency parsing with pointer networks, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 710–716, <http://dx.doi.org/10.18653/v1/N19-1076>, URL <https://www.aclweb.org/anthology/N19-1076>.
- [15] J. Nivre, An efficient algorithm for projective dependency parsing, in: *Proceedings of the Eighth International Conference on Parsing Technologies*, Nancy, France, 2003, pp. 149–160, URL <https://www.aclweb.org/anthology/W03-3017>.
- [16] R. McDonald, K. Crammer, F. Pereira, Online large-margin training of dependency parsers, in: *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, Association for Computational Linguistics, Ann Arbor, Michigan, 2005, pp. 91–98, <http://dx.doi.org/10.3115/1219840.1219852>, URL <https://www.aclweb.org/anthology/P05-1012>.
- [17] X. Zhang, J. Cheng, M. Lapata, Dependency parsing as head selection, in: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, 2017, pp. 665–676, URL <https://aclanthology.info/papers/E17-1063/e17-1063>.
- [18] T. Dozat, C. Manning, Deep biaffine attention for neural dependency parsing, in: *ICLR, 2017, OpenReview.net*.
- [19] L. Liu, X. Lin, S. Joty, S. Han, L. Bing, Hierarchical pointer net parsing, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 1007–1017, <http://dx.doi.org/10.18653/v1/D19-1093>, URL <https://www.aclweb.org/anthology/D19-1093>.
- [20] J. Nivre, M.-C. de Marneffe, F. Ginter, Y. Goldberg, J. Hajič, C.D. Manning, R. McDonald, S. Petrov, S. Pyysalo, N. Silveira, R. Tsarfaty, D. Zeman, Universal dependencies v1: A multilingual treebank collection, in: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, European Language Resources Association (ELRA), Portorož, Slovenia, 2016, pp. 1659–1666, URL <https://www.aclweb.org/anthology/L16-1262>.
- [21] M.P. Marcus, B. Santorini, M.A. Marcinkiewicz, Building a large annotated corpus of English: The Penn Treebank, *Comput. Linguist.* 19 (1993) 313–330.
- [22] N. Xue, F. Xia, F.d. Chiou, M. Palmer, The penn chinese treebank: phrase structure annotation of a large corpus, *Nat. Lang. Eng.* 11 (2005) 207–238, <http://dx.doi.org/10.1017/S135132490400364X>.
- [23] D. Fernández-González, C. Gómez-Rodríguez, Multitask pointer network for multi-representational parsing, *Knowl.-Based Syst.* 236 (2022) 107760, <http://dx.doi.org/10.1016/j.knsys.2021.107760>, URL <https://www.sciencedirect.com/science/article/pii/S0950705121009849>.
- [24] S. Yang, T. Gao, J. Wang, B. Deng, B. Lansdell, B. Linares-Barranco, Efficient spike-driven learning with dendritic event-based processing, *Front. Neurosci.* 15 (2021) <http://dx.doi.org/10.3389/fnins.2021.601109>, URL <https://www.frontiersin.org/articles/10.3389/fnins.2021.601109>.
- [25] S. Yang, T. Gao, J. Wang, B. Deng, M.R. Azghadi, T. Lei, B. Linares-Barranco, Sam: a unified self-adaptive multicompartmental spiking neuron model for learning with working memory, *Front. Neurosci.* 16 (2022) <http://dx.doi.org/10.3389/fnins.2022.850945>, URL <https://www.frontiersin.org/articles/10.3389/fnins.2022.850945>.
- [26] D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, 2014, *CoRR* abs/1409.0473.
- [27] X. Ma, E. Hovy, End-to-end sequence labeling via bi-directional lstm-cnns-crf, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics., 2016, pp. 1064–1074, <http://dx.doi.org/10.18653/v1/P16-1101>, URL <http://aclweb.org/anthology/P16-1101>.
- [28] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186, <http://dx.doi.org/10.18653/v1/N19-1423>, URL <https://www.aclweb.org/anthology/N19-1423>.
- [29] M. Haspelmath, M. Dryer, O. University, H. Bibiko, O.U. Press, M.P.I. für Evolutionäre Anthropologie, J. Hagen, D. Gil, B. Comrie, C. Schmidt, *The World Atlas of Language Structures*, Oxford linguistics, OUP Oxford, 2005, URL <https://books.google.es/books?id=7rkSDAAAQBAJ>.
