

InVeRo: Making Semantic Role Labeling Accessible with Intelligible Verbs and Roles

Simone Conia¹, Fabrizio Brignone², Davide Zanfardino², and Roberto Navigli¹

¹Sapienza NLP Group, Department of Computer Science, Sapienza University of Rome

²Babelscape, Italy

lastname@di.uniroma1.it, lastname@babelscape.com

Abstract

Semantic Role Labeling (SRL) is deeply dependent on complex linguistic resources and sophisticated neural models, which makes the task difficult to approach for non-experts. To address this issue we present a new platform named Intelligible Verbs and Roles (InVeRo). This platform provides access to a new verb resource, VerbAtlas, and a state-of-the-art pre-trained implementation of a neural, span-based architecture for SRL. Both the resource and the system provide human-readable verb sense and semantic role information, with an easy to use Web interface and RESTful APIs available at <http://nlp.uniroma1.it/invero>.

1 Introduction

Since its introduction (Gildea and Jurafsky, 2002), Semantic Role Labeling (SRL) has been recognized as a key task to enable Natural Language Understanding in that it aims at explicitly answering the "Who did What to Whom, When and Where?" question by identifying and labeling the predicate-argument structure of a sentence, namely, the actors that take part in the scenario outlined by a predicate. In fact, SRL has already proven to be useful in a wide range of downstream tasks, including Question Answering (Shen and Lapata, 2007; He et al., 2015), Information Extraction (Christensen et al., 2011), Situation Recognition (Yatskar et al., 2016), Machine Translation (Marcheggiani et al., 2018), and Opinion Role Labeling (Zhang et al., 2019).

Unfortunately, the integration of SRL knowledge into downstream applications has often been hampered and slowed down by the intrinsic complexity of the task itself (Navigli, 2018). Indeed, SRL is strongly intertwined with elaborate linguistic theories, as identifying and labeling predicate-argument relations requires well-defined predicate sense and semantic role inventories such as the popular PropBank (Palmer et al., 2005), VerbNet

(Kipper-Schuler, 2005), or FrameNet (Baker et al., 1998). The linguistic intricacies of such resources may, however, dishearten and turn away new practitioners. Regardless of which linguistic resource is used in the task, to further complicate the situation SRL has been usually divided into four subtasks – predicate identification, predicate sense disambiguation, argument identification and argument classification – but, to the best of our knowledge, recent state-of-the-art systems do not address all these four subtasks simultaneously without relying on external systems (Swayamdipta et al., 2017; He et al., 2018; Strubell et al., 2018; He et al., 2019). Therefore, obtaining predicate sense and semantic role annotations necessitates the tedious orchestration of multiple automatic systems, which in its turn further complicates the use of SRL in practice and in semantics-first approaches to NLP more generally.

In this paper, we present InVeRo (Intelligible Verbs and Roles), an online platform designed to tackle the aforementioned issues and make Semantic Role Labeling accessible to a broad audience. InVeRo brings together resources and tools to perform human-readable SRL, and it accomplishes this by using the intelligible verb senses and semantic roles of a recently proposed resource named VerbAtlas (Di Fabio et al., 2019) and exploiting them to annotate sentences with high performance. In more detail, the InVeRo platform includes:

- a Resource API to obtain linguistic information about the verb senses and semantic roles in VerbAtlas.
- a Model API to effortlessly annotate sentences using a state-of-the-art end-to-end pretrained model for span-based SRL.
- a Web interface where users can easily query linguistic information and automatically an-

notate sentences on-the-go without having to write a single line of code.

Notably, InVeRo also takes advantage of PropBank to get the best of both worlds, and provides annotations according to both resources, enabling comparability and fostering integration.

2 The InVeRo Platform

The InVeRo platform aims at making SRL more approachable to a wider audience, not only in order to promote advances in the area of SRL itself, but also to encourage the integration of semantics into other fields of NLP. The two main barriers to this objective are the complexity of i) the linguistic resources used in SRL which are, however, indispensable for the definition of the task itself, and ii) the complexity of the recently proposed techniques.

