

An Argument-Marker Model for Syntax-Agnostic Proto-Role Labeling

Juri Opitz and Anette Frank

Research Training Group AIPHES,
Leibniz ScienceCampus “Empirical Linguistics and Computational Language Modeling”
Department for Computational Linguistics
69120 Heidelberg
{opitz, frank}@cl.uni-heidelberg.de

Abstract

Semantic proto-role labeling (SPRL) is an alternative to semantic role labeling (SRL) that moves beyond a categorical definition of roles, following Dowty’s feature-based view of *proto-roles*. This theory determines *agenthood* vs. *patienthood* based on a participant’s instantiation of more or less typical agent vs. patient properties, such as, for example, *volition* in an event. To perform SPRL, we develop an ensemble of hierarchical models with self-attention and concurrently learned predicate-argument-markers. Our method is competitive with the state-of-the-art, overall outperforming previous work in two formulations of the task (multi-label and multi-variate Likert scale prediction). In contrast to previous work, our results do not depend on gold argument heads derived from supplementary gold tree banks.

1 Introduction

Deciding on a linguistically sound, clearly defined and broadly applicable inventory of semantic roles is a long-standing issue in linguistic theory and natural language processing. To alleviate issues found with classical thematic role inventories, Dowty (1991) argued for replacing categorical roles with a feature-based, composite notion of semantic roles, introducing the *theory of semantic proto-roles (SPR)*. At its core, it proposes two prominent, composite role types: *proto-agent* and *proto-patient*. Proto-roles represent multi-faceted, possibly graded notions of *agenthood* or *patienthood*. For example, consider the following sentence from Bram Stoker’s *Dracula* (1897):

- (1) *He opened it [the letter] and read it gravely.*

‘Davidsonian’ analyses based on SR and SPR of the event *open* are displayed in Figure 1. The SPR analysis provides more detail about the event and the roles of the involved entities. Whether an

<p>SR: $\exists e[open(e) \wedge agent(e, c^*) \wedge theme(e, l^*)]$ SPR: $\exists e[open(e) \wedge volition(e, c^*) \wedge aware(e, c^*) \wedge sentient(e, c^*) \wedge manipulated(e, l^*) \wedge changes-state(e, l^*) \wedge \dots]$</p>

Figure 1: Two different ‘Davidsonian’ event analyses of *open*. c^* and l^* refer to the *count* and *letter* entities.

argument is considered an *agent* or *patient* follows from the proto-typical properties the argument exhibits: e.g., *being manipulated* is proto-typical for patient, while *volition* is proto-typical for an agent. Hence, in both events of (1) the count is determined as agent, and the letter as patient.

Only recently two SPR data sets have been published. Reisinger et al. (2015) developed a property-based proto-role annotation schema with 18 properties. One Amazon Mechanical Turk crowd worker (selected in a pilot annotation) answered questions such as *how likely is it that the argument mentioned with the verb changes location?* on a 5-point Likert or responded *inapplicable*. This dataset (news domain) will henceforth be denoted by **SPR1**. Based on the experiences from the SPR1 annotation process, White et al. (2016) published **SPR2** which follows a similar annotation schema. However, in contrast to SPR1, the new data set contains doubly annotated data from the web domain for 14 refined properties.

Our work makes the following contributions: In Section §2, we provide an overview of previous SPRL work and outline a common weakness: reliance on gold syntax trees or gold argument heads derived from them. To alleviate this issue, we propose a span-based, hierarchical neural model (§3) which learns marker embeddings to highlight the predicate-argument structures of events. Our experiments (§4) show that our model, when combined in a simple voter ensemble, outperforms all previous works. A single model performs only

slightly worse, albeit having weaker dependencies than previous methods. In our analysis, we (i) perform ablation experiments to analyze the contributions of different model components. (ii) we observe that the small SPR data size introduces a severe sensitivity to different random initializations of our neural model. We find that combining multiple models in a simple voter ensemble makes SPRL predictions not only slightly better but also significantly more robust. We share our code with the community and make it publicly available.¹

2 Related Work

SPRL Teichert et al. (2017) formulate the semantic role labeling task as a multi-label problem and develop a conditional random field model (CRF). Given an argument phrase and a corresponding predicate, the model predicts which of the 18 properties hold. Compared with a simple feature-based linear model by Reisinger et al. (2015), the CRF exhibits superior performance by more than 10 pp. macro F1. Incorporating features derived from additional gold syntax improves the CRF performance significantly. For treating the task as a multi-label problem, the Likert classes $\{1, 2, 3\}$ and *inapplicable* are collapsed to $-$ and Likert classes $\{4, 5\}$ are mapped to $+$. Subsequent works, including ours, conform to this setup.

Rudinger et al. (2018) are the first to treat SPRL as a multi-variate Likert scale regression problem. They develop a neural model whose predictions have good correlation with the values in the testing data of both SPR1 and SPR2. In the multi-label setting, their model compares favourably with Teichert et al. (2017) for most proto-role properties and establishes a new state-of-the-art. Pre-training the model in a machine translation setting helps on SPR1 but results in a performance drop on SPR2. The model takes a sentence as input to a Bi-LSTM (Hochreiter and Schmidhuber, 1997) to produce a sequence of hidden states. The prediction is based on the hidden state corresponding to the head of the argument phrase, which is determined by inspection of the gold syntax tree.

Recently, Tenney et al. (2019) have demonstrated the capacities of contextualized word embeddings across a wide variety of tasks, including SPRL. However, for SPRL labeling they proceed similar to Rudinger et al. (2018) in the sense

that they extract the gold heads of arguments in their dependency-based SPRL approach. Instead of using an LSTM to convert the input sentence to a sequence of vectors they make use of large language models such as ELMo (Peters et al., 2018) or BERT (Devlin et al., 2018). The contextual vectors corresponding to predicate and the (gold) argument head are processed by a projection layer, self-attention pooling (Lee et al., 2017) and a two-layer feed forward neural network with sigmoid output activation functions. To compare with Rudinger et al. (2018), our basic model uses standard GloVe embeddings. When our model is fed with contextual embeddings a further observable performance gain can be achieved.

