Computational Management Science

Special Issue on
"Optimization in Learning from Data"

Guest Editors

Erhun Kundakcioglu, Marcello Sanguineti, Theodore Trafalis

Scope

In the last decades, data have become chief characters of modern life. Biological, financial, manufacturing, medical, and communications data pervade our world and witness the ubiquity of data, their analysis, and their management.

On one side, more and more applications produce data that have to be processed; on the other side, more and more applications are based on data that have already been processed: we are living in a society that is at the same time a huge data-producer and a huge data-consumer. Nowadays learning from data is a melting pot involving scientists from a large variety of areas: researchers from Computer Science, Operations Research, Engineering, Mathematics, Physics, Statistics, Biology, and Medicine provide challenging data mining problems and investigate innovative solution methodologies.

Some major underlying themes of modern data-driven and data-producing applications are:
- data richness (they elaborate and/or produce huge amounts of data);
- on-line adaptation (they require data-processing algorithms that must be updated fast as soon as new observations become available);
- high-dimensionality (each single datum often has a large dimension);
- structured nature (some kinds of data, e.g., images, bio-sequences, weighted graphs, cannot be represented simply as vectors of real numbers).

Classical statistical methods are simply not designed to cope with problems exhibiting such features. On the other hand, the last decades have seen major advances in optimisation: the development of new mathematical tools has kept up with the availability of powerful modelling languages and the dramatic improvement of computational capabilities. New methods of learning from data can strongly benefit from optimisation-based approaches and exploit them to deal with the kind of data arising in areas of science and technology of great current interest.

The goal of the special issue is two-fold:
- to contribute to the development of powerful solution methodologies and efficient algorithms for learning from data, based on optimisation theory and techniques;
- to present to the Operations Research community and to others interested in learning from data, challenging problems from data-intensive applications, which can be faced via optimisation tools.

Topics

Papers with an optimisation-oriented flavour are solicited from, but not limited to, the following topics:

Regression, classification, pattern recognition, clustering
Principal component analysis
Supervised, semi-supervised, and unsupervised learning
Reinforcement learning

On-line learning
Manifold learning
Support vector machines
Kernel methods
Neural networks
Inverse problems in learning from data
High-dimensional data, curse of dimensionality, and dimensionality reduction
Learning by labeled and unlabeled data
Data de-noising
Geometry of learning
Statistical learning and VC-dimension
Covering numbers in learning theory
Learning in reproducing kernel Hilbert spaces
Regularization
Sparse approximation
Empirical and expected errors minimization
Generalization and approximation errors
Approximation schemes and approximation rates
Computational complexity and optimal learning algorithms
Greedy learning algorithms
Primal and dual formulations
Learning over graphs: graph kernels, graph regularization, graph Laplacian, and spectral graph theory
Supervised graph inference and predictive graph mining
Robustness, sparseness, stability, and consistency of algorithms
Loss functions
Quadratic Programming
Regularized least-squares
Penalty functionals
Reoptimisation

Review Process

Manuscripts will be screened for topical relevance, and those that pass the screening process will undergo the standard peer-review process of Computational Management Science. The major criterion for acceptance is the quality and originality of the contribution.