

Improving Implicit Semantic Role Labeling by Predicting Semantic Frame Arguments

Quynh Ngoc Thi Do¹, Steven Bethard², Marie-Francine Moens¹

¹Katholieke Universiteit Leuven, Belgium

²University of Arizona, United States

quynhngocthi.do@cs.kuleuven.be

bethard@email.arizona.edu

sien.moens@cs.kuleuven.be

Abstract

Implicit semantic role labeling (iSRL) is the task of predicting the semantic roles of a predicate that do not appear as explicit arguments, but rather regard common sense knowledge or are mentioned earlier in the discourse. We introduce an approach to iSRL based on a predictive recurrent neural semantic frame model (PRNSFM) that uses a large unannotated corpus to learn the probability of a sequence of semantic arguments given a predicate. We leverage the sequence probabilities predicted by the PRNSFM to estimate selectional preferences for predicates and their arguments. On the NomBank iSRL test set, our approach improves state-of-the-art performance on implicit semantic role labeling with less reliance than prior work on manually constructed language resources.

1 Introduction

Semantic role labeling (SRL) has traditionally focused on semantic frames consisting of verbal or nominal predicates and *explicit arguments* that occur *within* the clause or sentence that contains the predicate. However, many predicates, especially nominal ones, may bear arguments that are left implicit because they regard common sense knowledge or because they are mentioned earlier in a discourse (Ruppenhofer et al., 2010; Gerber et al., 2009). These arguments, called *implicit arguments*, are resolved by another semantic task, implicit semantic role labeling (iSRL). Consider a NomBank (Meyers et al., 2004) annotation example:

[_{A0} *The network*] had been expected to have [_{NP} *losses*] [_{A1} *of \$20 million*] ... Those [_{NP} *losses*] may widen because of the short Series.

The predicate *loss* in the first sentence has two arguments annotated explicitly: A0, *the entity losing something*, and A1, *the thing lost*. Meanwhile, the other instance of the same predicate in the second sentence has no associated arguments. However, for a good reader, a reasonable interpretation of the second *loss* should be that it receives the same A0 and A1 as the first instance. These arguments are implicit to the second *loss*.

As an emerging task, implicit semantic role labeling faces a lack of resources. First, hand-crafted implicit role annotations for use as training data are seriously limited: SemEval 2010 Task 10 (Baker et al., 1998) provided FrameNet-style (Baker et al., 1998) annotations for a fairly large number of predicates but with few annotations per predicate, while Gerber and Chai (2010) provided PropBank-style (Palmer et al., 2005) data with many more annotations per predicate but covering just 10 predicates. Second, most existing iSRL systems depend on other systems (explicit semantic role labelers, named entity taggers, lexical resources, etc.), and as a result not only need iSRL annotations to train the iSRL system, but annotations or manually built resources for all of their sub-systems as well.

We propose an iSRL approach that addresses these challenges, requiring no manually annotated iSRL data and only a single sub-system, an explicit semantic role labeler. We introduce a predictive recurrent neural semantic frame model (PRNSFM), which can estimate the probability of a sequence of semantic arguments given a predicate, and can be trained on unannotated data drawn from the Wikipedia, Reuters, and Brown corpora, coupled with the predictions of the MATE (Björkelund et al., 2010) explicit semantic role labeler on these texts. The PRNSFM forms the foundation for our iSRL system, where we use its probability estimates over sequences of semantic arguments to predict *selectional preferences* for associating predicates with

their implicit semantic roles. Our PRNSFM-based iSRL model improves state-of-the-art performance, outperforming the only other system that depends on just an explicit semantic role labeler by 10 % F1, and achieving equal or better F1 score than several other models that require many more lexical resources.

Our work fits today’s interest in natural language understanding, which is hampered by the fact that content in a discourse is often not expressed explicitly because it was mentioned earlier or because it regards common sense or world knowledge that resides in the mind of the communicator or the audience. In contrast, humans easily combine relevant evidence to infer meaning, determine hidden meanings and make explicit what was left implicit in the text, using the *anticipatory power* of the brain that predicts or “imagines” circumstantial situations and outcomes of actions (Friston, 2010; Vernon, 2014) which makes language processing extremely effective and fast (Kurby and Zacks, 2015; Schacter and Madore, 2016). The neural semantic frame representations inferred by our PRNSFM take a first step towards encoding something like anticipatory power for natural language understanding systems.

The remainder of the paper is organized as follows: First, section 2 describes the related work. Second, section 3 proposes the predictive recurrent neural semantic frame model including the formal definition, architecture, and an algorithm to extract selectional preferences from the trained model. Third, in section 4, we introduce the application of our PRNSFM in implicit semantic role labeling. Fourth, the experimental results and discussions are presented in section 5. Finally, we conclude our work and suggest some future work in section 6.

