Introduction to the Special Issue on Foundations of Social Computing

We are pleased to bring you this special theme issue in the *ACM Transactions on Internet Technology* journal titled “Foundations of Social Computing”. By “social computing”, we broadly refer to computing-supported approaches that facilitate interactions among people and organizations. Social computing has emerged as an exciting multidisciplinary area of research, driven especially by the wealth of easily available information and the success of online social networks and social media.

In spite of the recent focus on social computing and work in related paradigms such as crowdsourcing, collective intelligence, and human computation, we lack a deeper conceptual understanding of social computing. What is *social* in social computing? What is being computed? Who is computing it? Our motivation for doing this theme issue was to attract articles that address these questions, especially via treatments of interaction and information.

After the usual rigorous process of evaluation, we selected the following six articles for publication in this special issue. All the articles focus either on social interaction or user-generated information, the understanding of which, we think, is crucial to answering the foregoing questions.

Baldoni et al., in “A Commitment-Based Infrastructure for Programming Socio-Technical Systems”, propose 2COMM, which is an agent-oriented middleware for realizing cross-organizational systems. The autonomy of the interacting organizations necessitates the decoupling of the design of their internal business processes, that is, their *agents*. To enable such a decoupling, the authors model cross-organizational business processes via commitment-based protocols and show how they can be realized in the 2COMM middleware.

Figueiredo et al., in “On the Dynamics of Social Media Popularity: A YouTube Case Study”, analyze how video popularity evolves over time and the referrers that lead users to videos. The article takes YouTube as a case study and provides an analysis of popularity dynamics and its implications for services such as advertising and search.

Lukasiewicz et al., in “Ontology-Based Query Answering with Group Preferences”, address the problem of finding answers to queries about group preferences. Their approach builds a group preference model over the consequences of a domain ontology from the preferences of individual users. The consequences are probabilistic, thus allowing for the modeling of uncertainty. The article presents algorithms for efficient group queries and a prototype implementation. The feasibility of the approach is demonstrated by applying it to preference models obtained from real users.

Authors' addresses: A. K. Chopra, School of Computing and Communications, Lancaster University, Lancaster LA1 4WA, United Kingdom; email: a.chopra1@lancaster.ac.uk; R. Ali, Bournemouth University, Poole, BH12 5BB, United Kingdom; email: rali@bournemouth.ac.uk; M. Vukovic, IBM Thomas J. Watson Research Center, Yorktown Heights, NY USA; email: maja@us.ibm.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 1533-5399/2014/12-ART22 \$15.00

DOI: <http://dx.doi.org/10.1145/2680536>

Silva et al., in “Revealing the City That We Cannot See”, propose a visualization technique that uses location-sharing systems, such as Foursquare, to understand the human dynamics of cities. The technique can identify similarities and differences of human dynamics across cities by clustering cities according to their citizens’ routines. The article provides insights into the use of social media data to understand city dynamics.

Vosecky et al., in “Integrating Social and Auxiliary Semantics for Multifaceted Topic Modeling in Twitter”, advocate the importance of external information sources, such as Web documents linked from Twitter posts, in understanding microblog contents. They propose a topic modeling approach for Twitter streams that makes use of such sources. They demonstrate the effectiveness of their approach in the context of tweet clustering.

Yang et al., in “SESAME: Mining User Digital Footprints for Fine-Grained Preference-Aware Social Media Search”, extract users’ sentiments about media items from their comments on them in order to deduce their personal preferences. These preferences are then used to generate a media item ranking via a novel algorithm. Their framework is validated using data from Foursquare and YouTube.

We hope you will find much of interest in our selection of articles.

Amit K. Chopra
Lancaster University

Raian Ali
Bournemouth University

Maja Vukovic
IBM Thomas J. Watson Research Center

Guest Editors