

Semantic Role Labeling Meets Definition Modeling: Using Natural Language to Describe Predicate-Argument Structures

Simone Conia^{1*}

Edoardo Barba^{1*}

Alessandro Scirè^{1,2}

Roberto Navigli¹

Sapienza NLP Group

Sapienza University of Rome

¹{first.lastname}@uniroma1.it

Babelscape, Italy

²scire@babelscape.com

Abstract

One of the common traits of past and present approaches for Semantic Role Labeling (SRL) is that they rely upon discrete labels drawn from a predefined linguistic inventory to classify predicate senses and their arguments. However, we argue this need not be the case. In this paper, we present an approach that leverages Definition Modeling to introduce a generalized formulation of SRL as the task of describing predicate-argument structures using natural language definitions instead of discrete labels. Our novel formulation takes a first step towards placing interpretability and flexibility foremost, and yet our experiments and analyses on PropBank-style and FrameNet-style, dependency-based and span-based SRL also demonstrate that a flexible model with an interpretable output does not necessarily come at the expense of performance. We release our software for research purposes at <https://github.com/SapienzaNLP/dsrl>.

1 Introduction

Commonly regarded as one of the key ingredients for Natural Language Understanding (Navigli, 2018), Semantic Role Labeling (Gildea and Jurafsky, 2002, SRL) aims at identifying “Who did What to Whom, Where, When, and How?” within a given sentence (Màrquez et al., 2008). More precisely, for each predicate in the sentence, the task requires: i) selecting its most appropriate sense from a predetermined linguistic inventory; ii) identifying its arguments, i.e., those parts of the sentence that are semantically related to the predicate; and, iii) assigning a semantic role to each predicate-argument pair, as shown in Figure 1. Due to the potential uses of these semantically rich structures, the research community has seen steady progress in the task, and SRL has been shown to be beneficial

* Equal contribution.

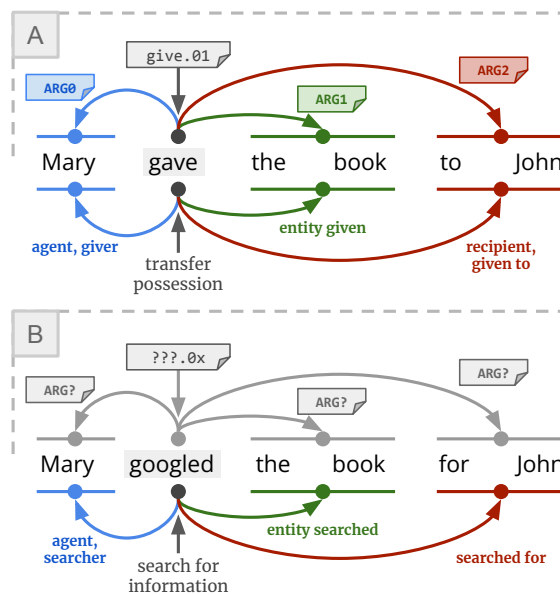


Figure 1: **A:** SRL annotations using predicate sense and semantic role labels (top) compared with their natural language definitions (bottom). **B:** the semantics of sense and role labels is undefined for out-of-inventory predicates (e.g., the inventories used for CoNLL-2009 and CoNLL-2012 do not include an entry for “google”), but we can still use valid natural language definitions.

for an increasingly wide range of applications in Natural Language Processing (NLP), such as Question Answering (Shen and Lapata, 2007), Information Extraction (Christensen et al., 2011), Machine Translation (Marcheggiani et al., 2018), and Summarization (Mohamed and Oussalah, 2019), as well as in Computer Vision for Situation Recognition (Yatskar et al., 2016) and Video Understanding (Sadhu et al., 2021), *inter alia*.

An important yet often overlooked aspect of SRL is that, since its conception, the formulation of the task has generally relied upon predetermined linguistic resources, such as FrameNet (Baker et al., 1998), PropBank (Palmer et al., 2005), VerbNet (Kipper Schuler, 2005) and, more recently,

VerbAtlas (Di Fabio et al., 2019), which provide the labels to be used for tagging predicates and their arguments with senses and semantic roles, respectively. Therefore, to this day, SRL has been framed predominantly as a classification task in which systems assign *discrete labels* to portions of a sentence (Figure 1A, top). Although recent systems have achieved impressive results on standard benchmarks (Hajič et al., 2009; Pradhan et al., 2012) in English (Shi and Lin, 2019; Marcheggiani and Titov, 2020) as well as in multilingual SRL (He et al., 2019; Conia et al., 2021), we observe and emphasize that relying upon discrete labels raises the following critical questions:

- The assumption that both predicate senses and semantic roles can be unequivocally categorized into distinct classes has long been – and still is – at the center of numerous discussions because the boundaries between meanings are not always clear-cut (Tuggy, 1993; Hanks, 2000); unsurprisingly, disambiguation approaches that are not tied to specific inventories have been gaining momentum (Bevilacqua et al., 2020; Barba et al., 2021a,b).
- FrameNet, PropBank, and VerbNet are heterogeneous, non-overlapping resources that have led, consequently, to specialized techniques that are more effective on PropBank’s rather than FrameNet’s labels, or vice versa.
- Relying on any predetermined inventory hinders the ability to generalize to out-of-inventory instances. For example, some rare senses or neologisms may not be covered by the inventory of choice, which, therefore, does not define either their possible senses, or their corresponding semantic roles (Figure 1B, top).¹

Furthermore, recent progress in NLP at large has primarily pursued state-of-the-art results without giving much importance as to why a system may have a predilection for one particular option over the alternatives, thus making it difficult for a human

¹In several linguistic inventories, semantic role labels are defined according to specific predicate senses, i.e., they are sense-specific. This is the case for PropBank, in which core arguments (ARG0 through ARG5) acquire meaning only with respect to a predicate sense, and for FrameNet, in which some frame elements are specific to the frame they belong to (e.g., *Ingestible* is only defined for the frame *Ingestion*). We note that this is not the case for some inventories such as VerbNet and VerbAtlas, whose semantic roles generalize across frames.

to interpret their output. And SRL is no exception to this.

In this paper, instead, we put forward a generalized formulation of Definition Modeling – the task of defining the meaning of a word or multiword expression in context – to reframe SRL as the task of describing sentence-level semantic relations between a predicate and its arguments using natural language definitions only. More specifically, our contributions can be summarized as follows:

1. We move away from discrete labels and introduce a novel formulation of SRL that reframes the problem as the task of using natural language to describe predicate-argument structures (Figure 1A, bottom).
2. We propose DSRL (Descriptive Semantic Role Labeling), a simple yet effective conditional generation model to produce such natural language descriptions, dropping discrete labels while also demonstrating how to use these descriptions to retrieve standard SRL labels and achieve competitive or even state-of-the-art results on gold benchmarks.
3. In contrast to previous work, our approach provides an interpretable output in natural language, can seamlessly produce descriptions according to different linguistic theories and annotation formalisms, and naturally admits descriptions for out-of-inventory instances (Figure 1B, bottom).
4. We provide an in-depth analysis of the strengths and pitfalls of our approach, showing where there is still room for improvement.

We hope that our semantically-driven descriptions in natural language, free of resource-specific labels that require expert knowledge of SRL, will not only enable easier integration of sentence-level semantics into downstream applications but also provide valuable insights to NLP researchers.

