

Chapter 5

Summary

Association rule hiding is a subarea of privacy preserving data mining that focuses on the privacy implications originating from the application of association rule mining to large public databases. In this first part of the book, we provided the basics for the understanding of the problem, which investigates how sensitive association rules can escape the scrutiny of malevolent data miners by modifying certain values in the original database, and presented some related problems on the knowledge hiding thread. Specifically, in the first two chapters we motivated the problem of association rule hiding, presented the necessary background for its understanding and derived the problem statement along two popular variants of the problem: *frequent itemset hiding* and *association rule hiding*. In Chapter 3, we provided a classification of the association rule hiding algorithms to facilitate the organization that we follow for the presentation of the methodologies in the rest of this book. Our proposed taxonomy partitions the association rule hiding methodologies along four orthogonal directions based on the employed hiding strategy, the data modification strategy, the number of rules that are concurrently hidden, and the nature of the algorithm. Elaborating on the last direction, we identified three classes of association rule hiding approaches, namely *heuristic-based*, *border-based* and *exact* approaches, and discussed the differences among them. Last, in Chapter 4 we examined the problem of knowledge hiding in the related research areas of clustering, classification and sequence mining. For each of these areas we briefly discussed some state-of-the-art approaches that have been proposed.