PyStruct - Learning Structured Prediction in Python

Andreas C. Müller   AMUELLER@AIS.UNI-BONN.DE
Sven Behnke   BEHNKE@CS.UNI-BONN.DE
Institute of Computer Science, Department VI
University of Bonn
Bonn, Germany

Abstract

Structured prediction methods have become a central tool for many machine learning applications. While more and more algorithms are developed, only very few implementations are available.

PyStruct aims at providing a general purpose implementation of standard structured prediction methods, both for practitioners and as a baseline for researchers. It is written in Python and adapts paradigms and types from the scientific Python community for seamless integration with other projects.

Keywords: structured prediction, structural support vector machines, conditional random fields, Python

1. Introduction

In recent years there has been a wealth of research in methods for learning structured prediction, as well as in their application in areas such as natural language processing and computer vision. Unfortunately only few implementations are publicly available—many applications are based on the non-free implementation of Joachims et al. (2009).

PyStruct aims at providing a high-quality implementation with an easy-to-use interface, in the high-level Python language. This allows practitioners to efficiently test a range of models, as well as allowing researchers to compare to baseline methods much more easily than this is possible with current implementations. PyStruct is BSD-licensed, allowing modification and redistribution of the code, as well as use in commercial applications. By embracing paradigms established in the scientific Python community and reusing the interface of the widely-used SCIKIT-LEARN library (Pedregosa et al., 2011), PyStruct can be used in existing projects, replacing standard classifiers. The online documentation and examples help new users understand the somewhat abstract ideas behind structured prediction.

2. Structured Prediction and Casting it into Software

Structured prediction can be defined as making a prediction $f(x)$ by maximizing a compatibility function between an input x and the possible labels y (Nowozin and Lampert, 2011). Most current approaches use linear functions, leading to:

$$f(x) = \arg\max_{y \in \mathcal{Y}} \ \theta^T \Psi(x, y). \tag{1}$$

Here, y is a *structured* label, Ψ is a joint feature function of x and y, and θ are parameters of the model. *Structured* means that y is more complicated than a single output class, for example a label for each word in a sentence or a label for each pixel in an image. Learning structured prediction means learning the parameters θ from training data.

Using the above formulation, learning can be broken down into three sub-problems:

1. Optimizing the objective with respect to θ.
2. Encoding the structure of the problem in a joint feature function Ψ.
3. Solving the maximization problem in Equation 1.

The later two problems are usually tightly coupled, as the maximization in Equation 1 is usually only feasible by exploiting the structure of Ψ, while the first is treated as independent. In fact, when 3. can not be done exactly, learning θ strongly depends on the quality of the approximation. However, treating approximate inference and learning as a joint optimization problem is currently out of the scope of the package, and we implement a more modular setup. PYSTRUCT takes an object-oriented approach to decouple the task-dependent implementation of 2. and 3. from the general algorithms used to solve 1.

Estimating θ is done in `learner` classes, which currently support cutting plane algorithms for structural support vector machines (SSVMs Joachims et al. (2009)), subgradient methods for SSVMs Ratliff et al. (2007), Block-coordinate Frank-Wolfe (BCFW) (Lacoste-Julien et al., 2012), the structured perceptron and latent variable SSVMs (Yu and Joachims, 2009). The cutting plane implementation uses the CVXOPT package (Andersen et al., 2012) for quadratic optimization. Encoding the structure of the problem is done using `model` classes, which compute Ψ and encode the structure of the problem. The structure of Ψ determines the hardness of the maximization in Equation (1) and is a crucial factor in learning. PYSTRUCT implements models (corresponding to particular forms of Ψ) for many common cases, such as multi-class and multi-label classification, conditional random fields with constant or data-dependent pairwise potentials, and several latent variable models. The maximization for finding y in Equation 1 is carried out using external libraries, such as OPENGM (Kappes et al., 2013), LIBDAI (Mooij, 2010) and others. This allows the user to choose from a wide range of optimization algorithms, including (loopy) belief propagation, graph-cuts, QPBO, dual subgradient, MPBP, TRWs, LP and many other algorithms. For problems where exact inference is infeasible, PYSTRUCT allows the use of linear programming relaxations, and provides modified loss and feature functions to work with the continuous labels. This approach, which was outlined in Finley and Joachims (2008) allows for principled learning when exact inference is intractable. When using approximate integral solvers, learning may finish prematurely and results in this case depend on the inference scheme and learning algorithm used.