2 Related Work

Linguistic resources for SRL. As mentioned above, SRL is generally associated with a linguistic theory and a corresponding linguistic resource, which defines an inventory of predicate senses and semantic roles² (Baker et al., 1998; Palmer et al.,

²More precisely, FrameNet delineates *frames* and *frame elements*, while VerbNet uses *classes* and *thematic roles*. Here-

2005; Kipper Schuler, 2005). These inventories are a rich and diverse source of expert-curated knowledge; however, aligning sense and semantic role labels across such resources using manual or automatic techniques (Giuglea and Moschitti, 2006; Palmer, 2009; Lopez de Lacalle et al., 2014; Stowe et al., 2021; Conia et al., 2021) is far from trivial due to their heterogeneous nature, variable degree of coverage, and different granularity. Perhaps it is this complexity that has led researchers towards the development of approaches that are effective mainly in just one of the task “styles”, usually PropBank-style SRL (Marcheggiani et al., 2017; Cai et al., 2018; Strubell et al., 2018; Shi and Lin, 2019; Blloshmi et al., 2021; Conia and Navigli, 2022, *inter alia*) or FrameNet-style SRL (Swayamdipta et al., 2017; Peng et al., 2018; Lin et al., 2021; Pancholy et al., 2021, *inter alia*). To sidestep this situation, recent studies have analyzed the feasibility of moving away from rigorous linguistic resources and have looked into capturing predicate-argument relations as question-answer pairs, with promising results in the production of questions through slot-filling templates and generative models (He et al., 2015; FitzGerald et al., 2018; Pyatkin et al., 2021). In this paper, instead, we reframe SRL as a generalization of Definition Modeling and directly generate human-readable descriptions of the semantic relations between a predicate and its arguments, replacing discrete labels with natural language definitions to overcome the heterogeneities of linguistic inventories.

Recent approaches in SRL. Independently of the linguistic inventory of choice, given the complexity of the task, early work often employed separate systems for each step of the SRL pipeline (Roth and Lapata, 2016; Marcheggiani et al., 2017). However, in recent years, researchers have successfully managed to develop end-to-end approaches (Cai et al., 2018; He et al., 2018), especially due to the increasing expressiveness of recent neural architectures. Since then, the attention of the community has mainly focused on when syntactic features are useful (Strubell et al., 2018) or can be dispensed with (Conia and Navigli, 2020). Further to this, several studies have also investigated the effectiveness of their proposed approaches on different annotation formalisms, namely, dependency- and

after, for simplicity, we follow PropBank and call them *senses* and *semantic roles*, respectively, independently of the resource.

span-based SRL (Li et al., 2019; Marcheggiani and Titov, 2020). Most recently, sequence-to-sequence models have found renewed traction by learning to directly generate predicate-argument structures as linearized sequences (Blloshmi et al., 2021; Paolini et al., 2021). Although the focus of our approach is to generate natural language descriptions, we stress that it can be flexibly employed to perform SRL in its traditional formulation, jointly tackling predicate sense disambiguation, argument identification and labeling in a syntax-agnostic fashion for both span- and dependency-based formalisms, the key difference being that our method also produces human-readable and, therefore, interpretable descriptions of the semantics of a sentence.

Definition Modeling. The task of Definition Modeling was originally concerned with producing a natural language definition for a given word and its corresponding embedding (Noraset et al., 2017). The formulation of the task was later generalized to take polysemy into account, as the same word may convey different meanings depending on the context it appears in. Although introduced a few years ago now, Definition Modeling has attracted significant interest (Ni and Wang, 2017; Ishiwatari et al., 2019) and has found success in semantic tasks (Huang et al., 2019; Bevilacqua et al., 2020) such as Word Sense Disambiguation (Bevilacqua et al., 2021, WSD) and Word-in-Context (Pilehvar and Camacho-Collados, 2019, WiC). Motivated by the success of Definition Modeling, we propose a novel generalization of its formulation, in which the objective is to use natural language not only to define a target word in context but also to describe its semantically-relevant sentential constituents.

3 Describing Predicate-Argument Structures using Natural Language

In this Section, we introduce our novel reformulation of the SRL task (Section 3.1), describe DSRL, a simple yet effective autoregressive approach for it (Section 3.2), and show how to use DSRL to perform standard SRL (Section 3.3).

3.1 Task Formulation

Taking inspiration from Definition Modeling, we propose addressing predicate sense disambiguation, argument identification, and argument classification in an end-to-end fashion as the task of describing the argument structure of a predicate p in a sentence s by generating a natural language

description t^p that defines not only p but also the semantic relations that connect p to its arguments $a_1, a_2, \dots, a_{|A|}$, where A is the set of arguments of p . For example, if we consider the predicate $p = \text{“gave”}$ in the sentence $s = \text{“Mary gave the book to John”}$, then a valid natural language description of p and its argument structure could be represented as $t^p = \text{“give: transfer. [Mary]{giver} gave [the book]{thing given} [to John]{entity given to}”}$. Indeed, such a sequence contains i) the predicate definition for predicate sense disambiguation, ii) all the arguments of p in s within square brackets for argument identification, along with iii) a definition of the semantic role of each argument within curly brackets.

3.2 Description Generation

To tackle our SRL formulation, we introduce a simple end-to-end autoregressive approach that, given an input sentence s and a predicate p in s , generates the natural language description t^p of its argument structure. In particular, we devise a sequence-to-sequence model whose input sequence s^p is defined as follows:

$$s^p = w_1 \dots w_i \dots \\ \langle p \rangle p_1 \dots p_k \langle /p \rangle \dots w_n$$

where w_i is the i -th word in the original sentence s , while $\langle p \rangle$ and $\langle /p \rangle$ are two special markers that indicate the beginning and the end, respectively, of the predicate p , with $k > 1$ if p is a multiword expression. Correspondingly, we instruct the model to generate a semantically-augmented sentence t^p in which: i) the sense definition of p is prepended to the original sentence, ii) the arguments of p are enclosed within square brackets, and, iii) each argument is followed by its semantic role definition within curly brackets. More formally:

$$t^p = p_1 \dots p_k : d_1^p \dots d_{k'}^p . \\ w_1 \dots [w_1^{a_1} \dots w_{m_1}^{a_1}] \{d_1^{a_1} \dots d_{m_1}^{a_1}\} \\ \dots [w_1^{a_2} \dots w_{m_2}^{a_2}] \{d_1^{a_2} \dots d_{m_2}^{a_2}\} \\ \vdots \\ \dots [w_1^{a_j} \dots w_{m_j}^{a_j}] \{d_1^{a_j} \dots d_{m_j}^{a_j}\} \dots w_n$$

where p_i is the i -th word of the predicate p , d_i^p is the i -th word of the definition of p , $w_i^{a_j}$ is the i -th word for the j -th argument of p , and $d_i^{a_j}$ is the i -th word of the definition of the semantic role

for the j -th argument of p , while k' , m_j and m_j' are the length of the definition of p , the length of the argument a_j , and the length of the definition of the semantic role for a_j , respectively. With this encoding, we then train our sequence-to-sequence model to learn the factorized probability $p(t^p | s^p)$ defined as follows:

$$p(t^p | s^p) = p(t_1^p | s^p) \prod_{j=2}^{|t^p|} p(t_j^p | t_{1:j-1}^p, s^p)$$

by minimizing the cross-entropy loss with respect to the generated natural language description.

3.3 From SRL to Natural Language and Back

Given a dataset annotated with predicate sense and role labels from an inventory that defines such labels in natural language, we note that it is always possible to convert such a dataset to our formulation.³ Moreover, although the main objective of our approach is to generate an output sequence that describes sentence-level semantics, in several scenarios, it is still useful to work with discrete labels for predicate senses and semantic roles, e.g., to assess the quality of the generated structures on gold benchmarks with their standard metrics. We stress that our formulation generalizes standard SRL; casting the descriptions generated by our model to standard SRL labels is only possible if the label inventory of choice defines a suitable sense for the target predicate, which is not the case in Figure 1B (top) as the verb “to google” is not covered by PropBank. If the predicate is covered by the inventory, we can easily select the sense or the role label \bar{y} whose natural language description $d^{\bar{y}}$ is most similar to the definition d' generated for the predicate p or for one of its arguments a_j . We select \bar{y} as follows:

$$\bar{y} = \operatorname{argmax}_{y \in Y} \sigma(f(d^y), f(d'))$$

where $\sigma(\cdot)$ is a similarity function (e.g., cosine similarity), $f(\cdot)$ provides a vector representation of a definition, Y is the set of labels, and d^y is the definition of y as provided by the inventory of choice. We note that, for simplicity, we do not apply any post-processing to enforce the validity of the generated output, leaving more complex strategies (e.g., constrained decoding) as future work.

³We also note that dependency-based annotations can be seen as span-based annotations and, thus, used directly as arguments in our natural language descriptions.

