Speaking of Relations: Connecting Statistical Relational Learning and Multi-Agent Systems

Marie desJardins MARIEDJ@CS.UMBC.EDU
Matthew E. Gaston MGASTO1@CS.UMBC.EDU
University of Maryland Baltimore County, Computer Science and Electrical Engineering Department,
1000 Hilltop Circle, Baltimore, MD 21250 USA

Abstract

We discuss the relationship between the fields of statistical relational learning (SRL) and multi-agent systems (MAS). We identify a number of SRL research problems that have analogies in MAS research, and vice-versa, and suggest how research from each area can be leveraged to provide solutions for the other.

1. Introduction

Statistical relational learning (SRL) refers to the problem of discovering patterns in complex relational networks. The vertices in these networks correspond to objects that are characterized by a set of attributes, and that are connected by links representing a variety of relationships among the objects. SRL approaches typically use centralized algorithms, which have a global view of the data.

Multi-agent systems (MAS) research focuses on the behaviors of complex "societies" of agents, who have a variety of skills and characteristics. MAS approaches are used in domains that do not admit of centralized solutions for a variety of reasons, including privacy, authority, communication restrictions, and localized availability of information.

A subfield of MAS, sometimes referred to as *networked multi-agent systems*, examines the case where agents are connected through some set of relationships, which may be geographic, trust-based, communication-based, or stemming from the limited knowledge that agents have of other agents in a large community. In this situation, agents can be thought of as having a social network: their behavior and interaction with other agents are influenced by their position in the network, and in turn, they may modify the network by adding or dropping connections with other agents.

Many problems in (networked) MAS can be seen as distributed forms of SRL problems. Conversely, SRL problems may benefit from approaches and formalisms that have been developed by the MAS community.

In this position paper, we use the categorization of link mining tasks provided by Getoor and Diehl (2006) to identify relevant problems and research in MAS that may be useful in solving SRL problems (and vice versa). It is our hope that this analysis will lead to more sharing of knowledge and ideas across these two largely disparate research communities.

2. Link Mining Analogies to MAS

Getoor and Diehl (2006) identify eight main problem areas in link mining: link-based object ranking, link-based object classification, group detection, entity resolution, link prediction, subgraph discovery, graph classification, and generative models for graphs. We follow this organization here, discussing relevant research and problem areas in MAS corresponding to each of the link mining topics.

2.1. Link-Based Object Ranking

The goal of link-based object ranking is to prioritize the objects in a graph based on its link structure. Link-based object ranking is a fundamental and widely used technique in link mining. Likewise, there are many multi-agent system analogies for link-based object ranking.

The emergence of social protocols and social conventions (Delgado, 2002) in multi-agent systems give rise to rankings of the agents in the system. Protocols and conventions may dictate the order in which agents accomplish tasks or a precedence ordering for the consumption of a shared resource. Precedence orderings

Presented at the ICML Workshop on Open Problems in Statistical Relational Learning, Pittsburgh, PA, 2006.

are also central to many distributed scheduling and execution problems (Horling et al., 2001). Issues related to social protocols and social conventions are found in both networked and non-networked multi-agent systems. Like many of the analogies between link mining and multi-agent systems, the emergence of social protocols is typically distributed, as opposed to the centralized computation of a ranking among the objects in a graph in link mining.

2.2. Link-Based Object Classification

In link-based object classification, objects are given labels based on their relative positions in the graph. Two closely related topics in multi-agent systems are role (or task) allocation (Walsh & Wellman, 1998; Nair et al., 2003) and reputation formation (Pujol et al., 2002). In role allocation, agents in an organization attempt to identify a utility-maximizing role to play in relationship to the other agents. In essence, distributed role allocation is the self-classification of each of the agents in a multi-agent system, where the labels are roles. Conversely, reputation formation can be thought of as each agent classifying the other agents in the organization in order to determine patterns of interaction, trustworthiness, or reliability.

It may also be possible to apply link-based object classification more directly to multi-agent systems. In many applications, agents in an organization may desire to predict the type, group membership, characteristics, or quality of other agents in the organization. As in link-based object classification, agents could use both interconnection patterns among the agents and observable features of the other agents in order to classify them.

