Mach Learn (2012) 89:233–255
DOI 10.1007/s10994-012-5309-4

Relational networks of conditional preferences

Frédéric Koriche

Received: 30 January 2012 / Revised: 6 May 2012 / Accepted: 14 June 2012 /
Published online: 18 July 2012

Abstract Much like relational probabilistic models, the need for relational preference models naturally arises in real-world applications involving multiple, heterogeneous, and richly interconnected objects. On the one hand, relational preferences should be represented into statements which are natural for human users to express. On the other hand, relational preference models should be endowed with a structure that supports tractable forms of reasoning and learning. Based on these criteria, this paper introduces the framework of *relational conditional preference networks (RCP-nets)*, that maintains the spirit of the popular "CP-nets" by expressing relational preferences in a natural way using the *ceteris paribus* semantics. We show that *acyclic* RCP-nets support tractable inference for optimization and ranking tasks. In addition, we show that in the online learning model, *tree-structured* RCP-nets (with bipartite orderings) are efficiently learnable from both optimization tasks and ranking tasks, using linear loss functions. Our results are corroborated by experiments on a large-scale movie recommendation dataset.

Keywords Conditional preferences · Relational networks · Preference optimization · Preference ranking · Online learning

1 Introduction

A recurrent issue in AI is the development of rational agents capable of tailoring their actions and recommendations to the preferences of human users. The spectrum of applications that depend on this ability is extremely wide, ranging from adaptive interfaces and configuration software, to recommender systems and group decision-making (Brafman and Domshlak 2009). In a nutshell, the crucial ingredients for addressing this issue are *representation*, *reasoning* and *learning*. In complex domains, we need a representation that offers a compact encoding of preference relations defined over large outcome spaces. We also need to be

Editors: S. Muggleton and J. Chen.

F. Koriche (✉)
LIRMM, CNRS UMR 5506, Université Montpellier II, Montpellier Cedex 5, France
e-mail: Frederic.Koriche@lirmm.fr