4 Experiments and Results

4.1 Data

We train and evaluate DSRL on three widely adopted benchmarks for English SRL, namely: i) CoNLL-2009 (Hajič et al., 2009) for dependency-based PropBank-style SRL, ii) CoNLL-2012 (Pradhan et al., 2012) for span-based PropBank-style SRL, and iii) FrameNet 1.7 (Baker et al., 1998) for span-based FrameNet-style SRL. While CoNLL-2009 is a collection of finance-related news from the Wall Street Journal, CoNLL-2012 is a more heterogeneous corpus comprising news, conversations, and magazine articles. FrameNet 1.7, instead, provides a relatively small dataset of annotated documents; following the literature (Swayamdipta et al., 2017; Peng et al., 2018), we include in the training set “exemplar” sentences extracted from partially annotated usage examples from the lexicon itself. We provide a broader look at the characteristics of each dataset in Appendix B and further details about semantic role definitions in Appendix D.

4.2 Implementation Details

We implement DSRL using Sunglasses.ai’s Classy.⁴ As our underlying sequence-to-sequence model, we use BART-large (Lewis et al., 2020), a Transformer-based neural network (400M parameters) pretrained with denoising objectives on massive amounts of unlabeled text.⁵ We do not modify its architecture except for the embedding layer, where we add the special tokens used to indicate predicates and their arguments,⁶ as described in Section 3.2. We train our model using RAdam (Liu et al., 2019) as the optimizer for a maximum of 500 000 steps with a batch size of 2048 tokens and a standard learning rate of 10^{-5} . We measure the F1 score on the validation set at the end of each training epoch, adopting an early stopping strategy to interrupt the training process if the F1 score does not improve for 10 consecutive epochs. We do not modify any of the hyperparameters of BART compared to its pretraining phase, and, more generally, we do not run any hyperparameter search due to the cost of fine-tuning the language model. The training process is carried out on a single GPU (a GeForce RTX 3090) and requires about 10 hours for FrameNet, 15 for CoNLL-2009

⁴<https://github.com/sunglasses-ai/classy>

⁵We use the model’s weights available from Huggingface Transformers Wolf et al. (2020).

⁶See Appendix F for further details on the special tokens.

and 20 for CoNLL-2012.

We recall that, in order to evaluate our system with standard scoring scripts,⁷ we have to cast our descriptions to the discrete labels of the target inventory (see Section 3.3). For this step, we compute the cosine similarity between the representation of a generated description and those of the possible senses or roles, using the sentence-level embeddings of SimCSE (Gao et al., 2021).⁸

4.3 Comparison Systems

We compare our results with the current state of the art in PropBank-style and FrameNet-style SRL. Following standard practice in PropBank-based SRL, we report the results achieved by our system using gold pre-identified (but not disambiguated) predicates, i.e., the position of a predicate (but not its sense label) is given as input to the system.

PropBank-style SRL. We consider Li et al. (2019), who first quantified the benefits of contextualized word representations in both dependency- and span-based PropBank-style SRL, later surpassed by Shi and Lin (2019), who used BERT instead of ELMo, and Conia and Navigli (2020), who designed and took advantage of complex language-agnostic components. We also take into account some studies for PropBank-style SRL that found success by leveraging syntactic features such as He et al. (2019), who devised a strategy to cleverly prune a sentence based on its syntactic dependency tree, and Marcheggiani and Titov (2020), who exploited graph convolutional networks to encode syntactic relations. Most recently, Biloshmi et al. (2021) proposed a simple and general approach to tackle SRL as a sequence-to-sequence task, in which, however, a system is still required to generate a linearized sequence of discrete labels.

FrameNet-style SRL. Although the research community has generally focused on PropBank-style SRL, especially due to the widespread adoption of PropBank in several CoNLL tasks (Carreras and Màrquez, 2005; Surdeanu et al., 2008; Hajič et al., 2009; Pradhan et al., 2012) and in other resources such as Abstract Meaning Representation (Banarescu et al., 2013, AMR), FrameNet-style SRL has also been at the center of notable studies such as Swayamdipta et al. (2017), who investigated the effect of joint learning of syntactic and

⁷[eval09.pl](#) for CoNLL-2009, [srl-eval.pl](#) for CoNLL-2012, and [fnSemScore.pl](#) for FrameNet.

⁸[princeton-nlp/sup-simcse-roberta-base](#).

semantic features, and Peng et al. (2018), who instead showed the advantages of learning from disjoint data sources. Finally, we also consider recent work by Pancholy et al. (2021), who developed a data augmentation strategy using frame relations, and the above-mentioned Marcheggiani and Titov (2020), who introduced a graph-based neural architecture to tackle FrameNet-style SRL.

4.4 Main Results

Here, we first evaluate the robustness of DSRL in achieving strong or even state-of-the-art results on standard benchmarks, and then its flexibility in performing dependency- and span-based, PropBank and FrameNet-style SRL. Remarkably, our model achieves even better results when jointly trained on dissimilar annotation formalisms and linguistic resources, despite their heterogeneous characteristics.

PropBank-style SRL. We first discuss the results obtained by DSRL on the gold standard benchmarks provided as part of the CoNLL-2009 and CoNLL-2012 Shared Tasks, annotated with PropBank sense and role labels. As can be seen in Table 1, we observe strong results in dependency-based SRL, reaching an F1 score of 92.5% in the English test set of CoNLL-2009. Therefore, despite having to cast our natural language descriptions to discrete labels, our approach performs in the same ballpark as the most recent state-of-the-art systems proposed by Conia and Navigli (2020) and Blloshmi et al. (2021); the fact that our approach is able to slightly outperform the latter (+0.1% in F1 score) is particularly meaningful, as they adopt the same pretrained language model (BART-large). We can observe the same behavior in span-based SRL, where our model – without any task-specific modifications – marginally surpasses (+0.1% in F1 score) that of Blloshmi et al. (2021) on the English test set of CoNLL-2012, as shown in Table 2. Thus, the key observation here is that a natural language output does not necessarily hurt performance.

FrameNet-style SRL. As shown in Appendix E, PropBank definitions for predicate senses and semantic roles are quite short, and therefore one may wonder whether our task reformulation is feasible in practice when using longer definitions from richer sources, such as FrameNet, in which the label definitions are up to three times longer. From our experiments, this is, indeed, the case: our

CoNLL-2009	<i>P</i>	<i>R</i>	F1
Li et al. (2019)	89.6	91.2	90.4
He et al. (2019)	90.4	91.3	90.9
Shi and Lin (2019)	92.4	92.3	92.4
Conia and Navigli (2020)	92.5	92.7	92.6
Fei et al. (2021)	–	–	92.2
Blloshmi et al. (2021)	92.9	92.0	92.4
Zhang et al. (2022)	93.0	91.0	92.0
This work _{CoNLL-2009}	92.9	92.1	92.5
This work _{ALL}	92.3	92.4	92.4

Table 1: Results (%) on precision (P), recall (R) and F1 score on the English test set of CoNLL-2009.

CoNLL-2012	<i>P</i>	<i>R</i>	F1
Li et al. (2019)	85.7	86.3	86.0
Shi and Lin (2019)	85.9	87.0	86.5
Marcheggiani and Titov	86.5	87.1	86.8
Conia and Navigli (2020)	86.9	87.7	87.3
Blloshmi et al. (2021)	87.8	86.8	87.3
This work _{CoNLL-2012}	88.6	86.1	87.4
This work _{ALL}	87.7	87.1	87.4

Table 2: Results (%) on precision (P), recall (R) and F1 score on the English test set of CoNLL-2012.