To summarize, previous state-of-the-art SPRL systems suffer from a common problem: they are dependency-based and their results rely on gold argument heads. Our approach, in contrast, does not rely on any supplementary information from gold syntax trees. In fact, our marker model for SPRL is agnostic to any syntactic theory and acts solely on the basis of argument spans which we highlight with position markers.

SRL The task of automatically identifying predicate-argument structures and assigning roles to arguments was firstly investigated by Gildea and Jurafsky (2002). Over the past years, SRL has witnessed a large surge in interest. Recently, very competitive end-to-end neural systems have been proposed (He et al., 2018a; Cai et al., 2018; He et al., 2018b). Strubell et al. (2018) show that injection of syntax can help SRL models and Li et al. (2019) bridge the gap between span-based and dependency-based SRL, achieving new state-of-the-art results both on the span based CoNLL data (Carreras and Màrquez, 2005; Pradhan et al., 2013) and the dependency-based CoNLL data (Surdeanu et al., 2008; Hajič et al., 2009). A fully end-to-end S(P)RL system has to solve multiple sub-tasks: identification of predicate-argument structures, sense disambiguation of predicates and as the main and final step, labeling their argument phrases with roles. Up to the present, SPRL works (including ours) focus on the main task and assume the prior steps as complete.

Research into SPRL is still in its infancy, especially in comparison to SRL. One among many reasons may be the fact that, in contrast to semantic roles, semantic proto-roles are multi-dimensional. This introduces more complexity:

¹<https://gitlab.cl.uni-heidelberg.de/opitz/sprl>

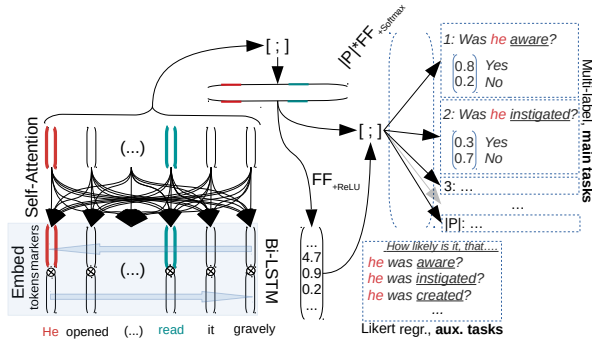


Figure 2: Model outline. *Input*: (i) a sequence of vectors representing the words and (ii) a sequence of vectors which serve to highlight predicate and argument. *Processing*: 1. element-wise multiplication of the two sequences (\otimes); 2. generation of hidden states with forward and backward Bi-LSTM reads (\Rightarrow & \Leftarrow); 3. self-attention mechanism builds a new sequence of hidden states by letting every hidden state attend to every other hidden state; 4. concatenation of the hidden states to generate a vector representation ($[:]$). *Output*: (i) use vector representation to output Likert scale auxiliary predictions (FF_{ReLU}) and (ii) concatenate auxiliary predictions to the vector representation ($[:]$) to finally (iii) compute the multi-label predictions at the top level ($|P| \cdot FF_{Softmax}$; P : set of proto role properties).

given a predicate and an argument, the task is no more to predict *a single* label (as in SRL), but a *list* of multiple labels or even a multi-variate Likert scale. Another reason may be related to the available resources. The published SPR data sets comprise significantly fewer examples. The design of annotation guidelines and pilot studies with the aim of in-depth proto-role annotations is a hard task. In addition, the SPR data were created, at least to a certain extent, in an experimental manner: one of the goals of corpus creation was to explore possible SPR annotation protocols for humans. We hope that a side-effect of this paper is to spark more interest in SPR and SPRL.

3 Attentive Marker Model

Since the work of Teichert et al. (2017), the SPRL problem has been phrased as follows: given a (*sentence, predicate, argument*) tuple, we need to predict for all possible argument properties from a given property-inventory whether they hold or not (regression: *how likely are they to hold?*).

Following previous work (Rudinger et al., 2018), the backbone of our model is a Bi-LSTM. To ensure further comparability, pretrained 300 dimensional GloVe embeddings (Pennington et al., 2014) are used for building the input sequence

(e_1, \dots, e_T). In contrast to Rudinger et al. (2018), we multiply a sequence of *marker embeddings* (m_1, \dots, m_T) element-wise with the sequence of word vectors: ($e_1 \cdot m_1, \dots, e_T \cdot m_T$) (\otimes , Figure 2). We distinguish three different marker embeddings that indicate the position of the argument in question (*red*, Figure 2), the predicate (*green*, Figure 2) and remaining parts of the sentence. This is to some extent similar to He et al. (2017) who learn two predicate indicator embeddings which are concatenated to the input vectors and serve the purpose of showing the model whether a token is the predicate or not. However, in SPRL we are also provided with the argument phrase. We will see in the ablation experiments that it is paramount to learn a dedicated embedding. Embedding multiplication instead of concatenation has the advantage of fewer LSTM parameters (smaller input dimension). Besides, it provides the model with the option to learn large coefficients of the word vector dimensions of predicate and argument vectors. This should immediately draw the model’s attentiveness to the argument and predicate phrases which now are accentuated.

The sequence of marked embeddings is further processed by a Bi-LSTM in order to obtain a sequence of hidden states $S = (s_1, \dots, s_T)$. In Figure 2, forward and backward LSTM reads are indicated by \Rightarrow and \Leftarrow .