Section 2.1 explains how InVeRo takes advantage of the intelligible verb senses and semantic roles of VerbAtlas to gently introduce non-expert users to SRL, while Section 2.2 details how the InVeRo model for SRL can make semantic role annotations accessible to everyone.

2.1 Intelligible Verb Senses and Roles

One of the most contentious points of discussion in SRL is how to formalize predicate-argument structures, that is, the semantic roles that actors can play in a scenario defined by a predicate. PropBank (Palmer et al., 2005), one of the most popular predicate-argument structure inventories, uses an enumerative approach where each predicate sense has a possibly different roleset, e.g., for the predicate *make*, the sense *make.01* (as in “making a product”) bears the semantic roles ARG0 (creator), ARG1 (creation), ARG2 (created from) and ARG3 (beneficiary), whereas *make.02* (as in “cause to be”) bears only ARG0 (impeller) and ARG1 (impelled). This exhaustive approach, however, requires an expert linguist to tell which roles share similar semantics across senses (e.g., ARG0 is an agent in both *make.01* and *make.02*) and which do not (e.g., ARG1 is a product in *make.01* but a result in *make.02*).

On the other hand, VerbAtlas (Di Fabio et al., 2019), a recently proposed predicate-argument structure inventory, in contrast to the enumerative approach of PropBank and the thousands of frame-specific roles of FrameNet, adopts a small set of explicit and intelligible semantic roles (AGENT, PRODUCT, RESULT, DESTINATION, . . . , THEME) inspired by VerbNet (Kipper-Schuler, 2005). As a result, in

VerbAtlas, whenever two predicate senses can bear the same semantic role, the semantics of this role is coherent across the two predicate senses by definition, resulting in readable labels for non-expert users. VerbAtlas also clusters predicate senses into so-called frames (COOK, DRINK, HIT, etc.) inspired by FrameNet (Baker et al., 1998), with the idea that senses sharing similar semantic behavior lie in the same frame. For non-expert users, this organization has the added advantage of explicitly linking predicate senses that are otherwise unrelated, like *make.01* and *create.01* in PropBank which, instead, are part of the same frame MOUNT-ASSEMBLE-PRODUCE in VerbAtlas and, therefore, also share the same semantic roles. In a bid to make SRL more accessible, the InVeRo platform adopts the intelligible verb senses and semantic roles of VerbAtlas.

2.2 An All-in-One Solution for SRL

As already mentioned in Section 1, the traditional SRL pipeline consists of four main steps: predicate identification, predicate sense disambiguation, argument identification and argument classification. While some of the above steps are considered easier than others, each of them features distinct peculiarities, which has driven recent works to focus on improving only specific aspects of the entire SRL pipeline. Instead, little attention has been paid to systems that can tackle all the above-mentioned steps at the same time. As a result, anyone wishing to take advantage of SRL annotations in another NLP task has to choose, mix and match multiple automatic systems in order to obtain sentences fully annotated with predicate sense and semantic role labels. Understandably, this has been a major deterrent for the integration of semantics into downstream applications.

As part of the InVeRo platform, not only do we introduce an all-in-one model that addresses the complete SRL pipeline with a single forward pass, but we also make this model available through a Web interface to let everyone label sentences with SRL annotations without the need to install any software. In other words, a user only has to provide a raw text sentence; the InVeRo all-in-one model for SRL takes care of the rest, making the predicate sense and role labeling process accessible and effortless.