FrameNet	<i>P</i>	<i>R</i>	F1
Swayamdipta et al. (2017)	70.5	66.7	68.6
Peng et al. (2018)	80.2	72.9	76.4
Marcheggiani and Titov	77.8	76.9	77.4
Pancholy et al. (2021)	72.1	70.2	71.1
This work _{FrameNet}	79.2	79.3	79.3
This work _{ALL}	79.9	79.9	79.9

Table 3: Results (%) on precision (P), recall (R) and F1 score on the English test set of FrameNet.

approach achieves state-of-the-art results in full-structure extraction (Baker et al., 2007) on the test set of FrameNet 1.7, obtaining 79.3 in F1 score (Table 3). We note that the results are not directly comparable with previous work, as DSRL employs a language model (BART) that is different from that of other approaches, e.g., Marcheggiani and Titov (2020) used RoBERTa. However, the results achieved by DSRL still indicate the performance that a generative approach can obtain in frame-semantic parsing (Das et al., 2014), which might be considered more complex than PropBank-based SRL. Indeed, predicates in FrameNet usually have a higher degree of polysemy, and the semantic roles are sparser, e.g., there are more than 2000 differ-

CoNLL-2009 (OOD)	P	R	F1
Li et al. (2019)	–	–	81.5
He et al. (2019)	–	–	82.2
Shi and Lin (2019)	–	–	92.4
Conia and Navigli (2020)	–	–	85.9
Blloshmi et al. (2021)	85.8	84.5	85.2
This work CoNLL-2009	86.4	84.8	85.6
This work ALL	86.1	86.4	86.3

Table 4: Results (%) on precision (P), recall (R) and F1 score on the English out-of-domain test set of CoNLL-2009.

ent semantic roles in FrameNet 1.7 compared to only 50-60 semantic roles in the PropBank releases used for the CoNLL-2009 and CoNLL-2012 shared tasks (see Table 8 in Appendix B).

Combining PropBank and FrameNet. The flexibility of our approach is evidenced by the fact that our model can benefit from learning to perform jointly dependency-based PropBank-style SRL on CoNLL-2009, span-based PropBank-style SRL on CoNLL-2012, and span-based FrameNet-style SRL on FrameNet 1.7, simply by enforcing two inventory-specific special tokens at the beginning of the decoding process, e.g., `<propbank><dep-srl> tp` or `<framenet><span-srl> tp`, where t^p is the target output, i.e., the semantically-augmented sentence described in Section 3.2. Using natural language descriptions instead of discrete labels as the common denominator across heterogeneous inventories yields similar – or even improved – results when training our model on the three resources at the same time, compared to training a separate model on each dataset, as reported in the last row of Tables 1, 2, and 3, removing the need for separate systems for different setups and empirically supporting the flexibility of our model in scaling across dissimilar formalisms (dependency- and span-based annotations) and linguistic theories (PropBank and FrameNet). Indeed, our model is able to leverage such features to achieve a new state of the art in the out-of-domain test set of CoNLL-2009. As shown in Table 4, when we train DSRL jointly on CoNLL-2009, CoNLL-2012, and FrameNet, we can observe a large improvement, achieving 86.3% in F1 score – +0.7% over training DSRL only on CoNLL-2009, and +1.1% over Blloshmi et al. (2021) – and setting a new state of the art on this out-of-domain benchmark, to the best of our knowledge.

5 Quantitative Analysis

5.1 Rare and Unseen Senses

The probability with which a word assumes one of its possible senses follows Zipf’s distribution (Kilgarriff, 2004), and thus it is very skewed towards the most frequent senses. Here, we analyze the bias that our system shows in predicting the most frequent predicate senses on the following partitions of the CoNLL-2009 and CoNLL-2012 test sets: i) **MFS**, all the instances containing predicates that are annotated with their most frequent sense; ii) **LFS**, all the instances containing predicates that are not annotated with their most frequent sense; iii) **UNSEEN**, all the instances containing predicates that are annotated with a sense that is not present in the training set.

As we can see from Table 5, the performance of our system on predicate sense disambiguation is strong in the MFS partition – more than 98.5% in both CoNLL-2009 and CoNLL-2012 – since the vast majority of predicates are annotated with their most frequent sense. This bias justifies the difference in F1 score between the MFS and LFS partitions, i.e., –11.9% and –9.3% on CoNLL-2009 and CoNLL-2012, respectively. As far as the UNSEEN partition is concerned, on the other hand, we observe that our approach seems to be capable of generating and retrieving senses that it has never seen at training time with a relatively low decrease in performance (–6.6% and –13.9% compared to the results on the LFS partition). Interestingly, the results on argument labeling are comparable between MFS and LFS predicates. However, there is still large room for improvement in the argument labeling of UNSEEN predicates, whose argument structure represents a more challenging zero-shot setting.

5.2 Data efficiency

Considering the large expense entailed in manually annotating text with sense and role labels, we deem it indispensable to also evaluate the flexibility of a system in terms of its scalability on fewer training instances. Therefore, we analyze the results of our model by gradually reducing the training set to 75%, 50%, 25%, and 10% of its original size, and compare this learning curve with that of GSRL (Blloshmi et al., 2021). Notwithstanding the significant differences between the two approaches, both show similar learning curves on CoNLL-2009 and CoNLL-2012 (Figure 2), confirming that manu-

	Dataset	ALL		MFS		LFS		UNSEEN	
		F1	Support	F1	Support	F1	Support	F1	Support
<i>Pred.</i>	CoNLL-2009	97.6	8986	98.9	8056 (89.7%)	87.0	838 (9.3%)	80.4	92 (1.0%)
	CoNLL-2012	96.7	62 002	98.5	50 607 (81.6%)	89.2	11 026 (17.8%)	75.3	369 (0.6%)
<i>Arg.</i>	CoNLL-2009	89.2	19 946	89.8	17 753 (89.0%)	87.2	1964 (9.8%)	58.5	229 (1.1%)
	CoNLL-2012	87.3	145 055	88.1	123 769 (85.3%)	83.9	20 350 (14.0%)	62.4	936 (0.6%)

Table 5: Predicate and argument labeling scores on the test sets of CoNLL-2009 and CoNLL-2012. We report the performance (F1) on the most frequent senses (MFS), least frequent senses (LFS) and unseen senses (UNSEEN). Support indicates the number of instances (percentage) of the corresponding class.

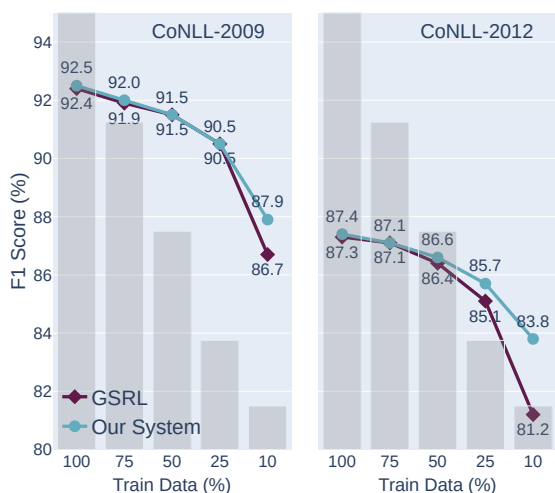


Figure 2: Performance comparison of our system and GSRL when down-sampling the training dataset to 10%, 25%, 50% and 75% of the total instances.

ally annotating more sentences eventually ceases to provide large improvements: in fact, the enormous effort of doubling the training instances of CoNLL-2012 by annotating other 100,000 predicates (from 50% to 100% of its original size) results in less than a 1.0% gain in F1 score. Interestingly, our system shows higher data efficiency in the lowest data regime, especially for span-based SRL with a 2.6% gain in F1 score over GSRL when they are both trained on 10% of the original dataset. We argue that our novel formulation better leverages the pretraining of the underlying language model in lower-data scenarios. However, when more training data is available, task-specific approaches are eventually able to close the gap.

Finally, we investigate whether our approach is still capable of handling multiple inventories at the same time in low-data regimes. To this end, we trained the model with several combinations of inventories on 10% of their training data. As we can see from Table 6, the model achieves im-

Training Data (10%)	CoNLL-09	CoNLL-12	FrameNet
CoNLL-2009 (C09)	87.9	–	–
CoNLL-2012 (C12)	–	83.8	–
FrameNet 1.7 (F17)	–	–	74.9
C09 + C12	88.4	84.2	–
C12 + F17	–	84.0	75.0
C09 + F17	87.9	–	75.1
C09 + C12 + F17	88.5	84.3	75.4

Table 6: Results of our model when trained on a random sample of 10% of the original training splits of CoNLL-2009, CoNLL-2012, and FrameNet 1.7 and their combinations.

proved results whenever it is trained on any two inventories, with the one trained jointly on CoNLL-2009, CoNLL-2012, and FrameNet performing best. Interestingly enough, the model is able to handle the CoNLL-2009 + FrameNet combination despite the different linguistic resources (PropBank vs FrameNet) and annotation formalism (dependency- vs span-based SRL).