2.3. Group Detection

In link mining, group detection is the process of clustering objects together based on similar patterns of connectivity and similar characteristics. In multi-agent systems, it is desirable to organize large groups of agents into coalitions or sub-groups in order to more efficiently govern coordination and cooperation (Abdallah & Lesser, 2004). These coalitions are typically formed by grouping agents with similar characteristics or interests together, either through a centralized mechanism or using distributed computation. A closely related topic is team formation (Nair et al., 2002; Gaston & desJardins, 2003), although in team formation agents typically form teams with complementary skills, rather than duplicative or similar characteristics.

Similar to both team formation and coalition formation is the phenomenon of emergent communities in multi-agent systems. The most obvious application of emergent communities is the distributed formation of peer-to-peer information retrieval systems (Yolum & Singh, 2003). In this application domain, agents use explicit information about other agents' interests and qualifications to manage and manipulate their peer relationships.

2.4. Entity Resolution

Entity resolution refers to the problem of identifying sets of nodes within a graph that actually refer to the same object. A similar problem arises in multi-agent systems, in which trust, reputation, and authentication are considerations. Ramchurn et al. (2004) mention but do not suggest any solutions for the problem of identifying "repeat offenders"—that is, individuals who have low levels of trust who leave a community, then rejoin with a new identity. A related problem is when one agent poses as multiple agents (a real-world example being individuals with multiple logins on ebay, posting positive feedback, or "bidding up" their own items). Although these are potentially important problems, the MAS community has not yet developed solutions; work on entity resolution from the database community may provide a source of ideas.

Yolum and Singh (2003) draw on work from link analysis and social network analysis to discover interest groups and communities in peer-to-peer systems. This is a slightly different problem from entity resolution—the equivalence classes being discovered are not of identical individuals, but of groups of individuals with similar properties.

2.5. Link Prediction

Link prediction is the task of inferring the existence of a link (relationship) in the graph that was not previously known. This class of problems appears in several different contexts in MAS. Norman et al. (2004) discuss the problem of "virtual organization formation," in which agents form subgroups (i.e., discover or create new links) within a larger community. Gaston and desJardins's work on agent-organized networks studies the dynamics of agent communities, in which agents locally add and remove links from the social network (Gaston & desJardins, 2005a).

Reputation networks (Pujol et al., 2002; Ramchurn et al., 2004) typically focus on the inverse problem—that of using the existing link structure (social network) to infer the reputation of agents—but these models could perhaps also be leveraged to infer relationships between agents.

A potentially interesting application of link prediction to MAS, which to our knowledge has not been explored, is that of inferring relationships between other agents. In particular, this might enable agents within a community to uncover collusion between other agents.

2.6. Subgraph Discovery

The problem of "sub-group" formation (e.g., team formation or coalition formation (Nair et al., 2002; Gaston & desJardins, 2003; Abdallah & Lesser, 2004)) is of central concern in MAS. In such settings, agents with particular skills, or resources, must form groups in order to accomplish some joint task (i.e., a task that has a specific set of skill requirements). When the agents are embedded in an organizational network structure, skill-constrained sub-group formation is very similar to subgraph discovery in networks. At first the problems appear to be different in that subgroup formation focuses on attributes of the agents, whereas subgraph discovery focuses on local patterns in the network. However, the two problems become closely related when the sub-group formation process in a multi-agent system is restricted by an interaction topology (Gaston & desJardins, 2003).

2.7. Graph Classification

Graph classification is the problem of assigning an entire graph to a specific category. This SRL problem has a direct analogy to organizational design for networked MAS (Horling & Lesser, 2004; Gaston & desJardins, 2005c). It is well known that the organizational network structure of a MAS directly affects the collective performance of that system. Therefore, it is desirable to identify (at least) two classes of networks for specific MAS environments: efficient and non-efficient network structures. If graph classification could be used to automatically recognize these classes, it would allow MAS designers to create efficient interaction topologies for specific MAS domains.

The problem of graph classification also brings about an interesting new problem relevant to MAS: determining when to join (or separate from) an open networked multi-agent system. If agents have a global view of the organizational network, they could use graph classification methods to determine when to join a specific networked organization. One application of this would be automated support for determining profitable participation in supply chain networks or other market-oriented environments.

