

Interested Followers? Using latent community structures to understand product diffusion patterns

Arash Yazdiha, Bas Donkers, Dennis Fok

Abstract

The structure of social networks plays a large role in the diffusion of new products and the success of social media campaigns. It is common practice to characterize influential members using their connection count. However, identical treatment of each connection is problematic as some connections may be important for one type of behavior but irrelevant for another. Diffusion literature suggests that considering multiple underlying networks aids in determining the relevance of specific connections. Such underlying networks can be thought of as different communities of people. However, direct observation of communities, or connections over multiple networks, is often not possible.

We propose a scalable stochastic variational inference methodology to uncover overlapping communities from a single observed network. Knowledge on the probabilistic community memberships of the individuals in the network can be used to measure the relevance of their connections. We apply our community membership model to a US social network that involves product adoption. We find multiple communities within the network, which allows us to define the degree to which friends have shared interests. We then show that friends who have stronger shared interests are more likely to adopt a product than friends with other interests. While it is already known that friends are more likely to adopt than non-friends, network structure is also informative about which friends are most likely to adopt a product.

Keywords: social networks, community structure, connection relevance, scalable inference, network metrics, product adoption