Expert Systems with Applications 42 (2015) 2884–2906

Contents lists available at ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

Utilising social recommendation for decision-making in distributed multi-agent systems

Brendan Neville [a,*], Maria Fasli [b], Jeremy Pitt [c]

[a] University of Southampton, United Kingdom
[b] University of Essex, United Kingdom
[c] Imperial College London, United Kingdom

ARTICLE INFO

Article history:
Available online 2 December 2014

Keywords:
Multi-agent systems
Self-organising software agents
Engineering emergent behaviour

ABSTRACT

Open multi-agent systems are typically formed from heterogeneous peers operating in a decentralised manner. Hence, their constituent agents must evaluate possible actions and opportunities based on local, subjective knowledge. When agents have insufficient personal experience, they may inevitably rely on their social connections to act as a source of relevant information or recommendations. We describe an agent-mediated electronic market for investigating social interaction within the context of evolving heterogeneous distributed networks. In our scenario, consumers look for appropriate services and this service choice is informed via peer recommendations. We define two alternative algorithms for selecting peers based on perceived similarity and we evaluate them on their ability to organise an overlay network such that it acts as a passive filter, tailoring the information that agents use to select services in the market. We use this scenario to explore the link between the peer selection algorithms and the emergent network topologies, as well as the impact of the peer selection algorithm on the agents' performance in choosing services based on peer recommendations. Our simulation results demonstrate a qualitative difference in the behaviour of the algorithms, with optimal algorithm selection relying on information regarding the preferences of the wider population of agents.

1. Introduction

Open multi-agent systems are typically formed from heterogeneous peers operating in a decentralised manner. For agent-mediated service[1] markets, this means service providers and service consumers seeking to maximise their owners' utility through trading. In order to achieve this, consumer agents are required to exhibit autonomy in discovering and selecting services, and producers must market their new services to the most relevant consumers. Dynamic environments, especially where services change over time, can undermine an agent's[2] ability to learn from their own experiences, making it necessary for them to cooperate and leverage their larger collective experience via recommendations, in order to make more informed decisions. While this lowers the agents' cost of learning, it introduces new challenges arising from component heterogeneity and scalability. In systems where the properties/requirements of agents are homogeneous, learning from peer experiences through recommendation is trivial. However for heterogeneous systems, while the agents' intentions may be cooperative, it would be detrimental to their performance to select services based upon the recommendations of peers which are not 'similar' with respect to their goals, attitudes, capabilities and requirements. Therefore it is essential that agents discover similar peers with which to interact, in order to ensure the subset of opinions on which they base decisions is relevant.

This is similar to making recommendations for consumer goods via collaborative filtering (Goldberg, Nichols, Oki, & Terry, 1992; Burke, 2002) techniques which rely on identifying similarities between users. In its most simple form, collaborative filtering is a two stage process of creating similarity estimates, before using these as weights when aggregating recommendations for the remaining items. The accuracy of these estimates can be improved by the addition of an intermediate step (Shardanand & Maes, 1995; Herlocker, Konstan, Borchers, & Riedl, 1999) that selects a subset of closely related users (their clique or neighbourhood) to act as the source of recommendations, as opposed to using recommendations from every user in the system. However, these systems rely upon

* Corresponding author.
E-mail addresses: b.j.neville@soton.ac.uk (B. Neville), mfasli@essex.ac.uk (M. Fasli), j.pitt@imperial.ac.uk (J. Pitt).

[1] We use the term 'service' in a generic way to describe products, services, etc.
[2] For brevity, where we use agents without a prefix of consumer or producer, we are referring to the consumer agents.

http://dx.doi.org/10.1016/j.eswa.2014.10.029

centralisation to provide a global perspective from which to calculate peer similarities; this centralisation is not always practical and has been identified as a potential weakness with respect to scalability and fault tolerance. More recently, and with the huge growth of social networks, social recommendation approaches based on collaborative filtering algorithms have been applied on the distributed setting of social networks (Tang, Hu, & Liu, 2013).

Human social systems are necessarily decentralised and feature aspects of self-organisation and adaptation, making them inherently scalable and robust. Therefore it has been proposed that networked applications should be modelled as agent societies (Artikis, Pitt, & Sergot, 2002; Sierra, Rodríguez-Aguilar, Noriega, Esteva, & Arcos, 2004) complete with social relations between components (Conte & Dellarocas, 2002). These social relations form an agent social network, consisting of localised neighbourhoods akin to the cliques formed in recommender systems.