From there, we take intuitions from Zheng et al. (2018) and compute the next sequence of vectors by letting every hidden state attend to every other hidden state, which is expressed by the following formulas:

$$\begin{aligned} h_{t,t'} &= \tanh(QS_t + KS_{t'} + \beta) \\ e_{t,t'} &= \sigma(v^T h_{t,t'} + \alpha) \\ a_t &= \text{softmax}(e_t) \\ z_t &= \sum_{t'} a_{t,t'} \cdot s_{t'} \end{aligned}$$

Q, K are weight matrices, β is a bias vector, α is a bias scalar and v a weight vector. Letting every hidden state attend to every other hidden state gives the model freedom in computing the argument-predicate composition. This is desirable, since arguments and predicates frequently are in long-distance relationships. For example, in Figure 3 we see that in the SPR1 data predicates and arguments often lie more than 10 words apart and a non-negligible amount of cases consists of distances of more than 20 words.

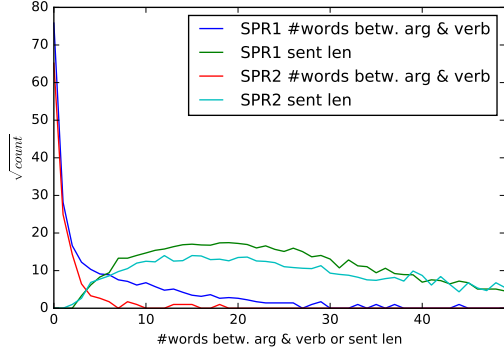


Figure 3: Distribution of the number of words between argument and verb (distance relationship) and sentence lengths in the data sets SPR1 and SPR2.

We proceed by concatenation, $\mathbf{z} = [z_1; \dots; z_T]$, and compute intermediate outputs approximating the property-likelihood Likert scales. This is achieved with weight matrix A and $ReLU$ activation functions (FF_{+ReLU} , Figure 2):

$$a = ReLU(A\mathbf{z}). \quad (1)$$

To perform multi-label prediction with $|P|$ possible labels we use $[a; \mathbf{z}]$ for computing the final decisions with $2|P|$ output neurons and $|P|$ separate weight matrices ($|P| * FF_{+softmax}$, Figure 2), one for each property $p \in P$:

$$o_p = softmax(W^p[a; \mathbf{z}]). \quad (2)$$

For example, given the 18 proto-role properties contained in SPR1, we learn $|P| = 18$ weight matrices and use the Softmax functions to produce 18 vectors of dimension 2 as outputs. The first dimension $o_{p,0}$ represents the predicted probability that property p does *not apply* ($o_{p,1}$: probability for p *applies*). For the regression task, we reduce the number of output neurons from $2|P|$ to $|P|$ and use $ReLU$ activation functions instead. We hypothesize that the hierarchical structure can support the model in making predictions on the top layer. E.g., if the argument is predicted to be most likely *not sentient* and very likely to be *manipulated*, the model may be less tempted to predict an *awareness* label at the top layer. The auxiliary loss for any datum is given as the mean square error over the auxiliary output neurons:

$$\ell' = \frac{\lambda'}{|P|} \sum_{p \in P} (a_p^* - a_p)^2 \quad (3)$$

In case of the multi-label formulation, our main loss for an example is the average cross entropy loss over every property:

$$\ell = -\frac{\lambda}{|P|} \sum_{p \in P} (o_{p,1}^* \log o_{p,1} + o_{p,2}^* \log o_{p,2}), \quad (4)$$

where $o_{p,0}^* = \mathbf{I}(\neg \mathbf{p})$ and $o_{p,1}^* = \mathbf{I}(\mathbf{p})$ i.e. the gold label indicator.

4 Experiments

Data We use the same data setup and split as Teichert et al. (2017); Rudinger et al. (2018); Tenney et al. (2019). For determining the gold labels, we also conform to prior works and (i) collapse classes in the multi-label setup from $\{NA, 1, 2, 3\}$ and $\{4, 5\}$ to classes ‘-’ and ‘+’ and (ii) treat NA as 1 in the Likert regression formulation. For doubly annotated data (SPR2), the Likert scores are averaged; in the multi-label setup we consider values ≥ 4 as ‘+’ and map lesser scores to ‘-’. More data and pre-processing details are described in the Supplement §A.1.

Baselines As baselines we present the results from previous systems: the state-of-the-art by Rudinger et al. (2018) is denoted in our tables as **RUD’18**, the linear feature-based classifier by Reisinger et al. (2015) as **REI’15** and the CRF developed by Teichert et al. (2017) as **TEI’17**. Like previous works, we use macro F1 as the global performance metric in the multi-label scenario and macro-averaged Pearson’s ρ (arithmetic mean over the correlation coefficients for each property, details can be found in the Supplement A.1) We refer to the system results as reported by Rudinger et al. (2018). The most recent work, which evaluates large language models on a variety of tasks including SPRL, is denoted by **TEN’19** (Tenney et al., 2019). In this case, we present the micro F1 results as reported in their paper.

Model instantiation We introduce four main models: (i) **Marker**: our basic, span-based single-model system. For (ii) **MarkerE**, we fit an ensemble of 50 **Markers** with different random seeds. Computationally, training 50 neural models in this task is completely feasible since neither SPR1 nor SPR2 contain more than 10,000 training examples (parallelized training took approximately 2 hours). The ensemble predicts unseen testing data by combining the models’ decisions in a simple majority vote when performing multi-label prediction

or, when in the regression setup, by computing the mean of the output scores (for every property). We also introduce (iii) **MarkerB**, and (iv) **MarkerEB**. These two systems differ in only one aspect from the previously mentioned models: instead of GloVe word vectors, we feed contextual vectors extracted from the BERT model (Devlin et al., 2018). More precisely, we use the transformer model *BERT-base-uncased*² and sum the inferred activations over the last four layers. The resulting vectors are concatenated to GloVe vectors and then processed by the Bi-LSTM.

We fit all models with gradient descent and apply early stopping on the development data (maximum average Pearson’s ρ for multi-variate Likert regression, maximum macro F1 for the multi-label task). Further hyper parameter choices and details about the training are listed in Appendix §A.2.