2 Related work

Language Modeling Language models, from n-gram models to continuous space language models (Mikolov et al., 2013; Pennington et al., 2014), provide probability distributions over sequences of words and have shown their usefulness in many natural language processing tasks. However, to our knowledge, they have not yet been used to model semantic frames. Recently, Peng and Roth (2016) developed two distinct models that capture semantic frame chains and discourse information while abstracting over the specific mentions of predicates and entities, but these models focus on discourse

processing tasks, not semantic frame processing.

Semantic Role Labeling In unsupervised SRL, Woodsend and Lapata (2015) and Titov and Khoddam (2015) induce embeddings to represent a predicate and its arguments from unannotated texts, but in their approaches, the arguments are words only, not the semantic role labels, while in our models, both are considered.

Low-resource Implicit Semantic Role Labeling

Several approaches have attempted to address the lack of resources for training iSRL systems. Laparra and Rigau (2013) proposed an approach based on exploiting argument coherence over different instances of a predicate, which did not require any manual iSRL annotations but did require many other manually-constructed resources: an explicit SRL system, WordNet super-senses, a named entity tagger, and a manual categorization of SuperSenseTagger semantic classes. Roth and Frank (2015) generated additional training data for iSRL through comparable texts, but the resulting model performed below the previous state-of-the-art of Laparra and Rigau (2013). Schenk and Chiarcos (2016) proposed an approach to induce prototypical roles using distributed word representations, which required only an explicit SRL system and a large unannotated corpus, but their model performance was almost 10 points lower than the state-of-the-art of Laparra and Rigau (2013). Similar to Schenk and Chiarcos (2016), our model requires only an explicit SRL system and a large unannotated corpus, but we take a very different approach to leveraging these, and as a result improve state-of-the-art performance.

3 Predictive Recurrent Neural Semantic Frame Model

Our goal is to use unlabeled data to acquire selectional preferences that characterize how likely a phrase is to be an argument of a semantic frame. We rely on the fact that current explicit SRL systems achieve high performance on verbal predicates, and run a state-of-the-art explicit SRL system on unlabeled data. We then construct a predictive recurrent neural semantic frame model (PRNSFM) from these explicit frames and roles.

Our PRNSFM views semantic frames as a *sequence*: a predicate, followed by the arguments in their textual order, and terminated by a special EOS symbol. We draw predicates from PropBank

verbal semantic frames, and represent arguments with their nominal/pronominal heads. For example, *Michael Phelps swam at the Olympics* is represented as [swam:PRED, Phelps:A0, Olympics:AM-LOC, EOS], where the predicate is labeled PRED and the arguments *Phelps* and *Olympics* are labeled A0 and AM-LOC, respectively. Our PRNSFM’s task is thus to take a predicate and zero or more arguments, and predict the next argument in the sequence, or EOS if no more arguments will follow.

We choose to model semantic frames as a sequence (rather than, say, a bag of arguments) because in English, there are often fairly strict constraints on the order in which arguments of a verb may appear. A sequential model should thus be able to capture these constraints and use them to improve its probability estimates. Moreover, a sequential model has the ability to learn the interaction between arguments in the same semantic frame. For example, considering a swimming event, if *Phelps* is A0, then *Olympics* is more likely to be the AM-LOC than *lake*.

Formally, for each t^{th} argument of a semantic frame f , we denote its word (e.g., *Phelps*) as $w_{f,t}$, its semantic label (e.g., A0) as $l_{f,t}$, where $w \in \mathbf{V}$, the word vocabulary, and $l \in \mathbf{L} \cup [\text{PRED}]$, the set of semantic labels. We denote the predicate word and label, which are always at the 0^{th} position in the sequence, in the same way as arguments: $w_{f,0}$ and $l_{f,0}$. We denote the sequence $[w_{f,0}, w_{f,1}, \dots, w_{f,t-1}]$ as $w_{f,<t}$, and the sequence $[l_{f,0}, l_{f,1}, \dots, l_{f,t-1}]$ as $l_{f,<t}$. Our model aims to estimate the conditional probability of the occurrence of $w_{f,t}$ as semantic role $l_{f,t}$ given the preceding words and their labels:

$$P(w_{f,t}:l_{f,t}|w_{f,<t}:l_{f,<t})$$

We use a recurrent neural network to learn this probability distribution over sequences of semantic frame arguments. For a semantic frame f with N arguments, at each time step $0 \leq t \leq N$, given the input $w_{f,t}:l_{f,t}$, the model computes the distribution $P(w_{f,t+1}:l_{f,t+1}|w_{f,<t+1}:l_{f,<t+1})$ and predicts the next most likely argument (or EOS). During training, model parameters are optimized by minimizing prediction errors over all time steps.

We consider two versions of this model that differ in input (\mathbf{V}_{in}) and output (\mathbf{V}_{out}) vocabularies.

3.1 Model 1: Joint Embedding LSTM

We adopt the standard recurrent neural network language model (Mikolov et al., 2010), which is a

natural architecture to deal with a sequence prediction problem.