Model Design. The InVeRo all-in-one system for SRL is based on the ideas put forward by He et al. (2018) in that, unlike other works that used word-

level BIO tagging schemes to label arguments (He et al., 2017; Strubell et al., 2018; Tan et al., 2018), it directly models span-level features. In particular, we follow He et al. (2018) by letting the neural model learn span-level representations from the word-level representations of the span start and span end words, while also adding a span-length specific trainable embedding. More formally, the span representation \mathbf{s}_{ij} from word i to word j is obtained as follows:

$$\mathbf{s}_{ij} = \mathbf{W}^s (\mathbf{e}_i^w \oplus \mathbf{e}_j^w \oplus \mathbf{e}_{j-i}^l) + \mathbf{b}^s$$

where \mathbf{e}_i^w and \mathbf{e}_j^w are the word representations of start and end of the span, \mathbf{e}_{j-i}^l is the span length embedding, and \oplus is the concatenation operation.

However, our approach features a few key differences that set the InVeRo model apart from the aforementioned works. First, it creates contextualized word representations from the inner states of BERT (bert-base-cased), a recent language model trained on massive amounts of textual data (Devlin et al., 2018). Differently from the recent work of Shi and Lin (2019), our model takes advantage of the topmost four layers of BERT and directly builds a word representation from its subword representations, similarly to Bevilacqua and Navigli (2020). More formally, given the BERT representations \mathbf{h}_{ij}^k at layer k of the m_i subwords w_{ij} in word w_i , with $1 \leq j \leq m_i$:

$$\begin{aligned} \mathbf{c}_{ij} &= \mathbf{h}_{ij}^{-1} \oplus \mathbf{h}_{ij}^{-2} \oplus \mathbf{h}_{ij}^{-3} \oplus \mathbf{h}_{ij}^{-4} \\ \mathbf{c}'_{ij} &= \text{ReLU}(\mathbf{W}^c \mathbf{c}_{ij} + \mathbf{b}^c) \\ \mathbf{e}_i^w &= \frac{1}{m_i} \sum_j \mathbf{c}'_{ij} \end{aligned}$$

Second, in contrast to other span-based SRL systems, our model integrates predicate disambiguation as an additional objective in a multitask fashion (Caruana, 1997). Third, our model is trained to jointly learn to label sentences with both VerbAtlas and PropBank so as to exploit the complementary knowledge of the two resources, and, at the same time, provide a means to directly compare the predicate sense and semantic role labels of two different inventories for the same input sentences.¹

Comparison with previous systems. Over the years, several SRL systems have been developed

¹We used the PropBank-to-VerbAtlas mappings available at <http://verbatlas.org/download> to remap CoNLL-2012.

and made available as prepackaged downloads, e.g. SENNA², or as online demos, e.g., AllenNLP’s SRL demo³. However, recent BERT-based online systems, such as AllenNLP’s SRL demo, do not perform predicate sense disambiguation (in addition to predicate identification, argument identification and argument classification), which is a crucial step in SRL, especially when considering that the PropBank roles ARG0, ARG1, through ARG5 become meaningful only if they are associated with a PropBank predicate sense (see Section 2.1).

Results. Thanks to the use of contextualized word representations from BERT, the joint exploitation of two complementary linguistic resources for SRL, and the introduction of a predicate sense disambiguation layer, our model achieves 84.0% in F₁ score in the standard argument identification and classification test split of the CoNLL-2012 dataset (Pradhan et al., 2012), significantly outperforming the previous state of the art among end-to-end models, currently represented by Strubell et al. (2018) with a 0.6% absolute improvement in F₁ score⁴ (84.0% against 83.4%). We note that this measure does not take into account the performance on predicate sense disambiguation, where our system achieves 86.1% in F₁ score, which is a significant absolute improvement (+5.7%) over the most-frequent-sense strategy (86.1% against 80.4%).

3 The InVeRo APIs

To foster the integration of semantics into a wider range of applications, the InVeRo platform introduces a set of RESTful APIs⁵ that offer i) easy-to-use abstractions to query resource-specific information in VerbAtlas (Section 3.1), and ii) out-of-the-box predicate and semantic role annotations from a state-of-the-art pretrained model (Section 3.2).

