

Abstract: There has been an exponential growth in the amount of multimedia content that is being generated on multimedia social networks such as YouTube, Flickr, Twitter etc. Unlike traditional multimedia content, content shared on these networks is annotated with a rich set of network and social metadata features which are generated as a byproduct of content creation and sharing on these networks. This social and network metadata (which we refer to as social signals) includes tags, description, geo-location, subscription, community membership, informal interactions via comments etc. As part of this thesis we explore how applications such as community detection, link prediction and image search can leverage social signals to their advantage. Sparsity, ambiguity and heterogeneity of these features pose challenges which are addressed using a wide array of machine learning methods such as probabilistic modelling, multi-task/multi-view learning and multiple kernel methods. Experiments on real world data sets confirm our hypothesis that incorporating social signals improves results for different multimedia tasks.