Model 1 consists of three layers (see Figure 1): an embedding layer that learns vector representations for input values; a long short-term memory (LSTM) layer that controls the sequential information receiving the vector representation as input; and a softmax layer to predict the next argument using the output of the LSTM layer as input.

This model treats the word and semantic label as a single unit in both input and output layers. The model, therefore, learns joint embeddings for the word and its corresponding semantic label. For example, if we take “Michael Phelps swam at the Olympics” as training data, the three input values would be *swam:PRED*, *Phelps:A0* and *Olympics:AM-LOC*, and the three expected outputs would be *Phelps:A0*, *Olympics:AM-LOC*, EOS. Since each word:label is considered as a single unit, the embedding layer will learn three vector representations, one for *swam:PRED*, one for *Phelps:A0*, and one for *Olympics:AM-LOC*. As can be seen, an important difference between our problem and the traditional language model is that we have to deal with two different types of information – word and label. By concatenating word and label, the standard recurrent neural network model can be applied directly to our data.

The detail of Model 1 is as following:

Embedding Layer is a matrix of size $|\mathbf{V}_{\text{in}}| \times d$ that maps each unit of input into an d -dimensional vector. The matrix is initialized randomly and updated during network training.

LSTM Layer consists of m standard LSTM units which take as input the output of the embedding layer, x_t , and produce an output h_t by updating at every time step $0 \leq t \leq T$:

$$\begin{aligned} i_t &= \text{sigmoid}(W_i x_t + U_i h_{t-1} + b_i) \\ \hat{C}_t &= \text{tanh}(W_c x_t + U_c h_{t-1} + b_c) \\ f_t &= \text{sigmoid}(W_f x_t + U_f h_{t-1} + b_f) \\ C_t &= i_t * \hat{C}_t + f_t * C_{t-1} \\ o_t &= \text{sigmoid}(W_o x_t + U_o h_{t-1} + b_o) \\ h_t &= o_t * \text{tanh}(C_t) \end{aligned}$$

where W_i, W_c, W_f, W_o are weight matrices of size $d \times m$; U_i, U_c, U_f, U_o are weight matrices of size $m \times m$; b_i, b_c, b_f, b_o are bias vectors of size m ; and $*$ is element-wise multiplication. As per the standard LSTM formulation, $i_t, \hat{C}_t, f_t, C_t, o_t$ represent

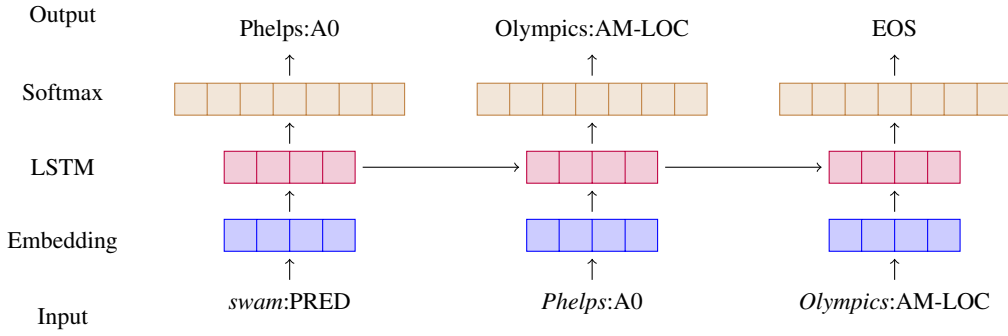


Figure 1: Model 1 – Joint Embedding LSTM

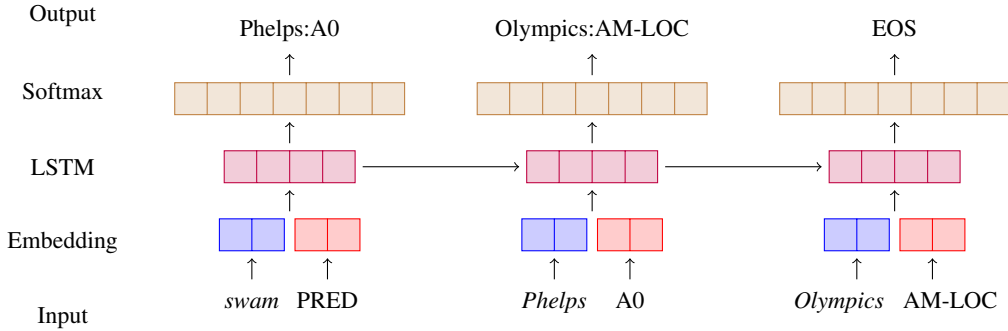


Figure 2: Model 2 – Separate Embedding LSTM

the input gate, states of the memory cells, activation of the memory cells’ forget gates, memory cells’ new state, and output gates’ values, respectively.