3.1 Resource API

The Resource API provides a RESTful interface to easily link predicate-level information, e.g., predicate lemmas and/or predicate senses, to VerbAtlas-specific features, e.g., frames and semantic roles. In particular:

²<https://ronan.collobert.com/senna>

³<https://demo.allennlp.org/semantic-role-labeling>

⁴Score computed with the official CoNLL-2005 script.

⁵<http://nlp.uniroma1.it/invero/api-documentation>

- the `/predicate` endpoint exposes functionalities to obtain frame-level information starting from a predicate lemma or a synset from WordNet 3.0 (Fellbaum et al., 1998) or BabelNet 4.0 (Navigli and Ponzetto, 2012);
- the `/frame` endpoint exposes functionalities to retrieve, for a given frame, its Predicate Argument Structure, and the WordNet/BabelNet synsets belonging to this frame.

Also included is a manually-curated PropBank-to-VerbAtlas alignment to remap existing corpora like the CoNLL-2009 and CoNLL-2012 datasets. In particular:

- the `/align/sense` endpoint returns, for a given PropBank predicate sense, its corresponding VerbAtlas frame, i.e., the VerbAtlas frame that generalizes the given PropBank predicate sense;
- the `/align/roles` endpoint returns, for a given PropBank predicate sense, e.g., *aim.01*, the alignment of each role in the PropBank argument structure of the given predicate sense to a VerbAtlas role, e.g., ARG0 → AGENT, ARG1 → THEME, and so on.

The online documentation provides an overview of the accepted parameters at the endpoints available in the Resource API.

3.2 Model API

To encourage the integration of SRL into downstream applications, the Model API offers a simple solution for out-of-the-box role labeling by providing an interface to a full end-to-end state-of-the-art pretrained model. Unlike most currently available models which focus on specific aspects of the entire SRL task, our solution jointly addresses in a single forward pass the whole traditional SRL pipeline, namely, i) predicate identification, ii) predicate sense disambiguation, iii) argument identification, and iv) argument classification. Furthermore, our model is fully self-contained as it does not require any of the additional linguistic information, from lemmatization to part-of-speech tags and syntactic parse trees, that are usually exploited by many systems. Our Model API is:

- **Easy to use:** an end user avoids the struggle of mixing and matching a set of automatic systems where each system independently addresses a different part of the SRL pipeline;

- **Fully self-contained:** the only input to the underlying model is a raw text sentence, dropping any dependency on external preprocessing tools;
- **State-of-the-Art:** the underlying model carries out SRL with high performances on the standard CoNLL-2012 benchmark dataset.

Usage. The Model API exposes a single endpoint named `/model/` which accepts GET requests with a single parameter named `sentence` containing the raw text sentence to label with semantic role annotations. The Model API returns a JSON response that contains, for each predicate it identifies in the sentence, the semantic role that each argument plays with respect to the identified predicate. For example, the response for the sentence “*Eliminating the income tax will benefit peasants*” contains:

```
[{
  "tokenIndex": 0,
  ...
}, {
  "tokenIndex": 5,
  "verbatim": {
    "frameName": "HELP_HEAL_CARE_CURE",
    "roles": [
      {"role": "agent", "score": 0.89, "span": [0, 4]},
      {"role": "beneficiary", "score": 1., "span": [6, 7]}
    ]
  },
  "propbank": {
    "frameName": "benefit.01",
    "roles": [
      {"role": "ARG0", "score": 0.97, "span": [0, 4]},
      {"role": "ARG1", "score": 1.00, "span": [6, 7]}
    ]
  }
},
}]
```

Our Model API also supports the more popular PropBank predicate sense and semantic role labels so as to provide a direct comparison with VerbAtlas and promote synergistic approaches that exploit both inventories to advance SRL.

