Data mining has become a booming topic of research in our community. There are many reasons behind this fact. First, the area is extremely fertile with research problems. Secondly, technology has made it so easy to gather data that now companies and organizations find themselves inundated with data, and eager to extract information from it.

Despite the extraordinary number of papers already published in data mining, we can still argue that the field is in its infancy. Much of the efforts of the database community have been devoted to the issue of finding association rules in datasets, and relatively less attention has been given to other classical data analysis methods such as classification.

This issue of the IEEE Data Engineering Bulletin brings a series of articles that focus on several aspects of data mining. The first paper, written by Surajit Chaudhuri, stresses the fact that our role in the database community must be that of finding techniques to make known data analysis techniques scale over large datasets. He offers guidelines to build **SQL-aware** data mining systems: that is systems that integrate well with an SQL backend.

Clustering, an important technique in data mining is the subject of the next two articles. The second paper in the issue, written by G.D. Ramkumar and Arun Swami, focuses in clustering as an important data mining tool. Clustering is usually performed by using distance functions to measure similarity between data points. However, distance functions are not always available in practice., The authors describe a new method for clustering without distance functions. The third paper in the issue, written by Eui-Hong Han, George Karypis, Vipin Kumar and Bamshad Mobasher presents a method to perform clustering in a high dimensional space.

The next paper in the issue, written by Charu Aggarwal and Philip Yu, surveys the itemset method for the discovery of association rules, discusses its strengths and weaknesses and proposes alternative techniques to the traditional generation methods.

The fourth paper, written by Claudio Bettini, X. Sean Wang and Sushil Jajodia, discusses a particular application domain: mining temporal relationships, such as "which pairs of events occur frequently every week." They introduce heuristics to find these temporal associations efficiently.

Our last paper, by Usama Fayyad, provides an overview and survey of knowledge discovery methods and problems. The paper also provides useful references to the knowledge discovery literature. Usama's paper should be of interest as it represents the views of a member of the knowledge discovery community, views that are not often heard in the database research community.

Data Mining is a fruitful area of research with an abundant wealth of practical applications. We hope you find the articles in this issue as a good and stimulating sample of the research being done currently in this area.