Abstract

A network provides powerful means of representing complex relationships between entities by abstracting entities as vertices, and relationships as edges connecting vertices in a graph. Beyond the presence or absence of relationships, a network may contain additional information that can be attributed to the entities and their relationships. Attaching these additional attribute data to the corresponding vertices and edges yields an attributed graph. Moreover, in the majority of real-world applications, such as online social networks, financial networks and transactional networks, relationships between entities evolve over time. Change detection in dynamic attributed networks is an important problem in many areas, such as fraud detection, cyber intrusion detection, and health care monitoring. It is a challenging problem because it involves a time sequence of attributed graphs, each of which is usually very large and can contain many attributes attached to the vertices and edges, resulting in a complex, high-dimensional mathematical object. In this survey we provide an overview of some of the existing change detection methods that utilize attribute information. We categorize these methods based on the levels of structure in the graph that are exploited to detect changes. These levels are vertices, edges, subgraphs, communities, and the overall graph. We focus attention on the strengths and weaknesses of these methods, including their performance and scalability. Furthermore, we discuss some publicly available dynamic network datasets and give a brief overview of models to generate dynamic attributed networks. Finally, we discuss the limitations of existing approaches identifying key areas for future research.