4 The InVeRo User Interface

Like many other linguistic resources in SRL, VerbAtlas may be daunting for inexperienced practitioners who may still face difficulties in finding their way with the formalisms defined in a linguistic resource for SRL. On top of the previously described APIs (Section 3) and in an effort to make VerbAtlas easier to interact with, the InVeRo platform includes a public-facing Web interface that provides a user-friendly environment to explore not only the functionalities offered by the resource, but

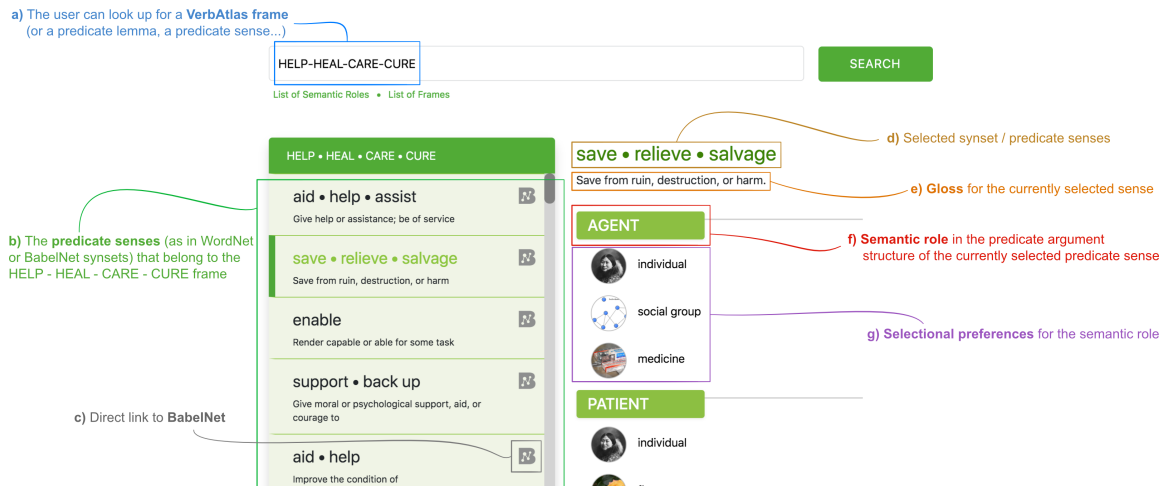


Figure 1: A look at the online interface when a user searches for resource-specific information about VerbAtlas. The user can a) search for a frame name, as in the Figure, or an individual predicate. The interface displays b) all the predicates belonging to the same frame, with each predicate c) directly linked to BabelNet. The right side displays the d) selected predicate with e) its WordNet gloss, f) the semantic roles of its predicate-argument structure, and g) the selectional preferences of each role.



Figure 2: A look at the online interface when a user inserts a sentence in the search bar. The system uses a pretrained model to display all the information of all the steps of a traditional SRL pipeline: predicate identification, predicate sense disambiguation, argument identification and argument classification.

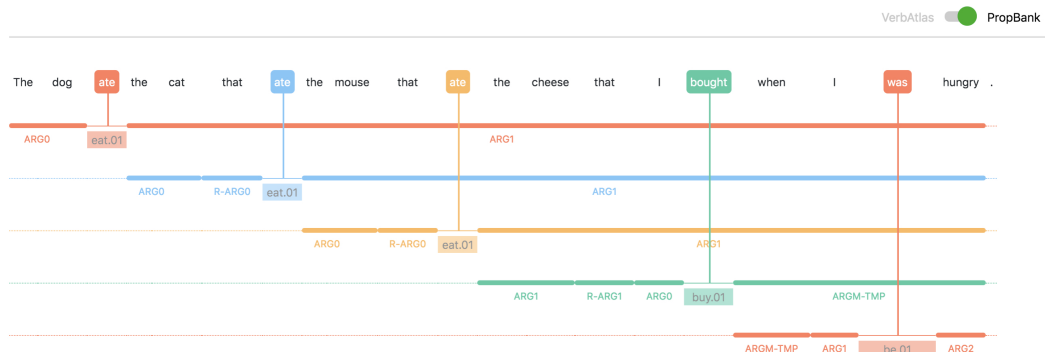


Figure 3: The interface can seamlessly switch between VerbAtlas and PropBank labels with a single click (the switch button at the top-right). Here we show the same sentence as in Figure 2 but labeled with PropBank predicates and roles, which enables comparison across the two annotation styles.

also to understand visually how an SRL system annotates a sentence in a live interactive demo. The Web interface mirrors the functionalities of both the Resource API and the Model API in a minimal unified view, letting users perform resource-specific queries or annotate sentences wherever they are without writing a single line of code.