6 Qualitative Analysis

6.1 Generation Examples

In Table 7, we provide some examples of the descriptions generated by our system. Given an input sentence, we compare its gold standard sequence (\hat{g}) with the one generated automatically (g). We find that, in some cases, the automatic descriptions are more contextual than the gold ones, occasionally overcoming the limitations of the linguistic inventories. In Example 1, for instance, the gold definition of the predicate *brandish.01* is only applicable to weapons; instead, the model-generated sequence is preferable as the entity brandished is a flag. In other cases, such as in Example 2, our approach generates more descriptive definitions, e.g., *depictor* instead of *agent*, and *thing described* rather than *theme*. Furthermore, we show some ex-

Ex. 1: Thousands of supporters, many <u>brandishing</u> flags ...	
Gold	brandish: wave weapons. ... brandishing [flags]{ weapon } ...
Pred	brandish: display, exhibit. ... brandishing [flags]{ entity displayed } ...
Ex. 2: [...] its unrealistic <u>depiction</u> of the characters' [...] private lives.	
Gold	depiction: show to be. ... [its]{ agent } [unrealistic]{instrument or manner} depiction [of]{ theme } ...
Pred	depiction: show to be. ... [its]{ depictor } [unrealistic] {instrument or manner} depiction [of]{ thing described } ...
Ex. 3: [...] he was " <u>nibbling</u> at" selected stocks during Friday's plunge.	
Gold	<i>n/a</i> (out of inventory)
Pred	nibble: eat lightly . [he]{eater} was nibbling [at selected stocks]{food} [during Friday's plunge]{time or duration}.
Ex. 4: Zaire's president Mobutu met with [...] a senior U.S. <u>envoy</u> .	
Gold	<i>n/a</i> (out of inventory)
Pred	envoy: stand for, correspond a senior [U.S.]{entity being substituted by the other} [envoy]{entity taking place of other}.

Table 7: Generation examples. Given an input sentence, we compare the gold and the system-generated sequence. Predicates are underlined.

amples in which the model generates semantically appropriate natural language descriptions for out-of-inventory, and thus unseen, predicates. Even in this setting, the model often generates semantically appropriate natural language descriptions. This is the case with Example 3, in which the model describes the semantics of *nibble.01* (unseen at training time) by taking advantage of a similar predicate, namely, *peck.01* (seen at training time). This is also true for noun predicates, as shown in Example 4.

6.2 Classes of Error

We identify three main classes of error: the first is directly connected to our system (*Disambiguation Errors*) and the other two (*Out-of-Inventory Descriptions* and *Retrieval Errors*) concern the noisy process we use to cast natural language descriptions to discrete class labels.

Disambiguation errors occur when the model generates a definition that does not describe the correct sense of a predicate in a given context. For example, the system provides the wrong definition for the predicate “bumble” in the following sentence *s*, misclassifying it as “speak quietly”:

s: Shane survived the week only to have an executive bumbling his way into a criminal investigation.

- **Gold**: speak or move in a confused way
- **Pred**: speak quietly

We note that, given the autoregressive nature of the model, producing a wrong sense definition often compromises the entire argument structure.

Out-of-inventory descriptions may be produced by our approach since it is not strictly tied to the vocabulary of a predefined linguistic resource. While our model can generate predicate-argument structures not present in the inventory, they can still provide correct semantic explanations. For instance, in the following sentence, the reference and the generated definitions convey the same semantics:

- **Gold**: dupe: *trick*. He meets [a French girl]{*tricker*} who dupes [him]{*tricked*} [into providing a home for her pet and then steals his car]{*induced action*}.
- **Pred**: dupe: *deceive*. He meets [a French girl]{*deceiver*} who dupes [him]{*victim*} [into providing a home for her pet and then steals his car]{*tricked into*}.

Associating “victim” to “tricked” is far from trivial, and such cases often result in **retrieval errors**, i.e., errors that are caused by the inability of the sentence embedding model – SimCSE in our case – to correctly capture the semantic similarity between the gold and generated definitions.

7 Conclusion

Recent progress in SRL has mainly revolved around the development of state-of-the-art systems which, however, are bound to specific predicate-argument inventories. In this paper, instead, we proposed a novel task formulation that takes a step towards putting interpretability and flexibility in the foreground: we reframed SRL as the task of describing the predicate-argument structure of a sentence using natural language only, which is human-interpretable by definition. Our experiments, supported by in-depth analyses, demonstrated that prioritizing interpretability does not come at the expense of performance. Furthermore, our approach is flexible enough to achieve competitive or even state-of-the-art results on popular gold standard benchmarks for SRL, showing that natural language can act as a bridge between heterogeneous linguistic resources, e.g., PropBank and FrameNet, and also annotation formalisms, e.g., dependency- or span-based SRL. We hope that our model will foster research in high-performance yet interpretable systems in NLP, and provide a means towards achieving easier integration of sentence-level semantics into downstream applications.

Limitations

Generation. Although our model achieves results on gold standard benchmarks that are on par or even better than the current state of the art, its generative nature certainly makes it slower than previous work based on discriminative approaches (He et al., 2019; Shi and Lin, 2019; Conia et al., 2021). Indeed, our model generates the entire semantically-augmented sentence, i.e., the input sentence with its predicate-argument structures in natural language, autoregressively. While this issue also affects our most direct competitor (Blloshmi et al., 2021), which generates discrete labels, this is still a limitation – or, more precisely, a weakness – we would like to remark. Indeed, before deploying our system in production environments, one should carefully weigh the advantages of our method against its slower inference times. The degree of slowdown will inevitably depend on the hardware, but we estimate that a generative approach could be several times slower than a discriminative one. However, this could also be a matter for further research on the topic; for example, non-autoregressive generative models are steadily narrowing the performance gap (Gu and Tan, 2022) while mitigating the weaknesses of current autoregressive approaches.

Evaluation. Section 6 and Table 7 provide a qualitative analysis of the behavior of our proposed approach on out-of-inventory instances, which may also include rare predicates or neologisms. We acknowledge that a quantitative analysis of how our model really performs on out-of-inventory instances would provide sounder evidence of the benefits of our approach. However, we do not possess the economic and human resources required to create a benchmark large enough for this purpose. We believe that such a benchmark could be a great contribution to the area of SRL, but the endeavor of annotating a significant number of out-of-inventory instances will require further study.

Multilinguality. Extending our work to multiple languages is still a challenge and may require more effort than current approaches, such as that proposed by Conia et al. (2021) which uses language-specific decoders on top of a shared cross-lingual encoder. One could consider pursuing a similar strategy, i.e., using a shared cross-lingual encoder and multiple language-specific autoregressive decoders. However, the main limitation here is the

availability and the structure of current linguistic inventories in other languages and, therefore, definitions in languages other than English. For instance, the Chinese PropBank inventory provided as part of the CoNLL-2009 Shared Task lacks definitions for the majority of the predicate senses, whereas the latest version is not freely distributed. Fortunately, the attention to multilingual SRL is increasing; for example, it would certainly be interesting to analyze the feasibility of our approach to the recently released global FrameNet project.

Ethics Statement

Pretrained language models have been shown to manifest undesirable biases, inherited from the corpora on which they have been trained using self-supervision strategies. We train our model starting from the weights of BART (Lewis et al., 2020) and, therefore, there is a high probability that these biases are also inherited, or even exaggerated, by our final models. However, we did not investigate such biases in this work; hence, we advise against using our model in a production environment without a careful analysis beforehand. Finally, we remark that the test sets of CoNLL-2009, CoNLL-2012, and FrameNet 1.7 also contain relatively old documents about economics, politics, and past events that do not reflect the current situation. Therefore, the results of such benchmarks are intended only as a basis for comparison with previous approaches and not as a measure of the performance of our model in real-world applications.