4.1 Multi-Label Prediction Results

News texts The results on newspaper data (SPR1) are displayed in Table 1 (left-hand side). Our basic ensemble (**MarkerE**) improves massively in the property *location*³ (+19.1 pp. F1). A significant loss is experienced in the property *changes possession* (-9.6). Overall, our ensemble method outperforms all prior works (REI’15: +17.7 pp. macro F1; TEI’17: +6.2, RUD’18: +1.0). Our ensemble method provided with additional contextual word embeddings (**MarkerEB**) yields another large performance gain. The old state-of-the-art, RUD’18, is surpassed by more than 6.0 pp. macro F1 (a relative improvement of 8.6%). With regard to some properties, the contextual embeddings provide massive performance gains over our basic **MarkerE**: *stationary* (+9.5 pp. F1), *makes physical contact* (+21.3), *change of location* (+14.1) and *created* (+17.3). The only loss is incurred for the property which asks if an argument is *destroyed* (-12.3). This specific property appears to be difficult to predict for all models. The best score in this property is achieved by **MarkerE** with only 26.6 F1.

Web texts On the web texts (SPR2), due to less previous works, we also use three label selection strategies as baselines: a *majority* label baseline, a *constant* strategy which always selects the positive label and a *random* baseline which samples a pos-

itive target label according to the occurrence ratio in the training data (*maj, constant & ran*, Table 2, left-hand side).

Our basic **MarkerE** method yields massive improvements over both baselines (more than +10 pp. F1) in 4 out of 14 proto-role properties. For *argument changes possession* and *awareness* the improvement over both baselines is more than +25 pp. F1 and for *sentient* more than +40 pp. However, in the *partitive* property, the constant-label baseline remains unbeaten by a large margin (-21.7 pp.). Overall, all Marker models yield large performance increases over the baselines. For example, **MarkerE** yields significant improvements both over *random* (+27.7 pp. macro F1), *constant* (+9.5) and *majority* (+45.0).

Intriguingly – while the contextual embeddings provide a massive performance boost on news texts (SPR1), – they appear not to be useful for our model on the web texts. The macro F1 score of **MarkerEB** is slightly worse (-1 pp.) than that of **MarkerE** and the micro F1 score is only marginally better (+0.9). The same holds true for the single-model instances: **Marker** performs better than **MarkerB** by 2.2 pp. macro F1 albeit marginally worse micro F1 wise by 0.5 pp.

Why exactly the contextual embeddings fail to provide any valuable information when labeling arguments in web texts, we cannot answer with certainty. A plausible cause could be overfitting problems stemming from the increased dimensionality of the input vectors. In fact, the contextual embeddings increase the number of word vector features by more than two times over the dimension of the GloVe embeddings. This inflates the number of parameters in the LSTM’s weight matrices. As a consequence, the likelihood of overfitting is increased – an issue which is further aggravated by the fact that SPR2 data are significantly fewer than SPR1 data. SPR2 contains less than five thousand predicate-argument training examples, roughly half the size of SPR1.

Another source of problems may be rooted in the target-label construction process for SPR2. This question does not arise when using SPR1 data since all annotations were performed by a single annotator. The SPR2 data, in contrast, contains for each predicate-argument pair, two annotations. In total, the data was annotated by many crowd workers – some of whom provided many and some provided few annotations. Perhaps, averaging Lik-

²<https://github.com/google-research/bert>

³i.e. *does the argument describe the location of the event?*

property	multi-label (ML), F1 score								regression (LR), ρ				
	previous works			ours					RUD'18	ours			
	REF'15	TEI'17	RUD'18	MarkerE	Marker	MarkerEB	MarkerB	MarkerE		Marker	MarkerEB	MarkerB	
awareness	68.8	87.3	89.9	89.4	88.1	93.4	91.6	0.897	0.880	0.868	0.922	0.915	
chg location	6.6	35.6	45.7	50.5	53.2	64.6	57.1	0.702	0.768	0.744	0.835	0.814	
chg state	54.6	66.1	71.0	66.6	68.2	71.5	66.1	0.604	0.651	0.621	0.701	0.653	
chg possession	0.0	38.8	58.0	48.4	55.5	55.7	59.5	0.640	0.609	0.576	0.652	0.582	
created	0.0	44.4	39.7	34.8	46.5	52.2	56.1	0.549	0.593	0.498	0.669	0.614	
destroyed	17.1	0.0	24.2	26.6	19.9	14.3	22.2	0.346	0.368	0.220	0.412	0.347	
existed after	82.3	87.5	85.9	87.0	82.9	88.5	84.8	0.619	0.612	0.571	0.695	0.669	
existed before	79.5	84.8	85.1	87.4	85.4	90.1	85.8	0.710	0.704	0.668	0.781	0.741	
existed during	93.1	95.1	95.0	95.2	94.3	95.8	94.0	0.673	0.675	0.626	0.732	0.714	
exists as physical	64.8	76.4	82.7	83.7	80.6	88.1	85.8	0.834	0.807	0.777	0.871	0.856	
instigated	76.7	85.6	88.6	86.6	86.4	88.9	87.6	0.858	0.856	0.842	0.879	0.860	
location	0.0	18.5	53.8	72.9	69.8	78.9	76.1	0.619	0.755	0.742	0.849	0.820	
makes physical contact	21.5	40.7	47.2	45.7	32.3	67.0	61.1	0.741	0.716	0.671	0.801	0.772	
manipulated	72.1	86.0	86.8	86.9	86.7	89.6	86.5	0.737	0.738	0.705	0.774	0.751	
pred changed arg	54.0	67.8	70.7	68.1	67.6	72.8	66.1	0.592	0.621	0.579	0.714	0.664	
sentient	42.0	85.6	90.6	89.5	88.3	95.4	92.2	0.925	0.904	0.887	0.959	0.951	
stationary	13.3	21.4	47.4	41.0	25.0	50.5	54.5	0.711	0.666	0.654	0.771	0.739	
volitional	69.8	86.4	88.1	88.3	86.7	91.6	90.1	0.882	0.873	0.863	0.911	0.903	
micro	71.0	81.7	83.3	83.6	82.0	86.8	83.5	-	-	-	-	-	
macro	55.4	65.9	71.1	72.1	69.3	77.5	73.8	0.706	0.711	0.673	0.774	0.743	

