ABSTRACT

A network is fundamentally composed of nodes, which represent the actors within it, and edges, which represent the relationships between these actors. A community in a network is a collection of entities that share a significant number of closely related interactions. Detecting these communities is a critical area of research within the field of data mining, as evidenced by extensive literature on the subject. Social networks that evolve over time present unique challenges; they demand algorithms that can accommodate temporal changes while operating within limited resource constraints. Static algorithms typically fall short in such dynamic environments, where there is a need for continuous updates to community structures, efficient computations, and immediate results.

Dynamic networks introduce new challenges because they are characterized by their ability to evolve, necessitating models that account for both their structural and temporal patterns. Research indicates that these networks, which reflect real-world scenarios, are organized into communities that frequently change over time. These communities can represent tightly-knit groups such as linked online documents, social circles, or even consumer purchasing patterns.

This thesis presents three novel techniques for the identification of communities within dynamic social networks. Prior to delving into the algorithms, an exposition on a network model tailored for dynamic networks is presented. The network is conceptualized as an evolving graph, depicted as a stream of events. Each event signifies a modification within the network, accompanied by corresponding timestamps.

The first method utilizes a tree-based algorithm, TCD2, which takes advantage of social network properties like connectedness and influence to detect communities. Nodes are categorized in two roles: cluster heads and members. Further they are organized in the tree where the root node stores cluster heads and their corresponding members are stored at

subsequent levels. TCD2 employs a tree structure to keep track of the network's changing community structures.

Subsequently, the thesis details a multi-objective optimization approach for community detection in dynamic networks, viewing each network as a temporal sequence of events characterized by the introduction of new edges. This method updates the community affiliations of nodes based on the recency of interactions, utilizing the Pareto front principle to determine the most fitting community.

Finally, a fuzzy logic approach to community detection in dynamic networks is presented, featuring a novel membership function based on insights from the literature. This function is critical in assigning nodes to communities at the decision-making step, enabling the algorithm to generate distinct communities at any given point in time.

Each of the aforementioned techniques has been thoroughly evaluated against the recognized metrics from the literature. Both quality and accuracy metrics are used in the evaluation of proposed algorithms. The experimental results on real world social networks along with synthetic networks validate the performance of algorithms. An empirical analysis has also been included that compares the effectiveness of the proposed algorithms with the current state-of-the-art algorithms, demonstrating their potential and performance. Each algorithm has the capability to delineate the community structure of a given dynamic network at a specific point in time.