

ABSTRACT: We consider the problem of tag prediction in collaborative tagging systems where users share and annotate resources on the Web. We put forward HAMLET, a novel approach to automatically propagate tags along the edges of a graph which relates similar documents. We identify the core principles underlying tag propagation for which we derive suitable scoring models combined in one overall ranking formula. Leveraging these scores, we present an efficient top-k tag selection algorithm that infers additional tags by carefully inspecting neighbors in the document graph. Experiments using real-world data demonstrate the viability of our approach in large-scale environments where tags are scarce.