

AllenNLP: A Deep Semantic Natural Language Processing Platform

**Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi,
Nelson Liu, Matthew Peters, Michael Schmitz, Luke Zettlemoyer**

Allen Institute for Artificial Intelligence

{mattg,joelg,markn,oyvindt,pradeepd,nelsonl,matthewp,michaels,lukez}@allenai.org

Abstract

This paper describes AllenNLP, a platform for research on deep learning methods in natural language understanding. AllenNLP is designed to support researchers who want to build novel language understanding models quickly and easily. It is built on top of PyTorch, allowing for dynamic computation graphs, and provides (1) a flexible data API that handles intelligent batching and padding, (2) high-level abstractions for common operations in working with text, and (3) a modular and extensible experiment framework that makes doing good science easy. It also includes reference implementations of high quality approaches for both core semantic problems (e.g. semantic role labeling (Palmer et al., 2005)) and language understanding applications (e.g. machine comprehension (Rajpurkar et al., 2016)). AllenNLP is an ongoing open-source effort maintained by engineers and researchers at the Allen Institute for Artificial Intelligence.

1 Introduction

Neural networks are now commonplace in natural language processing research. They have enabled significant performance gains on a wide range of tasks, but it can be surprisingly difficult to tune new models or replicate existing results. For example, deep BiLSTM models (Zhou and Xu, 2015; He et al., 2017) have achieved over 20% error reduction for span-based semantic role labeling, the first major improvement in accuracy for this task in over 10 years. However, they take over a week to train on modern GPUs and are sensitive to initialization and hyperparameter settings.

These types of challenges provide a barrier to entry for research on many problems, given the need for very large scale experimentation.

This paper describes AllenNLP, a platform for research on deep learning methods in natural language understanding, which aims to significantly lower such barriers to high quality work. AllenNLP is designed to support researchers who want to build novel language understanding models as quickly and easily as possible. It is easy to install and use (Section 3), and provides functionality for quickly evaluating a wide range of standard modeling choices (Section 4; e.g. to easily experiment with different word representations, including embeddings and character-level CNNs). It also provides reference implementations of a range of (often very challenging to train) models (Section 5), along with high quality pretrained reference models.

AllenNLP is an ongoing open-source effort maintained by engineers and researchers at the Allen Institute for Artificial Intelligence. In addition to supporting the broader research community, we will also use AllenNLP for our work on core language understanding challenges, and share all of our results publicly (Section 6).

2 Related Toolkits and Platforms

AllenNLP is built on PyTorch,¹ to support rapid model development, and adds functionality for data management and experimentation on common NLP problems.

Most existing deep learning toolkits are designed for general machine learning (Bergstra et al., 2010; Yu et al., 2014; Chen et al., 2015; Abadi et al., 2016; Neubig et al., 2017), and significant effort can be required to develop research infrastructure for particular model classes. More

¹<http://pytorch.org/>

specialized toolkits exist in some domains. For example, Caffe (Jia et al., 2014) includes strong reference models trained on ImageNet (Deng et al., 2009), significantly lowering the barrier to entry for computer vision research. AllenNLP provides a similar type of support, for semantic NLP problems.

Many existing NLP pipelines (Manning et al., 2014; Bird et al., 2009)² abstract away from specific modeling decisions and are designed to make it easy to predict linguistic structures (e.g. POS tags or syntactic parse trees). AllenNLP is designed to support the development of new models, which could be used in such a pipeline. In that sense, AllenNLP is closely related to SyntaxNet,³ but focuses more on semantic tasks and supports a wider range of neural architectures.

Finally, AllenNLP is related to toolkits for deep learning research in dialog (Miller et al., 2017) and machine translation (Klein et al., 2017). Those toolkits support learning general functions that map strings (e.g. foreign language text or user utterance) to strings (e.g. English text or system responses). AllenNLP, in contrast, is designed to support models that predict structured semantic representations of the input text, such as coreference clusters and semantic role edges.

3 Getting Started

The AllenNLP website⁴ provides tutorials, reference model implementations, pretrained models, and an open source code base. The AllenNLP platform is easy to download and install, either via Docker image or cloning the GitHub repository. It includes reference models (see Section 5) that can be easily run (to make predictions on new sentences) and retrained with different hyperparameter settings or on new data. These pretrained models have online demos,⁵ and provide examples of the framework functionality (Section 4).⁶ They also serve as baselines for future research.

