

Personalization takes data management to a new frontier where each user is able to get a tailored service according to the personal preferences, behavior, and surrounding context. This special issue includes ten articles that address various aspects of personalization, namely, preference queries, context-aware data management, and recommender systems.

The issue starts by two articles that lay out the foundation of preference queries in relational database systems. Preference queries present a major component in personalized data management where two different users querying a database using the same query may receive different results based on their personal preferences. The first article by Jan Chomicki gives a formal framework in which user preferences are formulated using first-order logic. The article also discusses the use of that logic in preference query evaluation and optimization in relational database systems. The second article by Werner Kießling *et al.* gives an overview of the Preference SQL system; a declarative extension of standard SQL by strict partial order preferences. Preference SQL enables a seamless preference application integration with SQL back-end systems.

The following article by Evaggelia Pitoura *et al.* calls for enhancing preference queries by considering the contextual information. Context may express conditions on situations external to the database or related to the data stored in the database. The article goes on outlining a model for expressing preferences and context to provide a more personalized query answer.

The fourth article in this issue by Georgia Koutrika identifies the differences between plug-in and native realization of preference queries in database management systems. This article aims at showing how tightly preferences are currently coupled with database queries and sharing a vision regarding opportunities and challenges in fully implementing preferences as first-class citizens inside the database engine by changing both the database query model and internal code.

The issue then includes three articles presenting system prototypes for preference and context-aware systems. Yannis Ioannidis *et al.* presents the PAROS system that offers personalized services to its users through a user model, profiling strategies which create instances of personalized user models, and adaptation strategies that adapt the system behavior based on the user profiles. Justin Levandoski *et al.* presents the CareDB system; a context and preference-aware database system that includes a generic and extensible query processing engine, a framework for handling expensive attributes, and a framework for supporting uncertain data. Cristiana Bolchini *et al.* gives an overview of various research related to context modeling and awareness within the Context-ADDICT project.

The last part of this issue includes three articles about recommender systems, which mainly aim to provide content that is likely to interest users, based on current and/or past user behavior. Gloria Chatzopoulou *et al.* presents the QueRIE system that recommends to its users a set of queries that can be posed to the underlying database system. Such system is mostly useful to those users who lack SQL expertise or familiarity with the database schema, e.g., users from the scientific community. Mohammad Khabbaz *et al.* presents the TopRecs⁺ system that extends current recommender systems from only recommending a single item to recommend package of items and to consider user-specified constraints. The issue is then concluded by an article from Sihem Amer-Yahia giving an overview of various projects at Yahoo! Labs in the context of recommendations within web search.

I sincerely hope that you will enjoy reading this issue, and find it interesting and thought-provoking.