Softmax Layer computes the probability distribution of the next argument given the preceding arguments at time step t :

$$P(w_{f,t+1}:l_{f,t+1}|w_{f,<t+1}:l_{f,<t+1}) = \text{softmax}(h_t W + b) \quad (1)$$

where W is a weight matrix of size $m \times |\mathbf{V}_{\text{out}}|$, and b is a bias vector of size $|\mathbf{V}_{\text{out}}|$. The predicted next argument is:

$$\underset{w_{f,t+1}:l_{f,t+1}}{\text{argmax}} P(w_{f,t+1}:l_{f,t+1}|w_{f,<t+1}:l_{f,<t+1})$$

The network is trained using the negative log-likelihood loss function.

3.2 Model 2: Separate Embedding LSTM

Model 2 shares the same basic structure as Model 1, but considers the word and the semantic label as two different units in the input layer. As shown in Figure 2, we use two different embedding layers, one for word values and one for semantic labels, and the two embedding vectors are concatenated before being passed to the LSTM layer. The LSTM

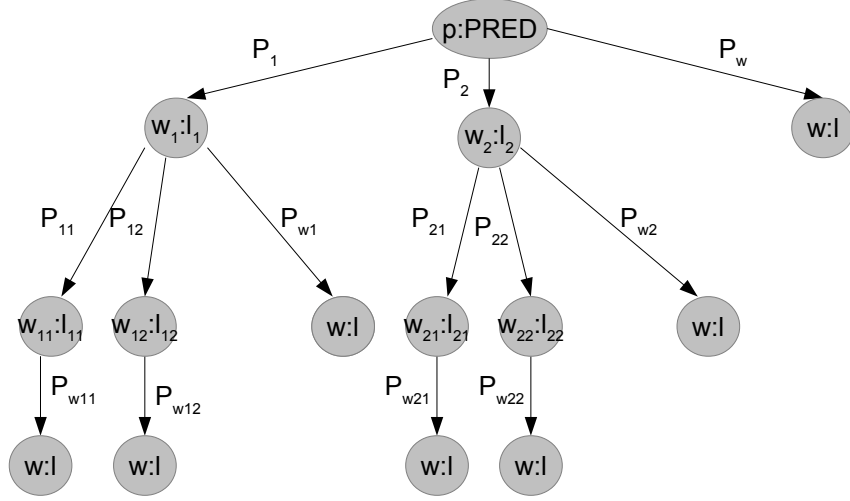
and softmax layers are then the same as in Model 1. For example, if we take “Michael Phelps swam at the Olympics” as training data, the three input words would be *swam*, *Phelps*, and *Olympics*, the three input roles would be PRED, A0 and AM-LOC, and the three expected outputs would be *Phelps:A0*, *Olympics:AM-LOC*, EOS. A total of six different vector representations will be learned: a word embedding for each of *swam*, *Phelps*, and *Olympics*, and a label embedding for each of PRED, A0 and AM-LOC.

In this model, the embedding layer for labels is initialized randomly (as in Model 1), but the embedding layer for word values is initialized with publicly available word embeddings that have been trained on a large dataset (Mikolov et al., 2013).

As compared to the joint-embedding Model 1, the separate-embedding Model 2 gives up a little power to represent the interaction between words and labels, but has a less sparse input vocabulary and gains the ability to incorporate pre-trained word embeddings.

3.3 Selectional Preferences

While the PRNSFM can predict the probability of an argument given the predicate and the preceding arguments, $P(w_{f,t}:l_{f,t}|w_{f,<t}:l_{f,<t})$, an iSRL



Time 0: $S_0 = [[p:PRED]]$

Time 1: $S_1 = [[p:PRED, w_1:l_1], [p:PRED, w_2:l_2]]$

Time 2: $S_2 = [[p:PRED, w_1:l_1, w_{11}:l_{11}], [p:PRED, w_1:l_1, w_{12}:l_{12}], [p:PRED, w_2:l_2, w_{21}:l_{21}], [p:PRED, w_2:l_2, w_{22}:l_{22}]]$

$$P(w:l|p) \sim P_w + P_{w_1} P_1 + P_{w_2} P_2 + P_{w_{11}} P_{11} P_1 + P_{w_{12}} P_{12} P_1 + P_{w_{21}} P_{21} P_2 + P_{w_{22}} P_{22} P_2$$

Figure 3: Selectional Preference Inference example: $k=2, T=3$. The possible sequences are represented as a tree. Each arrow label is the probability of the target node to be predicted given the path from the tree root to the parent of the target node.

system needs a *selectional preference* score representing the probability of a word w being the l argument of predicate p , $P(w:l|p:PRED)$. Thus, to convert our PRNSFM probabilities to selectional preferences, we need to marginalize over the possible argument sequences.