4 Designed for NLP Research

AllenNLP is a platform designed specifically for both deep learning and NLP research. It is built

²See also <https://spacy.io/>

³<https://github.com/tensorflow/models/tree/master/syntaxnet>

⁴<http://allennlp.org/>

⁵<http://demo.allennlp.org/>

⁶See <http://docs.allennlp.org/en/latest/> for the latest online documentation.

on top of PyTorch, allowing for dynamic computation graphs, and it provides (1) a flexible data API that handles intelligent batching and padding, (2) high-level abstractions for common operations in working with text, and (3) a modular and extensible experiment framework that makes doing good science easy.

4.1 Text Data Processing

AllenNLP's data processing API is built around the notion of `Fields`. Each `Field` represents a single input array to a model, and they are grouped together in `Instances` to create the input/output specification for a task. The `Field` API is flexible and easy to extend, allowing for a unified data API for tasks as diverse as tagging, semantic role labeling, question answering, and textual entailment. To represent the SQuAD dataset (Rajpurkar et al., 2016), for example, which has a question and a passage as inputs and a span from the passage as output, each `Instance` has a `TextField` for the question, a `TextField` for the passage, and two `IndexFields` representing the start and end positions of the answer in the passage.

Once a set of `Instance` objects has been created, the data API will automatically sort them into batches with similar sequence lengths, pad all sequences in each batch to the same length, and randomize the batches for input to a model. The only thing a user has to do is read data into a set of `Instance` objects with the desired fields, and the library takes care of the rest.

4.2 NLP-focused abstractions

AllenNLP provides a high-level API for building models, with abstractions designed specifically for natural language processing. These abstractions make it so that model code actually specifies a *class* of related models. Experiments with various architectures within this class is possible without changing a single line of model code.

The library has three key abstractions, dealing with (1) how text gets represented as vectors, (2) how vector sequences get modified to produce new vector sequences, and (3) how vector sequences get merged into a single vector.

TextFieldEmbedder: This abstraction takes input arrays generated by a `TextField` and returns a sequence of embedded vectors. Through the use of polymorphism and AllenNLP's experiment framework (see Section 4.3), any model that uses this abstraction can easily

switch between a wide variety of possible word representations. Deciding between pre-trained word embeddings, word embeddings concatenated with a character-level CNN encoding, or even using a pre-trained model to get token-in-context embeddings (Peters et al., 2017), is all done by configuring the `TextFieldEmbedder`, allowing for very easy controlled experimentation.

Seq2SeqEncoder: A very common paradigm in deep NLP is to take a sequence of word vectors and pass them through some kind of recurrent network to encode contextual information, getting a new sequence of vectors as output. There is a large number of ways to do this, including LSTMs (Hochreiter and Schmidhuber, 1997), GRUs (Cho et al., 2014), intra-sentence attention (Cheng et al., 2016), recurrent additive networks (Lee et al., 2017b), etc. AllenNLP’s `Seq2SeqEncoder` abstracts away the decision of which particular encoder to use, allowing the user to specify the encoder outside of model code. In this way, when a researcher is exploring new recurrent architectures, they can easily replace the LSTMs in *any model* that uses this abstraction with their new encoder, seeing the impact across a wide range of models and tasks.

Seq2VecEncoder: Similar to `Seq2SeqEncoders`, a common operation in NLP models is to merge a sequence of vectors into a single vector, using either a recurrent network with some kind of averaging or pooling, or using a convolutional network. This operation is encapsulated in AllenNLP by a `Seq2VecEncoder`. This abstraction again allows the model code to only describe a *class* of similar models, with particular instantiations of that model class being determined by configuration that happens later.

4.3 Experimental Framework

A main design goal of AllenNLP is to make it easy to do good science with controlled experiments. Because of the abstractions described in Section 4.2, large parts of the model architecture can be configured outside of model code, in addition to other training-related hyper-parameters. This makes it easy to specify the important decisions that define a new model, without having to code all of the implementation details from scratch.