Acknowledgments

The authors gratefully acknowledge the support of the ERC Consolidator Grant MOUSSE No. 726487 under the European Union’s Horizon 2020 research and innovation programme.



This work was supported in part by the MIUR under grant “Dipartimenti di Eccellenza 2018-2022” of the Department of Computer Science of Sapienza University of Rome.

References

- Collin Baker, Michael Ellsworth, and Katrin Erk. 2007. *SemEval-2007 task 19: Frame semantic structure extraction*. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*,

- pages 99–104, Prague, Czech Republic. Association for Computational Linguistics.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The Berkeley FrameNet project](#). In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Edoardo Barba, Tommaso Pasini, and Roberto Navigli. 2021a. [ESC: Redesigning WSD with extractive sense comprehension](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4661–4672, Online. Association for Computational Linguistics.
- Edoardo Barba, Luigi Procopio, and Roberto Navigli. 2021b. [ConSeC: Word sense disambiguation as continuous sense comprehension](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1492–1503, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Michele Bevilacqua, Marco Maru, and Roberto Navigli. 2020. [Generatory or “how we went beyond word sense inventories and learned to gloss”](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7207–7221, Online. Association for Computational Linguistics.
- Michele Bevilacqua, Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021. [Recent trends in Word Sense Disambiguation: A survey](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4330–4338. ijcai.org.
- Rexhina Blloshmi, Simone Conia, Rocco Tripodi, and Roberto Navigli. 2021. [Generating senses and roles: An end-to-end model for dependency- and span-based semantic role labeling](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3786–3793. International Joint Conferences on Artificial Intelligence Organization.
- 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, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Xavier Carreras and Lluís Màrquez. 2005. [Introduction to the CoNLL-2005 shared task: Semantic role labeling](#). In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 152–164, Ann Arbor, Michigan. Association for Computational Linguistics.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2011. [An analysis of open information extraction based on semantic role labeling](#). In *Proceedings of the Sixth International Conference on Knowledge Capture, K-CAP ’11*, page 113–120, New York, NY, USA. Association for Computing Machinery.
- Simone Conia, Andrea Bacciu, and Roberto Navigli. 2021. [Unifying cross-lingual semantic role labeling with heterogeneous linguistic resources](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 338–351, Online. Association for Computational Linguistics.
- Simone Conia and Roberto Navigli. 2020. [Bridging the gap in multilingual semantic role labeling: a language-agnostic approach](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1396–1410, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Simone Conia and Roberto Navigli. 2022. [Probing for predicate argument structures in pretrained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4622–4632, Dublin, Ireland. Association for Computational Linguistics.
- Dipanjan Das, Desai Chen, André F. T. Martins, Nathan Schneider, and Noah A. Smith. 2014. [Frame-semantic parsing](#). *Computational Linguistics*, 40(1):9–56.
- Andrea Di Fabio, Simone Conia, and Roberto Navigli. 2019. [VerbAtlas: A novel large-scale verbal semantic resource and its application to Semantic Role Labeling](#). 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)*, pages 627–637, Hong Kong, China. Association for Computational Linguistics.
- Hao Fei, Meishan Zhang, Bobo Li, and Donghong Ji. 2021. [End-to-end semantic role labeling with neural transition-based model](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelli-*

- gence, *EAAI 2021, Virtual Event, February 2-9, 2021*, pages 12803–12811. AAAI Press.
- Nicholas FitzGerald, Julian Michael, Luheng He, and Luke Zettlemoyer. 2018. [Large-scale QA-SRL parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2051–2060, Melbourne, Australia. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Daniel Gildea and Daniel Jurafsky. 2002. [Automatic labeling of semantic roles](#). *Computational Linguistics*, 28(3):245–288.
- Ana-Maria Giuglea and Alessandro Moschitti. 2006. [Semantic role labeling via FrameNet, VerbNet and PropBank](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 929–936, Sydney, Australia. Association for Computational Linguistics.
- Jiatao Gu and Xu Tan. 2022. [Non-autoregressive sequence generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 21–27, Dublin, Ireland. Association for Computational Linguistics.
- Jan Hajič, Massimiliano Ciaramita, Richard Johanson, Daisuke Kawahara, Maria Antònia Martí, Lluís Màrquez, Adam Meyers, Joakim Nivre, Sebastian Padó, Jan Štěpánek, Pavel Straňák, Mihai Surdeanu, Nianwen Xue, and Yi Zhang. 2009. [The CoNLL-2009 shared task: Syntactic and semantic dependencies in multiple languages](#). In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL 2009): Shared Task*, pages 1–18, Boulder, Colorado. Association for Computational Linguistics.
- Patrick Hanks. 2000. [Do word meanings exist?](#) *Comput. Humanit.*, 1-2(34):205–215.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018. [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, Mike Lewis, and Luke Zettlemoyer. 2015. [Question-answer driven semantic role labeling: Using natural language to annotate natural language](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 643–653, Lisbon, Portugal. Association for Computational Linguistics.
- Shexia He, Zuchao Li, and Hai Zhao. 2019. [Syntax-aware multilingual semantic role labeling](#). 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)*, pages 5350–5359, Hong Kong, China. Association for Computational Linguistics.
- Luyao Huang, Chi Sun, Xipeng Qiu, and Xuanjing Huang. 2019. [GlossBERT: BERT for word sense disambiguation with gloss knowledge](#). 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)*, pages 3509–3514, Hong Kong, China. Association for Computational Linguistics.
- Shonosuke Ishiwatari, Hiroaki Hayashi, Naoki Yoshinaga, Graham Neubig, Shoetsu Sato, Masashi Toyoda, and Masaru Kitsuregawa. 2019. [Learning to describe unknown phrases with local and global contexts](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3467–3476, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adam Kilgarriff. 2004. [How dominant is the commonest sense of a word?](#) In *Text, Speech and Dialogue, 7th International Conference, TSD 2004, Brno, Czech Republic, September 8-11, 2004, Proceedings*, volume 3206 of *Lecture Notes in Computer Science*, pages 103–112. Springer.
- Karin Kipper Schuler. 2005. [VerbNet: A broad-coverage, comprehensive verb lexicon](#). University of Pennsylvania.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- 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](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6730–6737. AAAI Press.
- ZhiChao Lin, Yueheng Sun, and Meishan Zhang. 2021. [A graph-based neural model for end-to-end frame semantic parsing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3864–3874, Online and Punta Cana,

- Dominican Republic. Association for Computational Linguistics.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*.
- Maddalen Lopez de Lacalle, Egoitz Laparra, and German Rigau. 2014. **Predicate Matrix: extending Sem-Link through WordNet mappings**. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 903–909, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Diego Marcheggiani, Jasmijn Bastings, and Ivan Titov. 2018. **Exploiting semantics in neural machine translation with graph convolutional networks**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 486–492, New Orleans, Louisiana. Association for Computational Linguistics.
- Diego Marcheggiani, Anton Frolov, and Ivan Titov. 2017. **A simple and accurate syntax-agnostic neural model for dependency-based semantic role labeling**. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 411–420, Vancouver, Canada. Association for Computational Linguistics.
- Diego Marcheggiani and Ivan Titov. 2020. **Graph convolutions over constituent trees for syntax-aware semantic role labeling**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3915–3928, Online. Association for Computational Linguistics.
- Lluís Màrquez, Xavier Carreras, Kenneth C. Litkowski, and Suzanne Stevenson. 2008. **Special issue introduction: Semantic role labeling: An introduction to the special issue**. *Computational Linguistics*, 34(2):145–159.
- Muhidin Mohamed and Mourad Oussalah. 2019. **Srl-esa-textsum: A text summarization approach based on semantic role labeling and explicit semantic analysis**. *Information Processing & Management*, 56(4):1356–1372.
- Roberto Navigli. 2018. **Natural language understanding: Instructions for (present and future) use**. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 5697–5702.
- Ke Ni and William Yang Wang. 2017. **Learning to explain non-standard English words and phrases**. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 413–417, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Thanapon Noraset, Chen Liang, Larry Birnbaum, and Doug Downey. 2017. **Definition modeling: Learning to define word embeddings in natural language**. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3259–3266. AAAI Press.
- Martha Palmer. 2009. **SemLink: Linking PropBank, VerbNet and FrameNet**. In *Proceedings of the Generative Lexicon Conference*, pages 9–15. GenLex-09, Pisa, Italy.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. **The Proposition Bank: An annotated corpus of semantic roles**. *Computational Linguistics*, 31(1):71–106.
- Ayush Pancholy, Miriam R L Petruck, and Swabha Swayamdipta. 2021. **Sister help: Data augmentation for frame-semantic role labeling**. In *Proceedings of The Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 78–84, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. **Structured prediction as translation between augmented natural languages**. In *9th International Conference on Learning Representations, ICLR 2021*.
- Hao Peng, Sam Thomson, Swabha Swayamdipta, and Noah A. Smith. 2018. **Learning joint semantic parsers from disjoint data**. 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 1492–1502, New Orleans, Louisiana. Association for Computational Linguistics.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. **WiC: the word-in-context dataset for evaluating context-sensitive meaning representations**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. **CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes**. In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Valentina Pyatkin, Paul Roit, Julian Michael, Yoav Goldberg, Reut Tsarfaty, and Ido Dagan. 2021. **Asking it all: Generating contextualized questions for any**