We approximate this marginalization by constructing a tree where the root is the predicate, p , the branches are likely sequences of arguments, and the leaves are the word and label for which we need to estimate a probability, $w:l$. Formally, we define this tree of possible sequences as:

$$S_t = \begin{cases} \{[p:PRED]\} & \text{if } t = 0 \\ \{[q, w_t:l_t] : q \in S_{t-1}, \\ w_t:l_t \in \mathbf{argmax}^k(q)\} & \text{if } 0 < t < T \\ \{[q, w:l] : q \in S_{t-1}\} & \text{if } t = T \end{cases}$$

where $w_{f,0}:l_{f,0} = p:PRED$; k and T are thresholds; and $\mathbf{argmax}^k(q)$ is the k word:label pairs that have the highest probability of being the next argument given the sequence q according to the PRNSFM.

We then estimate $P(w:l|p:PRED)$ as the sum of the probabilities of all the sequences encoded in

the tree. Formally:

$$P(w:l|p:PRED) \approx \sum_{0 \leq t \leq T} P(w:l|w_{f,<t+1}:l_{f,<t+1}) \\ \approx \sum_{0 \leq t \leq T} \sum_{q \in S_t} P(w:l|q) \times P(q)$$

where the probability of an argument sequence q is the product of the PRNSFM's estimates for each step in the sequence:

$$P(q) = P(w_t:l_t|w_{t-1}:l_{t-1}, \dots, p:PRED) \\ \times P(w_{t-1}:l_{t-1}|w_{t-2}:l_{t-2}, \dots, p:PRED) \\ \times \dots \times P(w_1:l_1|p:PRED) \quad (2)$$

An example of the calculation of $P(w:l|p:PRED)$ is shown in Figure 3.

Intuitively, the tree enumerates all possible argument sequences that start with the predicate, have zero or more intervening arguments, and end with the word and label of interest, $w:l$. The probability of $w:l$ given the predicate is the sum of the probabilities of all branches in this tree, i.e., of all possible sequences that end with $w:l$. In reality, we do not have the computational power to explore

all possible sequences, so we must limit the tree somehow. Thus, we only ask the PRNSFM for its top k predictions at each branch point, and we only explore sequences with a maximum length of T .

4 Implicit Semantic Role Labeling

As you will recall from previous sections, implicit semantic role labeling is the task of identifying discourse-level arguments of a semantic frame, which are missed by standard semantic role labeling, which operates on individual sentences. For instance, in “This house has a new owner. The sale was finalized 10 days ago.”, the semantic frame evoked by “sale” in the second sentence should receive “the house” as an implicit A1 semantic role. Humans easily resolve the object of the sale given the candidates (in our example: “house” and “owner”), but for a machine this is more difficult unless it has knowledge on what the likely objects of a sale are. This kind of knowledge of selectional preferences can be extracted from our trained PRNSFM.

The previous section described how to extract selectional preferences from our PRNSFM. However, that model is trained on verbal predicates, and the test data that we use (Gerber and Chai, 2010) contains nominal predicates. Thus, for each triple of a nominal predicate np , a word candidate w , and a label l , we approximate the selectional preference score of w being the implicit argument role l of np as:

$$P(w:l|np) = \max_{p \in V(np)} P(w:l|p:\text{PRED})$$

where $P(w:l|p)$ is the selectional preference score described in Section 3.3, and $V(np)$ is set of verbal forms of np . Here, we use the NomBank lexicon to get verbs associated with each nominal predicate, and then find instances of those verbs in the explicit SRL training data. For example, for the noun *funds*, $V(\textit{funds}) = \{\textit{funds}, \textit{fund}, \textit{funding}, \textit{funded}\}$.

We apply selectional preferences to iSRL following (Laparra and Rigau, 2013). For each nominal predicate np and implicit label l , the current and previous two sentences are designated the *context window*. Each sentence in the context window is annotated with the explicit SRL system. If any instances of np or $V(np)$ in the text have an explicit argument of type l , we deterministically predict the closest such argument as the implicit l argument of np . Otherwise, we run the PRNSFM over each word in the context window, and select the

word with the highest selectional preference score above a threshold s . If all the candidates’ scores are less than s , the system leaves the missing argument unfilled. We optimized this threshold on the development data, resulting in $s = 0.0003$.

As in Laparra and Rigau (2013), we apply a *sentence recency factor* to emphasize recent candidates. The selectional preference score x is updated as $x' = x - z + z \times \alpha^d$ where d is the sentence distance, and α and z are parameters. We set $z = 0.00005$ based on the development set and set $\alpha = 0.5$ as in (Laparra and Rigau, 2013).

5 Experiments

We evaluate the two PRNSFM models on the iSRL task. The tools, resources, and settings we used are as follows:

Semantic Role Labeling We used the full pipeline from MATE (<https://code.google.com/archive/p/mate-tools/>) (Björkelund et al., 2010) as the explicit SRL system, retraining it on just the CoNLL 2009 training portion.

Unannotated Data The unannotated data used in the experiments was drawn from Wikipedia (<http://corpus.byu.edu/wiki/>), Reuters (<http://about.reuters.com/researchandstandards/corpus/>), and Brown (<https://catalog.ldc.upenn.edu/ldc99t42>).