2.8. Generative Models for Graphs

The use of generative models in designing and understanding MAS is becoming increasingly important. One aspect of multi-agent systems is to understand the effects of real-world network structures on organizational performance, for which it is necessary to be able to generate various "life-like" network topologies to assess (Gaston & desJardins, 2005c; Delgado, 2002). Embedding agents in various realistic network topologies provides some evidence regarding the properties of network structures that promote high performance in specific domains.

Generating and modeling realistic network structures also helps understand the behavior of so-called agent-organized networks. Agent-organized networks are dynamic multi-agent networks that evolve based on the decisions of individual agents (Gaston & desJardins, 2005a; Gaston & desJardins, 2005b). Comparing the networks as they evolve with models of realistic network structures provides insight into how and why agent-organized networks are evolving.

3. Conclusions

SRL and MAS, particularly networked MAS, have many areas in common, although the two research communities have approached these problems from very different perspectives.

Some of the specific ways in which results from MAS research might contribute to SRL solutions include the use of emergent rankings derived from social protocols to perform link-based object ranking, applying MAS coalition and team formation methods to the problems of group detection and subgraph discovery in SRL, and using techniques from organization formation and agent-organized networks to perform link prediction. MAS research has, of necessity, focused on distributed methods that may be useful for scaling up SRL to large, complex networks.

Promising avenues for applying SRL results to MAS problems include using link-based object classification to facilitate role allocation and group membership identification, applying entity resolution techniques to identify repeat offenders and multiple identities in MAS communities, using link prediction methods to recognize collusion among agents, and applying graph classification to discover efficient networks for MAS problems.

References

Abdallah, S., & Lesser, V. (2004). Organization-based cooperative coalition formation. *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT).*

Delgado, J. (2002). Emergence of social conventions in complex networks. *Artificial Intelligence*, *141*, 171–185.

Gaston, M., & desJardins, M. (2003). Team formation in complex networks. *Proceedings of the 1st NAACSOS Conference.*

Gaston, M., & desJardins, M. (2005a). Agent-organized networks for dynamic team formation. *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems.*

Gaston, M., & desJardins, M. (2005b). Agent-organized networks for multi-agent production and exchange. *Proceedings of the Twentieth National Conference on Artificial Intelligence.*

Gaston, M., & desJardins, M. (2005c). Social network structures and their impact on multi-agent system dynamics. *Proceedings of the 18th International Florida Artificial Intelligence Research Society Conference.*

Getoor, L., & Diehl, C. (2006). Link mining: A survey. *SIGKDD Explorations*, *7*, 3–12.

Horling, B., Benyo, B., & Lesser, V. (2001). Using self diagnosis to adapt organizational structure. *5th International Conference on Autonomous Agents.*

Horling, B., & Lesser, V. (2004). *A survey of multiagent organizational paradigms* (Technical Report CS 04-45). University of Massachusetts.

Nair, R., Tambe, M., & Marsella, S. (2002). Team formation for reformation. *Proceedings of the AAAI Spring Symposium on Intelligent Distributed and Embedded Systems.*

Nair, R., Tambe, M., & Marsella, S. (2003). Role allocation and reallocation in multiagent teams: Towards a practical analysis. *Proceedings of the Second International Joint Conference on Agents and Multiagent Systems (AAMAS).*

Norman, T., Preece, A., Chalmers, S., Jennings, N., Luck, M., Dang, V., Nguyen, T., Deora, V., Shao, J., Gray, W. A., & Fiddian, N. (2004). Agent-based formation of virtual organizations. *Knowledge-Based Systems*, *17*, 103–111.

Pujol, J. M., Sanguesa, R., & Delgado, J. (2002). Extracting reputation in multi agent systems by means of social network topology. *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems* (pp. 467–474).

Ramchurn, S., Huynh, D., & Jennings, N. (2004). Trust in multi-agent systems. *The Knowledge Engineering Review*, *19*, 1–25.

Walsh, W. E., & Wellman, M. P. (1998). A market protocol for decentralized task allocation. *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98).*

Yolum, P., & Singh, M. P. (2003). Emergent personalized communities in referral networks. *Proceedings of the IJCAI Workshop on Intelligent Techniques for Web Personalization (ITWP).*