- semantic role**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1429–1441, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Michael Roth and Mirella Lapata. 2016. **Neural semantic role labeling with dependency path embeddings**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1192–1202, Berlin, Germany. Association for Computational Linguistics.
- Arka Sadhu, Tanmay Gupta, Mark Yatskar, Ram Nevatia, and Aniruddha Kembhavi. 2021. Visual semantic role labeling for video understanding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Dan Shen and Mirella Lapata. 2007. **Using semantic roles to improve question answering**. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 12–21, Prague, Czech Republic. Association for Computational Linguistics.
- Peng Shi and Jimmy Lin. 2019. **Simple BERT models for relation extraction and semantic role labeling**. *arXiv preprint arXiv:1904.05255*.
- Kevin Stowe, Jenette Preciado, Kathryn Conger, Susan Windisch Brown, Ghazaleh Kazeminejad, and Martha Palmer. 2021. **SemLink 2.0: Chasing lexical resources**. In *30th Annual Meeting of the Association for Computational Linguistics*, pages 222–227, Gronigen, Netherlands. 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, Brussels, Belgium. Association for Computational Linguistics.
- 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 *CoNLL 2008: Proceedings of the Twelfth Conference on Computational Natural Language Learning*, pages 159–177, Manchester, England. Coling 2008 Organizing Committee.
- Swabha Swayamdipta, Sam Thomson, Chris Dyer, and Noah A. Smith. 2017. Frame-Semantic Parsing with Softmax-Margin Segmental RNNs and a Syntactic Scaffold. *arXiv preprint arXiv:1706.09528*.
- David Tuggy. 1993. **Ambiguity, polysemy, and vagueness**. *Cognitive Linguistics*, 4(3):273–290.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li Zhang, Ishan Jindal, and Yunyao Li. 2022. **Label definitions improve semantic role labeling**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5613–5620, Seattle, United States. Association for Computational Linguistics.

A Data License

Both the CoNLL-2009 and CoNLL-2012 datasets are distributed by the Linguistic Data Consortium (LDC) and can be used under the LDC license.⁹ FrameNet 1.7 – the linguistic resource and its annotated dataset – is freely available upon request.¹⁰ We note that the original Shared Task of CoNLL-2012 was concerned with the task of Coreference Resolution; however, given its SRL annotations, it soon also became a popular benchmark for span-based SRL.

B Data Statistics

In Tables 8, 9, and 10, we provide an overview of the statistics of the train, validation and test sets, respectively, for the datasets we use in our experiments, namely, the English splits of CoNLL-2009, CoNLL-2012, and FrameNet 1.7. In particular, for each dataset, we report the number of sentences and their average length in tokens, with FrameNet having the longest sentences on average (+20% over CoNLL-2009 and +40% over CoNLL-2012). We also report the number of annotated predicates for each dataset; interestingly, each predicate in FrameNet features around 6 arguments per predicate, a value that is much larger than those of CoNLL-2009 and CoNLL-2012, which feature around 2.5 arguments per predicate. These are probably the reasons why the FrameNet dataset is particularly challenging, even for modern neural-based models.

Finally, we can also appreciate the heterogeneity between the characteristics of PropBank-style and FrameNet-style SRL. Indeed, FrameNet clusters predicate senses into frames, resulting in a smaller number of predicate classes (around 1,000) compared to PropBank (5,000 to 8,000). At the same time, the frame-specific semantic roles of FrameNet result in a much larger number of role classes compared to the coarse-grained semantic roles of PropBank.

C Training Sequence Statistics

In Table 11, we report the average length in characters of the sequences used to train our model. As we can see, FrameNet 1.7 features the longest sequences among the three datasets we take into account, in line with what we report in Appendix B.

⁹<https://www.ldc.upenn.edu/data-management/using/licensing>

¹⁰<https://framenet.icsi.berkeley.edu/fndrupal>

D Argument Modifiers Definitions

The English PropBank features two categories of semantic roles: core and adjunct. If we define a semantic role as the relationship between an action or event (predicate) and one of the participants (argument), then the former category includes all those semantic roles that mark an important participant in the event, one that is expected to take part in it. In PropBank, these core roles are identified using the labels ARG0, ARG1, . . . , ARG5, and their definitions change from predicate sense to predicate sense. Instead, the second category, namely the adjunct roles or argument modifiers, are general roles whose semantics is not specific to a particular predicate and, therefore, can be used to tag general arguments, e.g., the time of the action (ARGM-TMP) or the place of the event (ARGM-LOC). We use the PropBank guidelines to translate such labels into natural language. In Tables 12 and 13, we list the argument modifiers definitions that we use to train our model on CoNLL-2009 and CoNLL-2012, respectively.

While we aimed at creating argument modifier definitions that are homogeneous with the core role definitions, we remark that we did not perform a search for better definitions. As one can see, some of the definitions reported in Tables 12 and 13 are the natural language equivalent of the labels (e.g. ARGM-ADV and its definition “adverbial modifier”, ARGM-LVB and its definition “light verb”, or ARGM-PRD and its definition “secondary predication”, among others). We believe that a possible venue for future research is looking into how we can create better definitions for such semantic roles.

E Definitions Statistics

The length of the sequence that our model generates in output is certainly dependent on the length of the definitions we use to describe the sense of a predicate and its arguments. In this Appendix, we provide a broad look at the number of unique sense and role definitions that appear in the train, validation, and test sets of CoNLL-2009, CoNLL-2012 and FrameNet 1.7.

As we can see in Table 14, even though CoNLL-2009 and CoNLL-2012 are both tagged using PropBank labels, the number of distinct predicate sense definitions is quite different between the two datasets (1,317 unique definitions in the training set of CoNLL-2009 against 4,401 in CoNLL-2012). This difference is probably due to the narrower

	Sentences				Predicates		Arguments	
	Total _s	Distinct _s	Annotated	Avg. Len.	Total _p	Senses	Total _a	Roles
CoNLL-2009	39,279	38,770	37,847	24.4	179,014	8,237	393,699	52
CoNLL-2012	115,812	109,374	90,856	19.0	253,070	5,287	598,983	66
FrameNet	19,391	3,353	19,391	29.5	20,105	859	123,977	2,042

Table 8: Overview of the CoNLL-2009, CoNLL-2012, and FrameNet training datasets. For each dataset we report the number of sentences ($Total_s$), the number of sentences with at least an annotated predicate ($Annotated$), the average number of tokens per sentence ($Avg. Len.$), the number of predicates ($Total_p$) and predicate senses ($Senses$), and also the number of arguments ($Total_a$) and argument roles ($Roles$).

