

CIS Publication Spotlight

IEEE Transactions on Neural Networks and Learning Systems

Low-Rank Structure Learning via Non-convex Heuristic Recovery, by Y. Deng, Q. Dai, R. Liu, Z. Zhang, and S. Hu, *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 24, No. 3, March 2013, pp. 383–396.

Digital Object Identifier: 10.1109/TNNLS.2012.2235082

A nonconvex framework is proposed for learning the essential low-rank structure from corrupted data. Different from traditional approaches, which directly utilizes convex norms to measure the sparseness, this method introduces more reasonable nonconvex measurements to enhance the sparsity in both the intrinsic low-rank structure and the sparse corruptions. It includes how to combine the widely used L_p norm ($0 < p < 1$) and log-sum term into the framework of low-rank structure learning. Although the proposed optimization is no longer convex, it still can be effectively solved by a majorization–minimization (MM)-type algorithm, with which the nonconvex objective function is iteratively replaced by its convex surrogate and the nonconvex problem finally falls into the general framework of reweighed approaches. It is proved that

the MM-type algorithm can converge to a stationary point after successive iterations. The proposed model is applied to solve two typical problems: robust principal component analysis and low-rank representation. Experimental results on low-rank structure learning demonstrate that our nonconvex heuristic methods, especially the log-sum heuristic recovery algorithm, generally perform much better than the convex-norm-based method ($0 < p < 1$) for both data with higher rank and with denser corruptions.”

Qualitative Adaptive Reward Learning with Success Failure Maps: Applied to Humanoid Robot Walking, by J. Nassour, V. Hugel, F.B. Ouezdou, and G. Cheng, *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 24, No. 1, January 2013, pp. 81–93.

Digital Object Identifier: 10.1109/TNNLS.2012.2224370

“A learning mechanism is proposed to learn from negative and positive feedback with reward coding adaptively. It is composed of two phases: evaluation and decision making. In the evaluation phase, a Kohonen self-organizing map technique is used to represent success and failure. Decision making is based on an early warning mechanism that enables avoiding repeating past mistakes. The behavior to risk is modulated in



© CORBIS

order to gain experiences for success and for failure. Success map is learned with adaptive reward that qualifies the learned task in order to optimize the efficiency. The approach is presented with an implementation on the NAO humanoid robot, controlled by a bio-inspired neural controller based on a central pattern generator. The learning system adapts the oscillation frequency and the motor neuron gain in pitch and roll in order to walk on flat and sloped terrain, and to switch between them.”

IEEE Transactions on Fuzzy Systems

A Novel Approach to Filter Design for T–S Fuzzy Discrete-Time Systems with Time-Varying Delay, *IEEE Transactions on Fuzzy Systems*, Vol. 20, No. 6, December 2012, pp. 1114–1129.

Digital Object Identifier: 10.1109/TFUZZ.2012.2196522

“In this paper, the problem of l_2 – l_∞ filtering for a class of discrete-time Takagi–Sugeno (T–S) fuzzy time-varying delay systems is studied. The authors focused on the design of full- and reduced-order filters that guarantee the filtering error system to be asymptotically stable with a prescribed H_∞ performance. Sufficient conditions for the obtained filtering error system are proposed by applying an input-output

approach and a two-term approximation method, which is employed to approximate the time-varying delay. The corresponding full- and reduced-order filter design is cast into a convex optimization problem, which can be efficiently solved by standard numerical algorithms. Finally, simulation examples are provided to illustrate the effectiveness of the proposed approaches.”

Fuzzy c-Means Algorithms for Very Large Data, IEEE Transactions on Fuzzy Systems, Vol. 20, No. 6, December 2012, pp. 1130–1146.

Digital Object Identifier: 10.1109/TFUZZ.2012.2201485

“Very large (VL) data or big data are any data that we cannot load into our computer’s working memory. This is not an objective definition, but a definition that is easy to understand and practical, because there is a dataset too big for any computer we might use; hence, this is VL data for us. Clustering is one of the primary tasks used in the pattern recognition and data mining communities to search VL databases (including VL images) in various applications, and so, clustering algorithms that scale well to VL data are important and useful. This paper compares the efficacy of three different implementations of techniques aimed to extend fuzzy c-means (FCM) clustering to VL data. Specifically, we compare methods that are based on 1) sampling followed by noniterative extension; 2) incremental techniques that make one sequential pass through subsets of the data; and 3) kernelized versions of FCM that provide approximations based on sampling, including three proposed algorithms. Empirical results show that random sampling plus extension FCM, bit-reduced FCM, and approximate kernel FCM are good choices to approximate FCM for VL data. We conclude by demonstrating the VL algorithms on a dataset with 5 billion objects and presenting a set of recommendations regarding the use of different VL FCM clustering schemes.”

IEEE Transactions on Evolutionary Computation

Continuous Dynamic Constrained Optimization—The Challenges, by T. Nguyen and X. Yao, IEEE Transactions on Evolutionary Computation, Vol. 16, No. 6, December 2012, pp. 769–786.

Digital Object Identifier: 10.1109/TEVC.2011.2180533

“Many real-world dynamic problems have both objective functions and constraints that can change over time. Currently no research addresses whether current algorithms work well on continuous dynamic constrained optimization problems. There also is no benchmark problem that reflects the common characteristics of continuous dynamic optimization problems. This paper attempts to close this gap. The authors present some investigations on the characteristics that might make these problems difficult to solve by some existing dynamic optimization and constraint handling algorithms. A set of benchmark problems with these characteristics is presented. Finally, list of potential requirements that an algorithm should meet to solve these type of problems is proposed.”

