

# Call For Papers: Machine Learning Special Issue on Structured Prediction

Guest Editors:

Charles Parker, Yasemin Altun, and Prasad Tadepalli

Structured prediction, in the broadest sense, is the task of predicting a set of output variables from a set of input variables, where the values of any of the output variables may be correlated with any number of input variables or other output variables. There has been a great deal of recent work on this topic in journals and conferences, and it is now the subject of a new book, “Predicting Structured Objects”, based on a 2004 NIPS workshop on the subject. In addition, robust audiences for workshops at ICML in 2006 and 2007 show that the topic continues to generate interest.

Following this, we would like to invite submissions to a special issue of machine learning on “Structured Prediction”. Submissions on all areas of this broad field are encouraged, especially those that address the following concerns:

- **New Directions:** Recent work has taken tentative steps towards generalizing structured prediction to the unsupervised and semi-supervised cases. Conversely, other work has attempted to improve performance and expand theory on certain specialized instances of the structured prediction problem. What are the results of this work, both empirically and theoretically?
- **Theoretical Underpinnings:** Many of the current structured prediction algorithms have very different theoretical guarantees. Can the theory across different algorithms be unified? What happens when the preconditions of these guarantees, such as the requirement of exact inference under specified model parameters, are violated? Are there any other guarantees that can be made? Does the algorithm hold up empirically or fail?
- **Large-scale or Real-world Implementations:** Results have been reported in computational linguistics on a number of real world datasets. What other real world problems can be solved successfully using this formalism? In addition, many these methods appear to suffer from scaling issues. What can be done to alleviate these issues in practice? In what cases may approximations be appropriate and how can they be efficiently computed?

- **Empirical Evaluation:** Many results reported so far compare only two existing techniques against a single dataset. What do broad comparisons of many techniques over many datasets say about the relative effectiveness of these algorithms? Which methods are effective in which situations and when are they ineffective? What does this reveal about the connections between different algorithms?

## Important Dates

Titles and Abstracts due: **10 April 2008**  
Papers due: **20 April 2008**  
Author notification: **15 July 2008**  
Final versions due: **8 October 2008**

## Submission Information

Submissions must follow the publication guidelines set forth by the Machine Learning Journal. Templates and style files are available at:

`ftp://ftp.springer.de/pub/tex/latex/svjour/global.zip`

More information for authors is located at:

`http://pages.stern.nyu.edu/~fprovost/MLJ/info-for-authors.html`

Please do not send submissions directly to the editors, but submit them to Springer's online submission system:

`http://www.editorialmanager.com/mach`

Submissions are expected to be the unpublished, original work of the authors and not currently under consideration for publication by any other journal. Expansions of papers previously published at refereed workshops and conferences are acceptable, but this must be indicated in the submission, and the expansion must be significant.

If the submission is accepted, the author must provide Springer with exclusive commercial distribution rights for six months. Authors will retain copyright of their work and be free to distribute it as they see fit, provided the distribution is non-commercial until six months after publication. Papers are published electronically immediately upon final acceptance.

Further inquiries about the submissions should be directed to:

`clp-at-clparker.org`