Resource interface. Figure 1 shows the Web interface when a user inserts the name of a VerbAtlas frame in the search bar. Notice that, since the interface makes use of the Resource API, a user can also search for other resource-specific information such as individual predicates. Particular attention has been given to the visualization of a VerbAtlas frame (Figure 1, left side) which displays all the predicate senses that share similar semantic behavior. Each predicate sense is also conveniently linked to BabelNet 4.0 (Navigli and Ponzetto, 2012), a multilingual knowledge graph where users can find more information such as hypernyms, hyponyms, and semantically related concepts. Equally important is the visualization of a VerbAtlas predicate-argument structure (Figure 1, right side) which displays all the semantic roles that the currently selected predicate/frame can bear in a sentence.

Model interface. Figure 2 shows, instead, the online model interface when a user inserts a sentence with its corresponding predicate sense and semantic role labels from VerbAtlas. Notice how the user can quickly switch between the VerbAtlas and the PropBank predicate sense and semantic role annotations with just a single click, so that the two annotation styles can easily be compared one with the other (Figures 2 and 3). To the best of our knowledge, this is the first online demo where a neural model helps users visualize all the four steps of the traditional SRL pipeline for two different linguistic resources for SRL, VerbAtlas and PropBank, at the same time.

5 Conclusion and Future Work


Semantic Role Labeling is deeply dependent on complex linguistic resources and elaborate neural models: the combination of these two factors has made Semantic Role Labeling (SRL) difficult to approach for experts from other fields who are interested in exploring its integration into downstream applications. In this paper, we aim at ameliorating both of the issues by presenting the InVeRo platform. InVeRo features easy-to-use RESTful

APIs to effortlessly query VerbAtlas, a recently introduced linguistic resource for SRL, and to transparently use a pretrained state-of-the-art end-to-end system for the recent VerbAtlas-style and the more traditional PropBank-style approaches to SRL. Notably, the InVeRo system is fully self-contained as it tackles all the steps of the traditional SRL pipeline – predicate identification, predicate sense disambiguation, argument identification, and argument classification – and it does not require external tools such as lemmatizers, part-of-speech taggers or syntactic tree parsers: users just have to provide a raw text sentence to obtain its corresponding predicate and argument labels. Moreover, the InVeRo platform includes an online Web interface which repackages the APIs in a user-friendly environment. Thanks to this interface, users can easily obtain human-readable linguistic information about VerbAtlas, but also annotate entire sentences on-the-go without the need to install any software.

InVeRo is a growing platform: in the future, we plan to enhance our Model API by adding, alongside the already available state-of-the-art span-based model, the state-of-the-art dependency-based model of Conia and Navigli (2020a), so that users can easily switch between the two approaches and choose the one that best suits their needs. Thanks to BabelNet and recent advances in cross-lingual techniques for tasks where semantics is crucial (Barba et al., 2020; Blloshmi et al., 2020; Conia and Navigli, 2020b; Pasini, 2020; Scarlini et al., 2020), we also plan to provide support for multiple languages to enable SRL integration into multilingual and cross-lingual settings. We believe that the InVeRo platform can make SRL more accessible to the research community, and we look forward to the development of semantics-first approaches in an ever wider range of NLP 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.