Table 1 lists algorithms and models that are implemented in PYSTRUCT and compares them to other public structured prediction libraries: DLIB (King, 2009), SVM$^{\text{struct}}$ (Joachims et al., 2009) and CRFSUITE (Okazaki, 2007). We also give the programming language and the project license.

Package	Language	License	Algorithms				Models			
			CP	SG	BCFW	LV	ML	Chain	Graph	LDCRF
PyStruct	Python	BSD[1]	✓[1]	✓	✓	✓	×	✓	✓	✓
SVM$^{\text{struct}}$	C++	non-free	✓	×	×	✓	×	×	×	×
Dlib	C++	boost	✓	×	×	×	×	✓	✓	×
CRFsuite	C++	BSD	×	×	×	×	✓	✓	×	×

Table 1: Comparison of structured prediction software packages. CP stands for cutting plane optimization of SSVMs, SG for online subgradient optimization of SSVMs, LV for latent variable SSVMs, ML for maximum likelihood learning, Chain for chain-structured models with pairwise interactions, Graph for arbitrary graphs with pairwise interactions, and LDCRF for latent dynamic CRF (Morency et al., 2007). [1]PyStruct itself is BSD licensed, but uses the GPL-licensed package cvxopt for cutting-plane learning.

3. Usage Example: Semantic Image Segmentation

Conditional random fields are an important tool for semantic image segmentation. We demonstrate how to learn an n-slack support vector machine (Tsochantaridis et al., 2006) on a superpixel-based CRF on the popular Pascal data set. We use unary potentials generated using TextonBoost from Krähenbühl and Koltun (2012). The superpixels are generated using SLIC (Achanta et al., 2012).[1] Each sample (corresponding on one entry of the list X) is represented as a tuple consisting of input features and a graph representation.

```
model = crfs.EdgeFeatureGraphCRF(
    class_weight=inverse_frequency, symmetric_edge_features=[0, 1],
    antisymmetric_edge_features=[2], inference_method='qpbo')

ssvm = learners.NSlackSSVM(model, C=0.01, n_jobs=-1)
ssvm.fit(X, Y)
```

Listing 1: Example of defining and learning a CRF model.

The source code is shown in Listing 1. Lines 1-3 declare a model using parametric edge potentials for arbitrary graphs. Here class_weight re-weights the hamming loss according to inverse class frequencies. The parametric pairwise interactions have three features: a constant feature, color similarity, and relative vertical position. The first two are declared to be symmetric with respect to the direction of an edge, the last is antisymmetric. The inference method used is QPBO-fusion moves. Line 5 creates a learner object that will learn the parameters for the given model using the n-slack cutting plane method, and line 6 performs the actual learning. Using this simple setup, we achieve an accuracy of 30.3 on the validation set following the protocol of Krähenbühl and Koltun (2012), who report 30.2 using a more complex approach. Training the structured model takes approximately 30 minutes using a single i7 core.

---

1. The preprocessed data can be downloaded at http://www.ais.uni-bonn.de/download/datasets.html.

Figure 1: Runtime comparison of PYSTRUCT and SVM^struct for multi-class classification.

4. Experiments

While PYSTRUCT focuses on usability and covers a wide range of applications, it is also important that the implemented learning algorithms run in acceptable time. In this section, we compare our implementation of the 1-slack cutting plane algorithm (Joachims et al., 2009) with the implementation in SVM^struct. We compare performance of the Crammer-Singer multi-class SVM with respect to learning time and accuracy on the MNIST data set of handwritten digits. While multi-class classification is not very interesting from a structured prediction point of view, this problem is well-suited to benchmark the cutting plane solvers with respect to accuracy and speed.