- [30] A. Kulmizev, M. de Lhoneux, J. Gontrum, E. Fano, J. Nivre, Deep contextualized word embeddings in transition-based and graph-based dependency parsing - a tale of two parsers revisited, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 2755–2768, <http://dx.doi.org/10.18653/v1/D19-1277>, URL <https://www.aclweb.org/anthology/D19-1277>.
- [31] M.-C. de Marneffe, C.D. Manning, The stanford typed dependencies representation, in: *Coling 2008: Proceedings of the Workshop on Cross-Framework and Cross-Domain Parser Evaluation*, Coling 2008 Organizing Committee, Manchester, UK, 2008, pp. 1–8, URL <https://www.aclweb.org/anthology/W08-1301>.
- [32] Y. Zhang, S. Clark, A tale of two parsers: Investigating and combining graph-based and transition-based dependency parsing, in: *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Honolulu, Hawaii, 2008, pp. 562–571, URL <https://www.aclweb.org/anthology/D08-1059>.
- [33] W. Ling, C. Dyer, A.W. Black, I. Trancoso, Two/too simple adaptations of Word2Vec for syntax problems, in: *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics*, Association for Computational Linguistics, Denver, Colorado, 2015, pp. 1299–1304, <http://dx.doi.org/10.3115/v1/N15-1142>, URL <https://www.aclweb.org/anthology/N15-1142>.
- [34] R. Al-Rfou, B. Perozzi, S. Skiena, Polyglot: Distributed word representations for multilingual NLP, in: *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, Association for Computational Linguistics, Sofia, Bulgaria, 2013, pp. 183–192, URL <https://www.aclweb.org/anthology/W13-3520>.
- [35] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization, 2014, URL <http://arxiv.org/abs/1412.6980>. published as a conference paper at the 3rd International Conference for Learning Representations, San Diego.
- [36] X. Ma, E. Hovy, Neural probabilistic model for non-projective mst parsing, in: *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Asian Federation of Natural Language Processing, 2017, pp. 59–69, URL <http://aclweb.org/anthology/I17-1007>.
- [37] Z. Li, J. Cai, S. He, H. Zhao, Seq2seq dependency parsing, in: *Proceedings of the 27th International Conference on Computational Linguistics*, Association for Computational Linguistics, Santa Fe, New Mexico, USA, 2018, pp. 3203–3214, URL <https://www.aclweb.org/anthology/C18-1271>.
- [38] T. Ji, Y. Wu, M. Lan, Graph-based dependency parsing with graph neural networks, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 2475–2485, <http://dx.doi.org/10.18653/v1/P19-1237>, URL <https://www.aclweb.org/anthology/P19-1237>.
- [39] Z. Li, H. Zhao, K. Parnow, Global greedy dependency parsing, in: *the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-2020)*, 2020.
- [40] Y. Zhang, Z. Li, M. Zhang, Efficient second-order TreeCRF for neural dependency parsing, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, 2020b, pp. 3295–3305, <http://dx.doi.org/10.18653/v1/2020.acl-main.302>, Online, URL <https://www.aclweb.org/anthology/2020.acl-main.302>.
- [41] X. Wang, K. Tu, Second-order neural dependency parsing with message passing and end-to-end training, 2020, *arXiv:2010.05003*.
- [42] A. Mohammadshahi, J. Henderson, Recursive non-autoregressive graph-to-graph transformer for dependency parsing with iterative refinement, 2020, *arXiv:2003.13118*.
- [43] J. Zhou, H. Zhao, Head-driven phrase structure grammar parsing on penn treebank, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 2396–2408, <http://dx.doi.org/10.18653/v1/P19-1230>, URL <https://www.aclweb.org/anthology/P19-1230>.
- [44] R. McDonald, J. Nivre, Characterizing the errors of data-driven dependency parsing models, in: *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CONLL)*, Association for Computational Linguistics, Prague, Czech Republic, 2007, pp. 122–131, URL <https://www.aclweb.org/anthology/D07-1013>.