References

- 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, COLING-ACL '98, August 10-14, 1998, Université de Montréal, Montréal, Quebec, Canada. Proceedings of the Conference*, pages 86–90.
- Edoardo Barba, Luigi Procopio, Niccolò Campolungo, Tommaso Pasini, and Roberto Navigli. 2020. **MuLaN: Multilingual Label propagation for word sense disambiguation**. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3837–3844.
- Michele Bevilacqua and Roberto Navigli. 2020. **Breaking through the 80% glass ceiling: Raising the state of the art in Word Sense Disambiguation by incorporating knowledge graph information**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2854–2864, Online.
- Rexhina Blloshmi, Rocco Tripodi, and Roberto Navigli. 2020. **XL-AMR: Enabling Cross-Lingual AMR parsing with transfer learning techniques**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020*.
- Rich Caruana. 1997. **Multitask learning**. *Mach. Learn.*, 28(1):41–75.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2011. **An analysis of open information extraction based on semantic role labeling**. In *Proceedings of the 6th International Conference on Knowledge Capture (K-CAP 2011), June 26-29, 2011, Banff, Alberta, Canada*, pages 113–120.
- Simone Conia and Roberto Navigli. 2020a. **Bridging the gap in multilingual Semantic Role Labeling: A language agnostic approach**. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020*.
- Simone Conia and Roberto Navigli. 2020b. **Conception: Multilingually-enhanced, human-readable concept vector representations**. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. **BERT: pre-training of deep bidirectional transformers for language understanding**. *arXiv preprint arXiv:1810.04805*.
- 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 2019, Hong Kong, China, November 3-7, 2019*, pages 627–637.
- Christiane Fellbaum et al. 1998. **WordNet: An electronic database**. MIT Press, Cambridge, MA.
- Daniel Gildea and Daniel Jurafsky. 2002. **Automatic labeling of semantic roles**. *Comput. Linguistics*, 28(3):245–288.
- 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 364–369.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. **Deep semantic role labeling: What works and what's next**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 473–483.
- 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, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 643–653.
- 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 2019, Hong Kong, China, November 3-7, 2019*, pages 5349–5358.
- Karin Kipper-Schuler. 2005. **VerbNet: A broad-coverage, comprehensive verb lexicon**. University of Pennsylvania.
- 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, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 486–492.
- Roberto Navigli. 2018. **Natural Language Understanding: Instructions for (present and future) use**. In *Proc. of IJCAI 2018*, pages 5697–5702.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. **Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network**. *Artif. Intell.*, 193:217–250.

- Martha Palmer, Paul R. Kingsbury, and Daniel Gildea. 2005. [The Proposition Bank: An annotated corpus of semantic roles](#). *Comput. Linguistics*, 31(1):71–106.
- Tommaso Pasini. 2020. [The knowledge acquisition bottleneck problem in multilingual Word Sense Disambiguation](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 4936–4942.
- 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*, CoNLL 2012, pages 1–40, Stroudsburg, PA, USA.
- Bianca Scarlini, Tommaso Pasini, and Roberto Navigli. 2020. [With more contexts comes better performance: Contextualized sense embeddings for all-round Word Sense Disambiguation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020*.
- Dan Shen and Mirella Lapata. 2007. [Using semantic roles to improve question answering](#). In *EMNLP-CoNLL 2007, Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, June 28-30, 2007, Prague, Czech Republic*, pages 12–21.
- Peng Shi and Jimmy Lin. 2019. [Simple BERT models for relation extraction and semantic role labeling](#). *arXiv preprint arXiv:1904.05255*, abs/1904.05255.
- 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, Brussels, Belgium, October 31 - November 4, 2018*, pages 5027–5038.
- 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*, abs/1706.09528.
- Zhixing Tan, Mingxuan Wang, Jun Xie, Yidong Chen, and Xiaodong Shi. 2018. [Deep semantic role labeling with self-attention](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4929–4936.
- Mark Yatskar, Luke S. Zettlemoyer, and Ali Farhadi. 2016. [Situation recognition: Visual semantic role labeling for image understanding](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 5534–5542.
- Meishan Zhang, Peili Liang, and Guohong Fu. 2019. [Enhancing opinion role labeling with semantic-aware word representations from semantic role labeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 641–646.