Dataset for PRNSFM The first 15 million short and medium (less than 100 words) sentences from the unannotated data (described above) were annotated automatically by the explicit SRL system. The obtained annotations were then used together with the gold standard CoNLL 2009 SRL training data to train the PRNSFM.

Neural network training and inference Parameters were selected using the CoNLL 2009 development set. We set the dimensions of word and label embeddings in the PRNSFM to 50 and 16, respectively. The hidden sizes of LSTM layers are the same as their input sizes. Word embedding layers are initialized by Skip-gram embeddings learned by training the word2vec tool (Mikolov et al., 2013) on the unannotated data. Our models were trained for 120 epochs using the AdaDelta optimization algorithm (Zeiler, 2012). For fast selectional preference computing, we set $k = 1$ and $T = 4$ ¹.

¹We selected relatively small values for the parameters to reduce the training and prediction time. We tried some larger values of the parameters on a small dataset, but found that the

Evaluation We follow the evaluation setting in Gerber and Chai (2010); Laparra and Rigau (2013); Schenk and Chiarcos (2016)²: the method is evaluated on the evaluation portion of the nominal iSRL data by Dice coefficient metrics. For each missing argument position of a predicate instance, the system is required to either (1) identify a single constituent that fills the missing argument position or (2) make no prediction and leave the missing argument position unfilled. To give partial credit for inexact argument boundaries, predictions are scored by using the Dice coefficient, which is defined as follows:

$$Dice(predicted, true) = \frac{2 |predicted \cap true|}{|predicted| + |true|}$$

Predicted contains the tokens that the model has identified as the filler of the implicit argument position. *True* is the set of tokens from a single annotated constituent that truly fill the missing argument position. The model’s prediction receives a score equal to the maximum Dice overlap across any of the annotated fillers (AF)³:

$$Score(predicted) = \max_{true \in AF} Dice(predicted, true)$$

Precision is equal to the summed prediction scores divided by the number of argument positions filled by the model. Recall is equal to the summed prediction scores divided by the number of argument positions filled in the annotated data.

5.1 Experimental Setup

In the **baseline mode**, instead of using the PRNSFM, we only use the deterministic prediction by the explicit SRL system. We refer to this mode as *Baseline* in Table 1.

In the **main mode**, the joint embedding LSTM model (Model 1) and the separate embedding LSTM model (Model 2) were trained on the same dataset which is a combination of the automatic SRL annotations and the gold standard CoNLL

small values reported in the article achieved similar results with faster processing times.

² Following Schenk and Chiarcos (2016), we do not perform the alternative evaluation of Gerber and Chai (2012) that evaluates systems on the iSRL training set, since the iSRL training set overlaps with the CoNLL 2009 explicit semantic role training set on which MATE is trained.

³For iSRL, one implicit role may receive more than one annotated filler across a coreference chain in the discourse.

2009 training data as described in the previous section. We denote this mode as *gold CoNLL 2009 + unlabeled* in Table 1.

To evaluate how well the system acquires knowledge from unlabeled data, we also train the PRNSFM only on the gold standard CoNLL 2009 training data. We denote this mode as *CoNLL 2009* in Table 1.

In order to compare the performance of our sequential model to a non-sequential model, we train a skip-gram neural language model on the same unlabeled and labeled data as the PRNSFM in the main mode. The skip-gram model treats the predicates and arguments as a bag of labeled words rather than a sequence. The $P(w:l|p)$ is computed at the output layer of the skip-gram model by considering $w:l$ as the context of p . We denote this mode as *Skip-gram* in Table 1.

5.2 Results and Discussion

Table 1 shows the prior state-of-the-art and the performance of the baseline, skip-gram and our PRNSFM-based methods.

Our Model 2 achieves the highest precision and F1 score. This is notable because the first two models require many more language resources than just an explicit SRL system: Gerber and Chai (2010) use WordNet and manually annotated iSRL data, while Laparra and Rigau (2013) use WordNet, named entity annotations, and manual semantic category mappings. Schenk and Chiarcos (2016), like our approach, use only an explicit SRL system, but both our models strongly outperform their results. We assume that the difference here is caused by our proposed neural semantic frame model (PRNSFM). Schenk and Chiarcos (2016) measure the selectional preference of a predicate and a role as a cosine between a standard word2vec embedding for the candidate word, and the average of all word2vec embeddings for all words that appear in that role. Our algorithms are very different: we take a language modeling approach and leverage the sequence of semantic roles, we learn custom word/role embeddings tuned for SRL, and then marginalize over many possible argument sequences. We assume that the learned PRNSFM representations are better informed about semantic frames than simple word embeddings, which only capture knowledge of contextual words.