Table 1: SPR1 results. **bold**: better than all previous work; **bold**: overall best.

property	multi-label (ML), F1 score								regression (LR), ρ				
	baselines			ours					RUD'18	ours			
	maj	ran	const	MarkerE	Marker	MarkerEB	MarkerB	MarkerE		Marker	MarkerEB	MarkerB	
awareness	0.0	48.9	67.1	92.7	92.3	94.0	91.1	0.879	0.882	0.868	0.902	0.878	
chg location	0.0	12.0	21.7	28.6	35.1	38.0	18.2	0.492	0.517	0.476	0.563	0.507	
chg possession	0.0	5.5	6.6	33.3	33.3	35.6	41.1	0.488	0.520	0.483	0.549	0.509	
chg state	0.0	19.5	31.3	29.7	27.1	41.4	45.2	0.352	0.351	0.275	0.444	0.369	
chg state continuous	0.0	9.2	21.7	25.3	19.8	26.8	30.4	0.352	0.396	0.321	0.483	0.423	
existed after	94.1	86.1	94.1	94.0	92.4	94.0	94.5	0.478	0.469	0.403	0.507	0.476	
existed before	89.5	80.0	89.5	91.0	90.5	92.0	89.8	0.616	0.645	0.605	0.690	0.664	
existed during	98.0	96.2	97.0	98.0	97.8	98.1	98.1	0.358	0.374	0.280	0.354	0.301	
instigated	0.0	48.9	70.5	77.9	78.0	78.9	78.7	0.590	0.582	0.540	0.603	0.599	
partitive	0.0	10.4	24.2	2.5	16.5	9.2	2.4	0.359	0.283	0.213	0.374	0.330	
sentient	0.0	47.6	44.3	91.9	91.6	93.7	92.0	0.880	0.874	0.859	0.892	0.872	
volitional	0.0	39.1	61.8	88.1	87.2	89.7	88.5	0.841	0.839	0.825	0.870	0.854	
was for benefit	0.0	31.6	48.8	61.1	59.2	60.2	63.4	0.578	0.580	0.525	0.598	0.569	
was used	79.3	66.1	79.3	77.9	78.0	77.6	79.9	0.203	0.173	0.093	0.288	0.264	
micro	65.0	62.9	61.4	84.0	83.4	84.9	83.9	-	-	-	-	-	
macro	25.9	43.2	61.4	70.9	69.7	69.9	67.5	0.534	0.535	0.483	0.580	0.544	

Table 2: SPR2 results. **bold**: better than previous work and/or baselines; **bold**: overall best.

ert scale annotations of two random annotators is not the right way to transform SPR2 to a multi-label task. Future work may investigate new transformation strategies. For example, we can envision a strategy which finds reliable annotators and weighs the choices of those annotators higher than those of less reliable annotators. This should result in an improved SPR2 gold standard for both multi-label and multi-variate Likert scale SPRL systems.

4.2 Likert Scale Regression Results

News texts **MarkerE** achieves large performance gains for the properties *location* and *change of location* ($\Delta\rho$: +0.136 & $\Delta\rho$: +0.066, Table 1). This is in accordance with the results for these two properties in the multi-label predic-

tion setup. Our model is outperformed by RUD'18 in the property *stationary* ($\Delta\rho$: -0.045). All in all, **MarkerE** outperforms RUD'18 (Δ macro ρ : +0.005). When providing additional contextual word embeddings from the large language model, the correlations intensify for almost all role properties. Overall, the contextual embeddings in **MarkerEB** yield an observable improvement of +0.063 Δ macro ρ over **MarkerE** (which solely uses GloVe embeddings).

Web texts Our **MarkerE** model outperforms RUD'18 slightly by +0.001 $\Delta\rho$ (Table 2). However, if we compare with RUD'18's model setup which achieved the best score on the SPR1 testing data (pre-training with a supervised MT task, macro regression result SPR2: 0.521 ρ), we

Method	span or dep.	used by system				macro result				micro result		
		gold head	STL	ensembling	C-embeddings	SPR1		SPR2		SPR1 ML	SPR2 ML	
						ML	LR	ML	LR			
previous	REI'15	span	no	no	no	no	55.4	-	-	-	71.0	-
	TEI'17	dependency	yes (full parse)	no	no	no	65.9	-	-	-	81.7	-
	RUD'18	dependency	yes	no	no	no	69.3	0.697	-	0.534	82.2	-
	RUD'18 _(+MT pretrain)	dependency	yes	yes	no	no	71.1	0.706	-	0.521	83.3	-
	TEN'19	dependency	yes	no	no	yes (BERT concat)	-	-	-	-	84.7	83.0
		dependency	yes	no	no	yes (BERT lin. comb.)	-	-	-	-	86.1 [†]	83.8
ours	Marker	span	no	no	no	no	69.3	0.682	67.9 [†]	0.483	80.8	81.1
	MarkerE	span	no	no	yes	no	<u>72.1</u>	<u>0.711</u>	70.9	<u>0.535</u>	<u>83.6</u>	<u>84.0</u>
	MarkerB	span	no	no	no	yes (BERT sum)	73.5 [†]	0.741 [†]	67.5	0.544 [†]	84.9	83.9 [†]
	MarkerEB	span	no	no	yes	yes (BERT sum)	77.5	0.774	69.9	0.580	86.8	84.9

