Introduction to Special Issue on Computational Sustainability

Computational Sustainability is a newly emerging interdisciplinary field that aims to apply techniques from computer and information science and related disciplines, for instance, operations research, applied mathematics, and statistics, to the balancing of environmental, economic, and societal needs for sustainable development. Research in computational sustainability is inherently interdisciplinary, bringing together computational sciences and a variety of other disciplines as diverse as environmental sciences and engineering, biology, economics, and sociology.

This special issue of ACM *TIST* provides a sample of state-of-the-art research in computational sustainability, highlighting the role that artificial intelligence can play in this new field. The articles included in this issue cover several sustainability topics ranging from the monitoring of ecosystems to the tracking of animal movement, the management of renewable energy resources and the smart-grid, and the modeling of social behavior and social-ecological systems.

In "Submodularity and Its Applications in Optimized Information Gathering," Andreas Krause and Carlos Guestrin present a collection of algorithms for information gathering and sensing tasks. A general property exploited by the algorithms is submodularity, which allows for effective and efficient approximation algorithms that perform well in many large-scale, real-world information gathering and sensing tasks. In addition, the algorithms can also handle a variety of challenging settings, such as noisy communication links, adversarial behavior, and faulty sensors. The algorithms are general and can be applied to applications ranging from environmental monitoring using robotic sensors to the problem of deciding which blogs to read on the Web.

In "Sustainable Biomass Power Plant Location in the Italian Emilia-Romagna Region," Michela Milano, Marco Gavanelli, Massimiliano Cattafi and Paolo Cagnoli propose a mixed integer linear programming model for studying a biomass production-plant location problem. The article emphasizes the importance of considering indirect sources of energy used in biomass production. For example, while government policy may encourage biomass production, market prices may make it more economically attractive to transport biomass from faraway countries, thereby actually increasing global carbon emissions. A sustainable solution must consider such indirect effects. The authors provide a study based on real-world data from a potential biomass power generation facility in the Emilia-Romagna region of Italy.

In "Temporal Data Mining Approaches for Sustainable Chiller Management in Data Centers," Debprakash Patnaik, Manish Marwah, Ratnesh Sharma and Naren Ramakrishnan consider a data-driven approach to improving efficiency of ensembles of chillers in a data center. Data centers have become an object of scorn for environmentalists due to their increasing energy usage and their resulting large carbon footprints. Specifically, chillers have become energy hogs. This article tackles the efficient and sustainable operation of data centers with a novel chiller advisory and management system which uses temporal data mining on multivariate time-series data and characterizes sustainability measures of the patterns mined. The authors use motif mining, association analysis, and dynamic Bayesian network inference to help bridge the gap between low-level, raw, sensor streams, and the high-level operating regions and features needed for an operator to efficiently manage a data center. The developed system has been applied to a production data center managed by HP.

DOI 10.1145/1989734.1989735 http://doi.acm.org/10.1145/1989734.1989735

In "Agent-Based Homeostatic Control for Green Energy in the Smart Grid," by Sarvapali Ramchurn, Perukrishnen Vytelingum, Alex Rogers and Nicholas Jennings, the authors present an autonomous agent-based control mechanism that incorporates consumer preference in energy management demand on an electricity grid. They consider how to fully utilize green energy sources to reduce carbon emissions using a collection of methods, including signals on carbon content that are related to pricing, adaptively-tuned relative supply from the generator side, and the use of battery storage. Their methods consist of novel control mechanisms based on the use of autonomous agents that can both communicate with the grid and optimize their owner's energy consumption to satisfy their preferences.

In "Monitoring Global Forest Cover Using Data Mining," by Varun Mithal, Shyam Boriah, Ashish Garg, Michael Steinbach, Vipin Kumar, Chris Potter, Steve Klooster, and Juan Castilla-Rubio, the authors present an overview of data mining techniques for the characterization of forest cover changes and the discovery of the relationship between the changes and environmental variables. The work makes use of the enhanced vegetation index (EVI) and the fraction of photosynthetically active radiation (FPAR) information. Taking these data as time series, the article proposes a segmentation-based approach and a predictive model-based approach for solving the problem. The authors present an illustrative example of large-scale vegetation disturbances by forest fire, deforestation, and natural disasters, and explain how the change detection techniques can be used. The important role of forests in the climate change debate makes the topic highly relevant and important for a wide range of scientists.

In "MoveMine: Mining Moving Object Data for Discovery of Animal Movement Patterns," by Zhenhui Li, Jiawei Han, Ming Ji, Lu-An Tang, Yintao Yu, Bolin Ding, Jae-Gil Lee, and Roland Kaysr, the authors introduce a data mining system, MoveMine, to learn animal movement patterns based on large volumes of sensor data. The system integrates multiple data mining functions, including sophisticated pattern mining and trajectory analysis. Two movements are highlighted: periodic behavior and swarm patterns. The system can be used for a number of sustainability applications, including ecological studies, urban traffic control, and climatological forecast. A deep understanding of animal movement patterns and changes in the animal world will help the management of our ecosystems.

Finally, in "Spatiotemporal Correlations in Criminal Offense Records," by Jameson Toole, Nathan Eagle, and Joshua B. Plotkin, the authors present a set of analytical tools and test them on criminal offense records in a U.S. city. They demonstrate that their statistical analytical tools are able to identify spatiotemporal patterns on multiple scales and that there is significant correlation on the time scale of weeks. Interestingly, they identify clusters of neighborhoods whose crime rates are affected simultaneously by external forces.

We hope that this special issue will provide readers with a sample of research and applications in the new field of computational sustainability, exemplifying the exciting research currently under way in this field. We would like to thank the authors who submitted articles to this special issue. We are also grateful to the referees for their careful reviews and constructive comments which helped make this special issue a success.

Carla Gomes
Qiang Yang
Guest Editors