Results are shown in Figure 1. We report learning times and accuracy for varying regularization parameter C. The MNIST data set has 60 000 training examples, 784 features and 10 classes.[2] The figure indicates that PYSTRUCT has competitive performance, while using a high-level interface in a dynamic programming language.

5. Conclusion

This paper introduced PYSTRUCT, a modular structured learning and prediction library in Python. PYSTRUCT is geared towards ease of use, while providing efficient implementations. PYSTRUCT integrates itself into the scientific Python eco-system, making it easy to use with existing libraries and applications. Currently, PYSTRUCT focuses on max-margin and perceptron-based approaches. In the future, we plan to integrate other paradigms, such as sampling-based learning (Wick et al., 2011), surrogate objectives (for example pseudo-likelihood), and approaches that allow for a better integration of inference and learning (Meshi et al., 2010).

Acknowledgments

---

2. Details about the experiment and code for the experiments can be found on the project website.

References

Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *PAMI*, 2012.

Martin S. Andersen, Joachin Dahl, and Lieven Vandenberghe. CVXOPT: A Python package for convex optimization, version 1.1.5. Available at `http://cvxopt.org/`, 2012.

Thomas Finley and Thorsten Joachims. Training structural SVMs when exact inference is intractable. In *ICML*, 2008.

Thorsten Joachims, Thomas Finley, and Chun-Nam John Yu. Cutting-plane training of structural SVMs. *JMLR*, 77(1), 2009.

Jörg H Kappes, Bjoern Andres, Fred A Hamprecht, Christoph Schnörr, Sebastian Nowozin, Dhruv Batra, Sungwoong Kim, Bernhard X Kausler, Jan Lellmann, Nikos Komodakis, et al. A comparative study of modern inference techniques for discrete energy minimization problems. In *CVPR*, 2013.

Davis E. King. Dlib-ml: A machine learning toolkit. *JMLR*, 10, 2009.

Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected CRFs with Gaussian edge potentials. In *NIPS*, 2012.

Simon Lacoste-Julien, Mark Schmidt, and Francis Bach. A simpler approach to obtaining an O(1/t) convergence rate for projected stochastic subgradient descent. *arXiv preprint arXiv:1212.2002*, 2012.

Ofer Meshi, David Sontag, Tommi Jaakkola, and Amir Globerson. Learning efficiently with approximate inference via dual losses. In *ICML*, 2010.

Joris M. Mooij. libDAI: A free and open source C++ library for discrete approximate inference in graphical models. *JMLR*, 2010.

L-P Morency, Ariadna Quattoni, and Trevor Darrell. Latent-dynamic discriminative models for continuous gesture recognition. In *CVPR*, 2007.

Sebastian Nowozin and Christoph H. Lampert. *Structured Learning and Prediction in Computer Vision*. Now Publishers Inc., 2011.

Naoaki Okazaki. CRFsuite: A fast implementation of conditional random fields (CRFs), 2007. URL `http://www.chokkan.org/software/crfsuite/`.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in Python. *JMLR*, 2011.

Nathan Ratliff, J. Andrew (Drew) Bagnell, and Martin Zinkevich. (Online) subgradient methods for structured prediction. In *AISTATS*, 2007.

Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, Yasemin Altun, and Yoram Singer. Large margin methods for structured and interdependent output variables. *JMLR*, 6(2), 2006.

Michael Wick, Khashayar Rohanimanesh, Kedar Bellare, Aron Culotta, and Andrew Mc-Callum. Samplerank: Training factor graphs with atomic gradients. In *ICML*, 2011.

Chun-Nam John Yu and Thorsten Joachims. Learning structural SVMs with latent variables. In *ICML*, 2009.