

Kernel methods for conditional ranking from relational data

Willem Waegeman

(Department of Applied Mathematics, Biometrics and Process Control,
Ghent University), February 5, 2010

In domains like bioinformatics, information retrieval and social networks analysis, one can find learning tasks where the goal consists of inferring a ranking of objects, conditioned on a particular target object. I will present a general kernel framework for learning conditional rankings from various types of relational data, where rankings can be conditioned on unseen data objects. Conditional ranking from symmetric and reciprocal relations can in this framework be treated as two important special cases, for which the task simplifies to a traditional ranking task when the relations satisfy certain properties. Furthermore, I will discuss an efficient regularized least-squares algorithm for conditional ranking and relational learning, in which domain knowledge about the inferred relations can easily be included by means of specific kernel functions. Experiments on synthetic and real-world data illustrate that such an approach delivers state-of-the-art performance in terms of predictive power and computational complexity.