	Sentences				Predicates		Arguments	
	Total _s	Distinct _s	Annotated	Avg. Len.	Total _p	Senses	Total _a	Roles
CoNLL-2009	1,334	1,334	1,283	25.0	6,390	1,990	13,865	32
CoNLL-2012	15,680	15,086	12,600	19.4	35,297	2,912	83,362	48
FrameNet	2,272	326	2,272	35.2	2,382	394	17,347	893

Table 9: Overview of the CoNLL-2009, CoNLL-2012, and FrameNet validation datasets. For each dataset we report the number of sentences ($Total_s$), the number of sentences with at least an annotated predicate ($Annotated$), the average number of tokens per sentence ($Avg. Len.$), the number of predicates ($Total_p$) and predicate senses ($Senses$), and also the number of arguments ($Total_a$) and argument roles ($Roles$).

	Sentences				Predicates		Arguments	
	Total _s	Distinct _s	Annotated	Avg. Len.	Total _p	Senses	Total _a	Roles
CoNLL-2009	2,000	1,999	1,913	24.4	8,987	2,254	19,949	35
CoNLL-2012	27,897	26,698	21,863	19.2	62,012	3,489	145,078	50
FrameNet	6,714	1,247	6,714	27.2	6,872	620	34,454	1,354

Table 10: Overview of the CoNLL-2009, CoNLL-2012, and FrameNet test datasets. For each dataset we report the number of sentences ($Total_s$), the number of sentences with at least an annotated predicate ($Annotated$), the average number of tokens per sentence ($Avg. Len.$), the number of predicates ($Total_p$) and predicate senses ($Senses$), and also the number of arguments ($Total_a$) and argument roles ($Roles$).

Dataset	Avg. Len. (characters)
CoNLL-2009	83.1
CoNLL-2012	127.6
FrameNet 1.7	205.3

Table 11: CoNLL-2009, CoNLL-2012, and FrameNet training sequence statistics. For each dataset, we report the average length in characters of the sequence used for training the model.

domain of CoNLL-2009, which features a significant portion of sentences about finance from the Wall Street Journal, whereas CoNLL-2012 covers a more varied set of domains. Although the number

of unique sense definitions is different, the average length of these definitions between CoNLL-2009 and CoNLL-2012 is close, suggesting homogeneous definitions despite the use of two different versions of the English PropBank. This is not the case when comparing the average length of the PropBank definitions used for CoNLL-2009 and CoNLL-2012 with those of FrameNet. Indeed, predicate sense definitions in FrameNet are two to three times longer on average than PropBank’s. However, the experimental results reported in Tables 3 and 6 show that our proposed generative model is still able to produce longer sense definitions.

We can observe a similar picture in Table 15 for the definitions of the semantic roles. Interest-

Argument Modifier	Definition
AM-ADV	adverbial modifier
AM-CAU	cause or reason
AM-DIR	direction or source
AM-DIS	discourse connective
AM-EXT	amount or extent
AM-LOC	location or position
AM-MNR	instrument or manner
AM-MOD	modal auxiliary
AM-NEG	negation marker
AM-PNC	purpose, not cause
AM-PRD	secondary predication
AM-TMP	time or duration

Table 12: CoNLL-2009 argument modifiers definitions. We provide descriptions for argument modifiers when they are not specified in the given predicate role set.

ingly, the difference between CoNLL-2009 and CoNLL-2012 in the average length of the semantic role definitions is even narrower, whereas the difference in length between PropBank-style and FrameNet-style role definitions widens even further, with FrameNet using role definitions that are almost four times longer than PropBank’s. The difference in length between the predicate sense and semantic role definitions between FrameNet and PropBank can be explained by the fact that, in the former resource, the definitions are richer and more detailed. For example, the agent of the predicate *provide* is defined just as “giver” in PropBank, whereas in FrameNet is defined as “person that begins in possession of the theme and causes it to be in the possession of the recipient”.

F Special Tokens

As mentioned in Section 3.2, we use some special tokens to instruct the model on some task-specific functions. For example, we pre-identify a predicate in an input sentence by surrounding its tokens with the special tokens `<p>` and `</p>`, indicating the start and the end of a predicate, respectively. Table 16 lists all the special tokens we use in our model in addition to the standard ones (e.g., `<s>` and `</s>` to indicate the start and end of the generated sequence).

We note that some of these special tokens can be used in combination. For example, combining `<propbank<span-sr1>` informs the model that we want it to generate a sentence anno-

Argument Modifier	Definition
ARGM-ADJ	adjectival modifier
ARGM-ADV	adverbial modifier
ARGM-CAU	cause or reason
ARGM-COM	comitative
ARGM-DIR	direction or source
ARGM-DIS	discourse connective
ARGM-EXT	amount or extent
ARGM-GOL	goal or destination
ARGM-LOC	location or position
ARGM-LVB	light verb
ARGM-MNR	instrument or manner
ARGM-MOD	modal auxiliary
ARGM-NEG	negation marker
ARGM-PNC	purpose, not cause
ARGM-PRD	secondary predication
ARGM-PRP	purpose or motivation
ARGM-TMP	time or duration

Table 13: CoNLL-2012 argument modifiers definitions. We provide descriptions for argument modifiers when they are not specified in the given predicate role set.

tated with PropBank-style definitions according to the span-based formalism; instead, combining `<framenet<span-sr1>` will result in a sentence annotated with FrameNet-style definitions using a span-based formalism.

For reference, we also provide a few examples of how these special tokens are inserted in an input or output sequence in Table 17, using sentences from the training set of CoNLL-2012.

For the implementation, we simply add these special tokens to the input and output vocabulary of the underlying language model (i.e., BART). The embeddings corresponding to the special tokens are randomly initialized and updated during training.

	Train		Validation		Test	
	Distinct _d	Avg. Len. _d	Distinct _d	Avg. Len. _d	Distinct _d	Avg. Len. _d
CoNLL-2009	1,317	16.0	1,207	16.5	1,317	16.0
CoNLL-2012	4,401	19.5	2,393	18.4	2,864	18.7
FrameNet	3,750	46.7	882	47.3	1,982	48.1

Table 14: CoNLL-2009, CoNLL-2012, and FrameNet predicate definitions statistics. For each dataset and split we report the number of distinct definitions (*Distinct_d*), and their average length in characters (*Avg. Len._d*).

	Train		Validation		Test	
	Distinct _d	Avg. Len. _d	Distinct _d	Avg. Len. _d	Distinct _d	Avg. Len. _d
CoNLL-2009	1,255	15.6	1,161	15.4	1,255	15.6
CoNLL-2012	5,002	16.9	2,477	16.4	3,032	16.4
FrameNet	2,167	58.7	634	58.2	1,184	57.0

Table 15: CoNLL-2009, CoNLL-2012, and FrameNet role definitions statistics. For each dataset and split we report the number of distinct definitions (*Distinct_d*), and their average length in characters (*Avg. Len._d*).

Used in	Special Token(s)	Description
Input	<p></p>	indicate the start/end of a predicate
Output	<reference-to>	argument referring to another one (e.g., R-Arg1)
Output	<continuation-of>	continuation of another argument (e.g., C-Arg1)
Output	<propbank>	Instructs the model to perform PropBank-style SRL
Output	<framenet>	Instructs the model to perform FrameNet-style SRL
Output	<span-srl>	Instructs the model to perform span-based SRL
Output	<dep-srl>	Instructs the model to perform dependency-based SRL

Table 16: List of the special tokens and their use in our model. For each special token, we indicate whether it is used in the input or in the output sequence. Some of these special tokens can be used in combination, e.g., <propbank><dep-srl> to instruct the model to perform PropBank-style dependency-based SRL.

Special Token	Example
<p></p>	Not all those who <p> wrote </p> oppose the changes.
<p></p>	A <p> record </p> date has not been set.
<reference-to>	...It was [during this year]{time or duration} [that] { <reference-to> time or duration } [the Japanese...
<continuation-of>	[Japan]{helper}, [in terms of ...]{ adverbial modifier } , [it]{ <continuation-of> helper } could have helped...

Table 17: Examples of how the special tokens are inserted in the input or output sequence.