

## Call for Papers

### IEEE Transactions on Audio, Speech, and Language Processing

#### Special Issue on Large-Scale Optimization for Audio, Speech, and Language Processing

Large-scale optimization algorithms are finding a broad range of applications in modern computing. In the fields of data mining and signal processing, they have become ubiquitous. This special issue creates a forum for researchers working in various areas of audio, speech, and language processing to come together with optimization researchers and share ideas for improving the use of optimization approaches in these areas.

Firmer understanding of the relationships between standard optimization methods and the specialized approaches currently used in speech and language processing form a basis for future work at the intersection of these areas. Further leveraging of algorithmic ideas from optimization, cross-fertilization with existing algorithms, recognition of special structures and challenges, and adaptation to novel parallel computing environments will lead to significant advances in the state of the art.

Pattern recognition in audio, speech, and language processing requires estimation of parameters in statistical models via some optimization criteria. At large scale, these problems present challenges that cannot be resolved by naïve application of well known optimization techniques. Second-order algorithms cannot be applied directly to very large data sets, and even conventional first-order algorithms are impractical when they require repeated sweeps through the data. It is difficult even to obtain a well defined optimization formulation of the pattern recognition task. For example, likelihood criteria usually are inadequate if the training data do not represent all possible variations in patterns.

Significant progress in pattern recognition was achieved by introducing discrimination criteria for training, but overtraining remains a danger. A major challenge is to couple fast optimization techniques for these very large data sets with formulation techniques that prevent overtraining and degradation of pattern-recognition accuracy. Such formulations would allow prior information to be incorporated into the model, along with regularization techniques.

At the system level, fusion of decisions is commonly used in such speech and language processing problems as speaker / language recognition, and speech recognition / machine translation. As the optimization takes place at the system level, it involves many parameters of different types. This special issue provides a forum for authors to share their findings across different speech and language applications.

Over the past decades, a variety of specialized approaches have been proposed to solve pattern recognition problems. Recent success has been obtained with adaptations of conventional optimization approaches, including L-BFGS, inexact Newton methods, coordinate descent, and stochastic gradient methods. One goal of the special issue is to build on these successes, identifying further relevant optimization techniques, hybridizing these approaches, analyzing their convergence properties, and specializing them to specific pattern recognition problems in audio, speech, and language processing. The community would benefit from a broader view that incorporates recent advances in large-scale optimization methods.

Other important avenues of research could include algorithms that use parallel computing architectures, including GPUs. Recent developments in optimization, machine learning, and computational statistics could be leveraged here. Another possibility is multilevel algorithms, in which part of the parameter search can be performed in reduced spaces, potentially improving robustness and efficiency.

Because of the size of the data set, special attention must be paid to data handling and movement the type of data structures. These factors must be considered in adapting and implementing optimization approaches to pattern recognition problems effectively.

In light of the important research already performed in this exciting space, we invite papers describing various aspects of large-scale optimization in audio, speech, and language processing. All submissions must have specific connection to audio, speech, and language processing. Within this scope, topics of particular interest include, but are not limited to, the following.

- Hybrid modeling and optimization
- Computational studies of algorithm performance on large data sets
- Stochastic and semi-stochastic optimization methods
- Algorithms for parallel architectures, including clusters and GPUs
- High dimensional MCMC methods
- Sparse and regularized optimization
- Inverse methods
- Optimization techniques in discriminative training
- System-level optimization for fusion of decisions
- Multilevel optimization algorithms
- Effective handling of data in optimization algorithms.

The authors are required to follow the Author's Guide for manuscript submission to the IEEE Transactions on Audio, Speech, and Language Processing at <http://www.signalprocessingsociety.org/publications/periodicals/taslp/taslp-author-information>

- \* Submission deadline: Oct. 30, 2012
- \* Notification of first round review: Dec. 30, 2012
- \* Notification of acceptance: Mar. 30, 2013
- \* Final manuscripts due: Apr. 30, 2013
- \* Date of publication: Aug. 1, 2013

For further information, please contact the guest editors:

Dimitri Kanevsky (kanevsky@us.ibm.com)  
Xiaodong He (xiaohe@microsoft.com)  
Georg Heigold (heigold@google.com)  
Haizhou Li (hli@i2r.a-star.edu.sg)  
Hermann Ney (ney@informatik.rwth-aachen.de)  
Stephen Wright (swright@cs.wisc.edu)