The Automatic Design of Multiobjective Ant Colony Optimization Algorithms, by M. Lopez-Ibanez and T. Stutzle, IEEE Transactions on Evolutionary Computation, Vol. 16, No. 6, December 2012, pp. 861–875.

Digital Object Identifier: 10.1109/TEVC.2011.2182651

“Multiobjective optimization problems are problems with several, often conflicting, objectives that must be optimized. Without any a priori preference information, the Pareto optimality principle establishes a partial order among solutions, and the output of the algorithm becomes a set of nondominated solutions rather than a single one. Various ant colony optimization (ACO) algorithms have been proposed in recent years for solving such problems. This paper proposes a formulation of algorithmic components that suffices

to describe most multiobjective ACO algorithms proposed so far. The proposed framework facilitates the application of automatic algorithm configuration techniques.”

IEEE Transactions on Computational Intelligence and AI in Games

Monte Carlo Tree Search for the Hide-and-Seek Game Scotland Yard, by Pim Nijssen and Mark H.M. Winands, IEEE Transactions on Computational Intelligence and AI in Games, Vol. 4, No. 4, December 2012, pp. 282–294.

Digital Object Identifier: 10.1109/TCIAIG.2012.2210424

“This paper develops a strong Monte-Carlo Tree Search player for Scotland Yard, an interesting asymmetric imperfect information 2-player strategy game. The game involves one player controlling five detectives trying to capture a “hider.” A novel combination of techniques are used including determinization, location categorization and coalition reduction, the latter of which aims to optimally balance the tendencies for detectives to behave in glory hunting versus parasitic ways.”

IEEE Transactions on Autonomous Mental Development

A Unified Account of Gaze Following, by H. Jasso, J. Triesch, G. Deák, and J.M. Lewis, IEEE Transactions on Autonomous Mental Development, Vol. 4, No. 4, December 2012, pp. 257–272.

Digital Object Identifier: 10.1109/TAMD.2012.2208640

“Gaze following, the ability to redirect one’s visual attention to look at what another person is seeing, is foundational for imitation, word learning, and theory-of-mind. Previous theories have suggested that the development of gaze following in human infants is the product of a basic gaze following mechanism, plus the gradual incorporation of several distinct new mechanisms that improve the skill, such as spatial inference, and the

ability to use eye direction information as well as head direction. In this paper, we offer an alternative explanation based on a single learning mechanism. From a starting state with no knowledge of the

implications of another organism's gaze direction, our model learns to follow gaze by being placed in a simulated environment where an adult caregiver looks around at objects. Our infant model

matches the development of gaze following in human infants as measured in key experiments that we replicate and analyze in detail.”



CALL FOR PAPERS

General Game Systems

IEEE Transactions on Computational Intelligence and AI in Games (T-CAIAG)

Special issue: General Game Systems

Special issue editors: Cameron Browne, Nathan Sturtevant and Julian Togelius

General game playing (GGP) involves the development of AI agents for playing a range of games well, rather than specialising in any one particular game. Such systems have potential benefits for AI research, where the creation of general intelligence remains one of the open grand challenges.

GGP was first proposed in the 1960s and became a reality in the 1990s with the METAGAME system for general Chess-like games. The specification of the *game description language* (GDL) and annual AAAI GGP competitions followed in the first decade of this century, providing a platform for serious academic study into this topic. The recent advent of *Monte Carlo tree search* (MCTS) methods has allowed the development of truly competitive GGP agents, and there is exciting new research into applying GGP principles to general video games.

The field of general games research is now becoming fully rounded, with the development of complete *general game systems* (GGS) for playing, analysing and/or designing new games. These include not only GGP, but also any system that attempts to model a range of games; the definition is itself kept deliberately broad. The key feature of such systems is their generality, but the issue of representation remains an obstacle to true universality while they rely on formal descriptions of target domains.

The purpose of this special issue is to draw together the various research topics related to AI and CI in general games, to give an indication of where the field currently stands and where it is likely to head. It will explore questions such as: *How good and how general are existing systems, and how good and how general can they become? What have we learnt about AI and CI from studying general games? How do we apply existing GGP expertise to general video games?* We invite high quality work on any aspect of general games research in any genre of game—digital or physical—including play, analysis and design. Topics include but are not limited to:

- ☐ General game playing
- ☐ General game description and representation
- ☐ General game design and optimisation
- ☐ Generalised Monte Carlo tree search (MCTS) approaches
- ☐ Real-time, nondeterministic and imperfect information extensions to GGP
- ☐ General video game playing
- ☐ Framing issues and constraints on generality
- ☐ Bridging the gap between academic and commercial applications

Authors should follow normal *T-CAIAG* guidelines for their submissions, but identify their papers for this special issue during the submission process. Submissions should be 8 to 12 pages long, but may exceed these limits in special cases. Short papers of 4 to 6 pages are also invited. See <http://www.ieee-cis.org/pubs/tciaig/> for author information.

Deadline for submissions: May 3, 2013

Notification of Acceptance: July 5, 2013

Final copy due: October 4, 2013

Publication: December 6, 2013

Digital Object Identifier 10.1109/MCI.2013.2247901