Table 3: Main results, system properties and requirements of SPRL systems. Overall best system is marked in **bold**, best system using GloVe is underlined, best single-model system is marked by †. STL: supervised transfer-learning (e.g., RUD'18: pre-training on MT task). C-embeddings: contextual word embeddings (BERT-base). ML: multi-label prediction; LR: multi-variate Likert regression. *BERT concat*: last four BERT layers are concatenated. *BERT lin. comb.*: optimized linear combination of last four BERT layers (our BERT based models sum the last four layers). The Table is further discussed in the Section *Discussion* §4.3.

achieve a significantly higher macro-average ($\Delta\rho$: +0.014). Yet again, when our model is provided with contextual word embeddings, a large performance boost is the outcome. In fact, **MarkerEB** outperforms **MarkerE** by +0.045 Δ macro ρ and RUD'18's best performing configuration by +0.046. This stands in contrast to the multi-labeling results on this type of data, where the contextual embeddings did not appear to provide any performance increase. As previously discussed (§4.1), this discrepancy may be rooted in the task generation process of SPR2 which requires transforming two annotations per example to one ground truth (the two annotations per example stem from two out of, in total, 50 workers).

4.3 Discussion

Leaving the contextual word embedding inputs aside, the performance differences of our Marker models to RUD'18 may seem marginal for many properties. Albeit our **MarkerE** yields an observable improvement of 1.0 pp. macro F1 in the multi-label setup, in the regression setup the performance gains are very small ($\Delta\rho$ SPR1: +0.005, $\Delta\rho$ SPR2: +0.001, Table 1 & 3). In addition, our model as a single model instance (**Marker**) is outperformed by RUD'18's approach both in the regression and in the multi-label setup. However, it is important to note that the result of our system has substantially fewer dependencies (Table 3).

Firstly, our model does not rely on supplementary gold syntax – in fact, since it is span-based, our model is completely agnostic to any syntax. Besides our approach, only REI'15 does not depend on supplementary gold syntax for the dis-

played results. However, *all* of our models outperform REI'15's feature-based linear classifier in *every* property (+17 pp. macro F1 in total) except for *destroyed* (where **MarkerEB** performed slightly worse by -2.8 pp. F1). Also, results of SRL systems on semantic role labeling data show that span-based SRL systems often lag behind a few points in accuracy (cf. Li et al. (2019), Table 1). When provided with the syntactic head of the argument phrase, a model may immediately focus on what is important in the argument phrase. When solely fed the argument-span, which is potentially very long, the model has to find the most important parts of the phrase on its own and is more easily distracted. Additionally, identifying the head word of an argument may be more important than getting the boundaries of the span exactly right. In other words, span-based SPRL models may be more robust when confronted with slightly erroneous spans compared to dependency-based models which may be vulnerable to false positive heads. However, this hypothesis has to be confirmed or disproven experimentally in future work.

Further information about differences of various SPRL approaches is displayed in Table 3. In sum, despite having significantly fewer dependencies on external resources, our approach proves to be competitive with all methods from prior works, including neural systems. When combined in a simple ensemble, our model outperforms previous systems. When we feed additional contextual word embeddings, the results can be boosted further by a large margin. In the following section, we show that ensembling SPRL models has

Data	Ensembling worse		Ensembling better		NS
	$p \in [0.005, 0.05)$	$p < 0.005$	$p \in [0.005, 0.05)$	$p < 0.005$	
SPR1	0	1	5	9	3
SPR2	1	1	5	2	5

Table 4: McNemar significance test results of **MarkerEB** against **MarkerB**. Counts of properties for which a significance category applies (NS: #properties with insignificant difference).

another advantage besides predictive performance gains. Namely, it decreases SPRL model sensitivities towards different random initializations. As a result, we find that a simple neural voter committee (ensemble) offers more robust SPRL predictions compared to a randomly selected committee-member (single model).

4.4 Analysis

Ensembling increases accuracy To investigate whether ensembling improves proto-role labeling results significantly, we conduct McNemar significance tests (McNemar, 1947) comparing the predictions of **MarkerB** and **MarkerEB**. The significance test results summarized in Table 4 are unambiguous: for many proto-role properties, ensembling helps to improve performance significantly (SPR1: $\frac{14}{18}$ cases; SPR2: $\frac{7}{14}$ cases; significance level: $p < 0.05$). However, for few cases, ensembling resulted in significantly worse predictions (SPR1: *change of location*; SPR2: *change of location*, *instigated* and *partitive*; significance level: $p < 0.05$). For the rest of the properties, differences in predictions remain insignificant.

Ensembling increases robustness Additionally, we find that ensembling increases the robustness of our neural models. Consider Figure 4, where we display the performance difference of a n -voter ensemble to the same ensemble after one additional voter joined ($n+1$ -voter ensemble). The difference fluctuates wildly while the ensemble is still small. This suggests that a different random seed yields significantly different predictions. Hence, our a single neural Marker model is very vulnerable to the quirks of random numbers. However, when more voters join, we see that the predictions become notably more stable. An outcome which holds true for both data sets and two different ensemble model configurations (**MarkerE** and **MarkerEB**). We draw two conclusions from this experiment: (i) a single neural SPRL model is extremely sensitive to different random initial-

Data	ablated component					
	MarkerEB	SelfAtt	mark.	pred-mark.	arg-mark.	hier.
SPR1	77.5	73.1	50.7	76.3	60.1	76.7
SPR2	69.9	67.3	59.3	68.2	63.3	68.9

Table 5: Multi-labeling F1 macro scores for different MarkerEB model configurations over SPR1 and SPR2.

izations. Finally, (ii) a simple voter ensemble has the potential to alleviate this issue. In fact, when we add more voters, the model converges towards stable predictions which are less influenced by the quirks of random numbers.