Empirical studies of human social networks (Travers, Milgram, Travers, & Milgram, 1969; Granovetter, 1973; Friedkin, 1982) have shown that the overall network structure (the topology) and strength of peer relationships, has a significant effect on how people discover and exploit useful information in day to day tasks. It has also been demonstrated that these relationships tend to form between people who are similar with respect to their values and characteristics (Lazarsfeld & Merton, 1954); this tendency is termed homophily. Researchers in opinion dynamics (DeGroot, 1974) have developed top-down mathematical models of how social influence within groups of peers can affect the beliefs or opinions of peers over time. These have been expanded to include the notion that peer similarity is required for social influence to occur (Krause, 2000; Deffuant, Neau, Amblard, & Weisbuch, 2000). These social network dynamics will be present in distributed systems; our understanding of and ability to harness them is central to our ability to design reliable and predictable networked applications.

While these mathematical models and their corresponding top-down simulations are useful as an aid to intuition (Axelrod, 1997), they do not fully capture how emergent properties of distributed systems are the result of the individual interactions of heterogeneous components. Specifically, how the design of the interaction protocols, agent learning and peer selection algorithms will affect an autonomous agent's behaviour and performance. This motivates the need for an agent-based model that incorporates the features of homophily and social influence into a form that can be used as a specification for integrating recommendation into distributed applications, as well as the basis for further study of social dynamics.

In this article, we describe a simulation model for investigating recommendation within the context of an evolving network of heterogeneous service provider and consumer agents in an electronic market. We present an agent-based model for recommendation and choice based on homophilic neighbourhood selection that controls the diffusion of opinions in the global network leading to localised/tailored perception of the expected utility from offered services. We define two alternative algorithms for selecting peers based on perceived similarity and we evaluate them on their ability to organise an overlay network such that it acts as a passive filter, thus, tailoring the information that agents use to select services in the market. Our simulation experiments provide insights into the link between the choice and parameterisation of the neighbourhood selection algorithm, the emergent network topologies and the effect of the selection algorithm on the agents' performance in choosing services in the market.

The remainder of this paper is structured as follows: Section 2 provides the background to the current work in the areas of collaborative filtering and opinion dynamics while it also points to the differences between prior work and the work presented in this paper. Section 3 provides the description of the abstract scenario that we will be using and Section 4 describes the initial network formation process. The two types of agents in our scenario and their functionality are explained in Section 5. Finally, Section 6 presents a series of experiments and the ensuing results. The paper closes with a summary, conclusions and discussion of future work.

2. Background

Sharing recommendations and opinions among users has been investigated both in the context of recommender systems and in particular, collaborative-based filtering systems and in social network formation and opinion dynamics. In the following subsections, we present related work in these areas and we also describe the links and differences between these works and the work presented hereafter.

2.1. Collaborative-filtering based recommendations

In the collaborative filtering domain, users in a system submit opinions (typically in the form of numerical ratings within a specific range) on items that can be books, movies or other objects. Such systems are inherently centralised in that the system collects and stores all such opinions and ratings. The purpose of such a system is to provide a recommendation on an item to a target user by evaluating this user's pairwise similarity with other users in the system, based on their previously expressed opinions on items. This is done on the assumption that similar users will share tastes and preferences.

In memory-based collaborative filtering approaches the similarity between users is calculated based on their ratings of items using some measure such as cosine similarity or Pearson correlation. The neighbourhood selection algorithm chooses a subset of highly similar users to act as a source of recommendations from which to suggest new items. This has been shown to improve the quality of recommendations in comparison to just using weighted values from all users (Shardanand & Maes, 1995), and it also improves the scalability of collaborative filtering by reducing the number of comparisons and calculations required; such centralised collaborative filtering systems may have thousands of users rating hundreds of items each, hence making pairwise comparisons between users and items is a very time-consuming task.

Herlocker et al. (1999) investigated the algorithms used for each stage of collaborative filtering in the context of a centralised recommendation task. Of particular relevance to this paper, they compare two methods for neighbourhood selection, best-*n*-neighbours (analogous to our *TopX* algorithm) and correlation thresholding (analogous to our *Threshold* algorithm). Based on decision making accuracy and system coverage, they conclude that the best-*n*-neighbours is the best method to select similar peers.

There has been a significant body of work in collaborative filtering algorithms with the most recent developments being the development of model-based approaches and in particular the use of matrix factorization techniques to learn from data (Koren, Bell, & Volinsky, 2009). Such methods, which were originally applied in centralised settings, have been shown to yield superior results to other collaborative filtering techniques (Koren, 2009).

However, these methods are deployed in a centralised fashion. Using centralised services in open agent systems has a number of disadvantages, e.g. trust, fault tolerance, scalability. The alternative is to adapt the neighbourhood selection methods seen in collaborative filtering to fit with a distributed architecture, whereby, the chosen cliques can be represented as a web of social relationships between peers.