Table 1 also shows that training on large unlabeled data results in a marked improvement com-

Method	PRNSFM training data	iSRL data	SRL system	WordNet	NER system	P	R	F1
Gerber and Chai (2010)		✓	✓	✓		44.5	40.4	42.3
Laparra and Rigau (2013)			✓	✓	✓	47.9	43.8	45.8
Schenk and Chiarcos (2016)			✓			33.5	39.2	36.1
Baseline			✓			75.3	17.2	28.0
Skip-gram	gold CoNLL 2009 + unlabeled		✓			26.3	32.3	29.0
Model 1: Joint Embedding	gold CoNLL 2009 + unlabeled		✓			48.0	38.2	42.6
Model 2: Separate Embedding	gold CoNLL 2009 + unlabeled		✓			52.6	41.0	46.1
Model 1: Joint Embedding	gold CoNLL 2009		✓			39.2	34.1	36.5
Model 2: Separate Embedding	gold CoNLL 2009		✓			40.2	36.0	38.0

Table 1: Implicit role labeling evaluation.

pared to training on only the CoNLL 2009 labeled data, providing evidence that the models have acquired linguistic knowledge from the unlabeled data. Although the automatically annotated data used to train the PRNSFM can be noisy, using a large amount of data has smoothed out the noise.

Moreover, the better performance of our models over the standard skip-gram neural language model proves the effectiveness of modeling semantic frames as sequential data. The intuition here is that explicit semantic arguments have typical orderings in which they occur, so a sequential model should be a good fit for this problem. Modeling this sequential aspect of the problem is effective, but requires us to marginalize out positional information to compute selectional preferences, since implicit semantic arguments can occur anywhere in the discourse and do not have a typical position.

Among our two models, Model 2, which learns separate vector representations for words and semantic roles, is better than Model 1, which learns a single vector representation of each (word, semantic role) pair. The separate representation of words and roles means that Model 2 can share information across multiple occurrences of a word even if the semantic roles of that word are different, and this model can use publicly available embeddings pre-trained from even larger unannotated corpora when initializing its embeddings.

Gerber and Chai (2012) report an inter-annotator agreement of 64.3% using Cohen’s kappa measure on the annotated NomBank-based iSRL data. This value is borderline between low and moderate agreement indicating the sheer complexity of the

Predicate	Baseline	2010	2013	2016	2017
sale	36.2	44.2	40.3	37.2	52.8
price	15.4	34.2	53.3	27.3	29.0
investor	9.8	38.4	41.2	33.9	43.1
bid	32.3	21.3	52.0	40.7	35.5
plan	38.5	64.7	40.7	47.4	76.8
cost	34.8	62.9	53.0	36.9	44.4
loss	52.6	83.3	65.8	58.9	72.8
loan	18.2	37.5	22.2	37.9	38.6
investment	0.0	30.8	40.8	36.6	23.5
fund	0.0	15.4	44.4	37.5	42.8

Table 2: A comparison on F1 scores (%). 2010: (Gerber and Chai, 2010), 2013: (Laparra and Rigau, 2013), 2016: Best model from (Schenk and Chiarcos, 2016), 2017: Our best model (Model 2).

annotation task, and explaining the relatively low performance of the iSRL systems.

In Table 2, we compare the F1 scores over all the ten predicates of our Model 2 to other state-of-the-art systems⁴. Our system obtains relatively high scores (> 50%) on three predicates including “sale”, “plan” and “loss”. These three are the most frequent predicates (among the 10 defined in the nominal iSRL dataset) in the CoNLL 2009 training data – they occur 1016, 318 and 275 times in verbal forms, respectively. In contrast, irregular predicates such as “bid” or “loan” usually have low performance. This is possibly caused by the de-

⁴As an overly conservative estimate, we take a t-test over the 10 predicate-level F1 scores as can be seen in Table 2. Comparing against Model 2, this yields p=0.28 for Gerber and Chai (2010), p=0.46 for Laparra and Rigau (2013), and most importantly p=0.058 for Schenk and Chiarcos (2016).

pendence of our PRNSFM on the performance of the explicit semantic role labeling system on verbal predicates.

It is important to consider how iSRL can be extended beyond the 10 annotated predicates of Gerber and Chai (2010). Our models do not require any handcrafted iSRL annotations for training, and thus can be applied to all predicates observed in large unannotated data on which they are trained.

However, as other work in iSRL, our approach still relies on a resource-heavy SRL system to learn selectional preferences. It would be interesting to investigate in further studies whether this SRL system can be replaced by a low-resource system (Collobert et al., 2011; Connor et al., 2012).

6 Conclusion and Future Work

We have presented recurrent neural semantic frame models for learning probability distributions over semantic argument sequences. By modeling selectional preferences from these probability distributions, we have improved state-of-the-art performance on the NomBank iSRL task while using fewer language resources. In the future, we believe that our semantic frame models are valuable in many language processing tasks that require discourse-level understanding of language, such as summarization, question answering and machine translation.

Acknowledgment

This work is carried out in the frame of the EU CHIST-ERA project “MULTimodal processing of Spatial and TEmporal expReSSions” (MUSTER), and the “MACHine Reading of patient recordsS” project (MARS, KU Leuven, C22/015/016).