Model Ablations All ablation experiments are conducted with **MarkerEB** in the multi-label formulation. We proceed by ablating different components in a leave-one-out fashion: (i) the self-attention components of the ensemble model are removed (SelfAtt in Table 5); (ii) we abstain from highlighting (a) the arguments and predicates, (b) only the predicates and (c) only the arguments (mark. pred-mark. and arg-mark. in Table 5); Finally, (iii) we remove the hierarchical structure and do not predict auxiliary outputs (hier. in Table 5).

From all ablated components, removing simultaneously both predicate and argument-markers hurts the model the most (SPR1: -26.8 pp. macro F1; SPR2: -10.6). Only ablating the argument-marker also causes a great performance loss (SPR1: -17.4, SPR2: -6.6). On the other hand, when only the predicate marker is ablated, the performance decreases only slightly (SPR1: -1.2, SPR2: -1.7). In other words, it appears to be of paramount importance to provide our model with indicators for the argument position in the sentence, but it is of lesser importance to point at the predicate. The self-attention component can boost the model’s performance by up to +4.4 pp. F1 on SPR1 and +2.6 on SPR2. The hierarchical structure with intermediate auxiliary Likert scale outputs leads to gains of approximately +1 pp. macro F1 in both data sets. This indicates that indeed the finer Likert scale annotations provide auxiliary information of value when predicting the labels at the top layer, albeit the performance difference appears to be rather small.

5 Conclusion

In our proposed SPRL ensemble model, predicate-argument constructs are highlighted with concurrently learned marker embeddings and self-

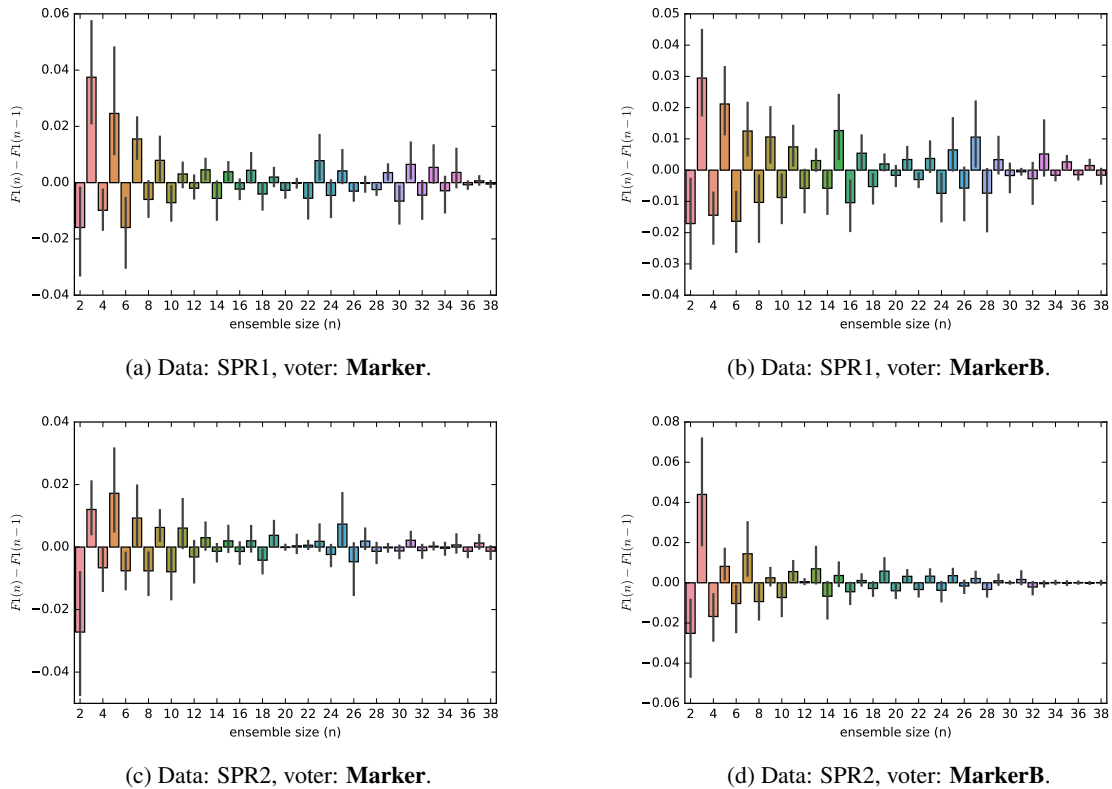


Figure 4: Adding more voters leads to convergence in SPRL predictions. x -axis: number of voter models partaking in the ensemble. y -axis: F1 mean difference over all proto-roles from the ensemble with x voters compared to the ensemble with $x - 1$ voters. Thin bars represent standard deviations.

attention enables the model to capture long-distance relationships between arguments and predicates. The span-based method is competitive with the dependency-based state-of-the-art which uses gold heads. When combined in a simple ensemble, the method overall outperforms the state-of-the-art on newspaper texts (multi-label prediction macro F1: +1.0 pp.). When fed with contextual word embeddings extracted from a large language model, the method outperforms the state-of-the-art by 6.4 pp. macro F1. Our method is competitive with the state-of-the-art for Likert regression on texts from the web domain. In the multi-label setting, it outperforms all baselines by a large margin. Furthermore, we have shown that a simple Marker model voter ensemble is very suited for conducting SPRL, for two reasons: (i) results for almost every proto-role property are significantly improved and (ii) considerably more robust SPRL predictions are obtained.

We hope that our work sparks more research into semantic proto-role labeling and corpus creation. Dowty’s feature-based view on roles allows us to analyze predicate-argument configurations in

great detail – an issue which we think is located in the marrow of computational semantics.

Acknowledgements

This work has been supported by the German Research Foundation as part of the Research Training Group Adaptive Preparation of Information from Heterogeneous Sources (AIPHES) under grant no. GRK 1994/1 and by the Leibniz ScienceCampus “Empirical Linguistics and Computational Language Modeling”, supported by the Leibniz Association under grant no. SAS-2015-IDS-LWC and by the Ministry of Science, Research, and Art of Baden-Württemberg.