This architecture design is done in AllenNLP

through a simple configuration file that specifies, e.g., which text representations and encoders to use in an experiment. Mapping from strings in the configuration file to instantiated objects in code is done through the use of a *registry*, which allows users of the library to add new implementations of any of the provided abstractions.

5 Reference Models

AllenNLP includes reference implementations of widely used language understanding models. The models show how to use much of the framework functionality presented in Section 4. They also have verified performance levels that closely match the original results, and can serve as comparison baselines for future research.

At the time of launch, AllenNLP includes reference implementations for three tasks:

- **Semantic Role Labeling** (SRL) models predict the verbal predicate argument structure of a sentence (Palmer et al., 2005). The AllenNLP toolkit contains a deep BiLSTM SRL model (He et al., 2017) that is state of the art for PropBank SRL, at the time of publication.
- **Machine Comprehension** (MC) systems take an evidence text and a question as input, and predict a span within the evidence that answers the question. AllenNLP includes a reference implementation of the BiDAF MC model (Seo et al., 2017) which was state of the art for the SQuAD benchmark (Rajpurkar et al., 2016) in early 2017.
- **Textual Entailment** (TE) systems take a pair of sentences and predict whether the facts in one necessarily imply that the other is true. AllenNLP includes a reference implementation of the decomposable attention TE model (Parikh et al., 2016) which was state of the art for the SNLI benchmark (Bowman et al., 2015) in late 2016.

Additional models are currently under development and should be released soon, including: end-to-end neural coreference (Lee et al., 2017a), and semi-supervised learning for named entity recognition (Peters et al., 2017). We also expect the number of tasks and reference implementations to grow steadily over time.⁷

⁷The most up-to-date list of reference models is maintained online: <http://allennlp.org/models>

6 Semantics Research with AllenNLP

The AllenNLP toolkit is designed for use by a wide variety of NLP researchers, and will be used extensively in many of the language understanding research efforts currently under way at the Allen Institute for Artificial Intelligence and the University of Washington.

These efforts will focus core semantic problems, likely including (1) efforts to generalize semantic role labeling to all words (not just verbs), (2) models for general coreference resolution (e.g. entities, events, bridging, etc.), (3) semantic parsers that build relatively complete meaning representations (e.g. mapping language to code), and (4) approaches for semi-supervised learning of improved word representations. We will also emphasize building models that work well on text genres that are typically out of domain, such as science texts. Finally, we will focus on making large new datasets for these generalized semantic understanding tasks (e.g. QA-SRL (He et al., 2015)), annotating data in a variety of domains, and releasing all of our developed resources for broad use.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. 2010. Theano: A cpu and gpu math compiler in python. In *Proc. 9th Python in Science Conf.* pages 1–7.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*.
- Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the Computer Vision and Pattern Recognition*. IEEE, pages 248–255.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and whats next. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- 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 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pages 643–653.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM international conference on Multimedia*. ACM, pages 675–678.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M Rush. 2017. OpenNMT: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810*.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017a. End-to-end neural coreference resolution. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2017b. Recurrent additive networks. *arXiv preprint arXiv:1705.07393*.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language

- processing toolkit. In *Proceedings of the Association for Computational Linguistics (ACL) (System Demonstrations)*.
- Alexander H Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, and Jason Weston. 2017. Parlai: A dialog research software platform. *arXiv preprint arXiv:1705.06476* .
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, et al. 2017. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980* .
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics* 31(1):71–106.
- Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Matthew E Peters, Waleed Ammar, Chandra Bhagavathula, and Russell Power. 2017. Semi-supervised sequence tagging with bidirectional language models. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension .
- Dong Yu, Adam Eversole, Mike Seltzer, Kaisheng Yao, Zhiheng Huang, Brian Guenter, Oleksii Kuchaiev, Yu Zhang, Frank Seide, Huaming Wang, et al. 2014. An introduction to computational networks and the computational network toolkit. *Microsoft Technical Report MSR-TR-2014-112* .
- Jie Zhou and Wei Xu. 2015. End-to-end learning of semantic role labeling using recurrent neural networks. In *Proceedings of the Association for Computational Linguistics (ACL)*.