References

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The berkeley framenet project](#). In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 1*. Association for Computational Linguistics, Stroudsburg, PA, USA, ACL '98, pages 86–90. <https://doi.org/10.3115/980845.980860>.

Anders Björkelund, Bernd Bohnet, Love Hafdell, and Pierre Nugues. 2010. [A high-performance syntactic and semantic dependency parser](#). In *Coling 2010: Demonstrations*. Coling 2010 Organizing Committee, Beijing, China, pages 33–36. <http://www.aclweb.org/anthology/C10-3009>.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. [Natural language processing \(almost\) from scratch](#). *J. Mach. Learn. Res.* 12:2493–2537. <http://dl.acm.org/citation.cfm?id=1953048.2078186>.

M. Connor, C. Fisher, and D. Roth. 2012. Starting from Scratch in Semantic Role Labeling: Early Indirect Supervision. *Cognitive Aspects of Computational Language Acquisition*.

Karl Friston. 2010. [The free-energy principle: a unified brain theory?](#) *Nature Reviews Neuroscience* 11(2):127–138. <https://doi.org/10.1038/nrn2787>.

Matt Gerber, Joyce Y. Chai, and Adam Meyers. 2009. [The role of implicit argumentation in nominal srl](#). In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, PA, USA, NAACL '09, pages 146–154. <http://dl.acm.org/citation.cfm?id=1620754.1620776>.

Matthew Gerber and Joyce Chai. 2010. [Beyond NomBank: A Study of Implicit Arguments for Nominal Predicates](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Uppsala, Sweden, pages 1583–1592. <http://www.aclweb.org/anthology-new/P/P10/P10-1160.bib>.

Matthew Gerber and Joyce Y. Chai. 2012. Semantic role labeling of implicit arguments for nominal predicates. *Computational Linguistics* 38(4):755–798.

C.A. Kurby and J.M. Zacks. 2015. *Situation Models in Naturalistic Comprehension*, Cambridge University Press, pages 59–76.

Egoitz Laparra and German Rigau. 2013. [Impar: A deterministic algorithm for implicit semantic role labelling](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers*. pages 1180–1189. <http://aclweb.org/anthology/P/P13/P13-1116.pdf>.

A. Meyers, R. Reeves, C. Macleod, R. Szekely, V. Zielinska, B. Young, and R. Grishman. 2004. The nombank project: An interim report. In A. Meyers, editor, *HLT-NAACL 2004 Workshop: Frontiers in Corpus Annotation*. Association for Computational Linguistics, Boston, Massachusetts, USA, pages 24–31.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). *CoRR* abs/1301.3781. <http://arxiv.org/abs/1301.3781>.

Tomas Mikolov, Martin Karafit, Lukas Burget, Jan Cernock, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In Takao Kobayashi, Keikichi Hirose, and Satoshi Nakamura, editors, *INTERSPEECH*. ISCA, pages 1045–1048.

- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics* 31(1):71–106. <http://www.cs.rochester.edu/gildea/palmer-propbank-cl.pdf>.
- Haoruo Peng and Dan Roth. 2016. Two discourse driven language models for semantics. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. <http://aclweb.org/anthology/P/P16/P16-1028.pdf>.
- 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. <http://www.aclweb.org/anthology/D14-1162>.
- Michael Roth and Anette Frank. 2015. Inducing Implicit Arguments from Comparable Texts: A Framework and its Applications. *Computational Linguistics* 41:625–664.
- Josef Ruppenhofer, Caroline Sporleder, Roser Morante, Collin Baker, and Martha Palmer. 2010. Semeval-2010 task 10: Linking events and their participants in discourse. In *Proceedings of the 5th International Workshop on Semantic Evaluation*. Association for Computational Linguistics, Stroudsburg, PA, USA, SemEval '10, pages 45–50. <http://dl.acm.org/citation.cfm?id=1859664.1859672>.
- D.L. Schacter and K.P. Madore. 2016. Remembering the past and imagining the future: Identifying and enhancing the contribution of episodic memory. *Memory Studies* 9(3):245–255.
- Niko Schenk and Christian Chiarcos. 2016. Unsupervised learning of prototypical fillers for implicit semantic role labeling. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*. pages 1473–1479. <http://aclweb.org/anthology/N/N16/N16-1173.pdf>.
- Ivan Titov and Ehsan Khoddam. 2015. Unsupervised induction of semantic roles within a reconstruction-error minimization framework. In *Proceedings of the North American chapter of the Association for Computational Linguistics (NAACL)*.
- David Vernon. 2014. *Artificial Cognitive Systems: A Primer*. The MIT Press.
- Kristian Woodsend and Mirella Lapata. 2015. Distributed representations for unsupervised semantic role labeling. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, pages 2482–2491. <http://aclweb.org/anthology/D15-1295>.
- Matthew D. Zeiler. 2012. ADADELTA: an adaptive learning rate method. *CoRR* abs/1212.5701. <http://arxiv.org/abs/1212.5701>.