References

- Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. [A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware?](#) In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2753–2765. Association for Computational Linguistics.
- Xavier Carreras and Lluís Màrquez. 2005. Introduction to the conll-2005 shared task: Semantic role la-

- beling. In *Proceedings of the ninth conference on computational natural language learning (CoNLL-2005)*, pages 152–164.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- David Dowty. 1991. [Thematic proto-roles and argument selection](#). *Language*, 67:547–619.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational linguistics*, 28(3):245–288.
- Jan Hajič, Massimiliano Ciaramita, Richard Johansson, Daisuke Kawahara, Maria Antònia Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Štěpánek, et al. 2009. The conll-2009 shared task: Syntactic and semantic dependencies in multiple languages. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–18. Association for Computational Linguistics.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018a. [Jointly predicting predicates and arguments in neural semantic role labeling](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 364–369, Melbourne, Australia. Association for Computational Linguistics.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. [Deep semantic role labeling: What works and what’s next](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 473–483, Vancouver, Canada. Association for Computational Linguistics.
- Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018b. [Syntax for semantic role labeling, to be, or not to be](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 2061–2071.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end neural coreference resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark. Association for Computational Linguistics.
- Zuchao Li, Shexia He, Hai Zhao, Yiqing Zhang, Zhuosheng Zhang, Xi Zhou, and Xiang Zhou. 2019. [Dependency or span, end-to-end uniform semantic role labeling](#). *CoRR*, abs/1901.05280.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. Towards robust linguistic analysis using ontonotes. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152.
- Drew Reisinger, Rachel Rudinger, Francis Ferraro, Craig Harman, Kyle Rawlins, and Benjamin Van Durme. 2015. [Semantic proto-roles](#). *Transactions of the Association for Computational Linguistics*, 3:475–488.
- Rachel Rudinger, Adam Teichert, Ryan Culkin, Sheng Zhang, and Benjamin Van Durme. 2018. [Neural-davidsonian semantic proto-role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 944–955. Association for Computational Linguistics.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038.
- Mihai Surdeanu, Richard Johansson, Adam Meyers, Lluís Màrquez, and Joakim Nivre. 2008. The conll-2008 shared task on joint parsing of syntactic and semantic dependencies. In *Proceedings of the Twelfth Conference on Computational Natural Language Learning*, pages 159–177. Association for Computational Linguistics.
- Adam R. Teichert, Adam Poliak, Benjamin Van Durme, and Matthew R. Gormley. 2017. [Semantic proto-role labeling](#). In *AAAI*, pages 4459–4466. AAAI Press.

Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.

Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016. Universal decompositional semantics on universal dependencies. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723.

Guineng Zheng, Subhabrata Mukherjee, Xin Luna Dong, and Feifei Li. 2018. Opentag: Open attribute value extraction from product profiles. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1049–1058. ACM.

A Supplemental Material

A.1 Notes

Calculation of macro F1 The global performance metric for multi-label SPRL is defined as ‘macro F1’. To ensure full comparability of results, we use the same formula as prior works (Rudinger et al., 2018):

$$\frac{2 \cdot P_{macro\ avg.} \cdot R_{macro\ avg.}}{P_{macro\ avg.} + R_{macro\ avg.}}, \quad (5)$$

where P and R are *Precision* and *Recall* and *macro avg.* means the unweighted mean of these quantities computed over all proto-role properties. The above macro F1 metric, though not explicitly displayed in the prior work papers, has been confirmed by the main authors (email).

Calculation of Pearson’s ρ Person’s ρ quantifies the linear relationship between two random variables X and Y . Computed over a sample $\{(x_i, y_i)\}_1^n$ it is calculated with the following formula:

$$\rho = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}} \quad (6)$$

Given $|P|$ proto-role properties and corresponding correlation coefficients $\rho_1, \dots, \rho_{|P|}$, the macro Pearson’s ρ is calculated as

$$macro\ \rho = \frac{\sum_{i=1}^{|P|} \rho_i}{|P|}. \quad (7)$$

hyper-parameter	choice
λ (main loss)	1
λ' (aux. losses)	0.2
optimizer	Adam (Kingma and Ba, 2014)
optimizer param.	$\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 1e^{-07}$
learning rate	0.001
Bi-LSTM units	$2 \cdot 64$
max. seq length	30
padding	pre
Marker embeddings, init	$U(-0.05, +0.05)$
Fixed embeddings, init	GloVe 300d (Pennington et al., 2014)

Table 6: Hyper parameter configuration.

Data split of SPR1 (Teichert et al., 2017) re-framed the SPRL task as a multi-label problem. Previously the task was to answer, given a predicate and an argument, one specific proto-role question (binary label or single output regression). Now we need to predict all proto-role questions at once (multi-label or multi-output regression). In order to allow this formulation of the task, the authors needed to redefine the original train-dev-test split of SPR1 (recent works, including ours, all use the re-defined split).

Reported Numbers In the EMNLP publication of Rudinger et al. (2018) we found a few minor transcription errors in the result tables (confirmed by email communication with the main authors, who plan to upload an errata section). In the case of transcription errors, we took the error-corrected numbers which were sent to us via email.

A.2 Hyperparameters & Preprocessing

The hyper parameter configuration of our model are displayed in Table 6. *Sequence pruning*: Consider that $I = \{i\}$ is the index of the predicate and $J = \{j, \dots, k\}$ are the indices corresponding to the argument. As long as the input sequence length is longer than maximum length (30, cf. Table 6), we clip left tokens so that the index of the token $m < \min I \cup J$, then we proceed to clip tokens to the right so that $m > \max I \cup J$, for the very rare cases that this was not sufficient we proceed to clip tokens with $m \notin I \cup J$ (the marker sequences are adjusted accordingly). The clipping strategy ensures that predicate and argument tokens are present in every input sequence. Sequences which are shorter than 30 words are *pre-padded* with zero vectors.