Chapter 7
Bayesian Agent Adaptation in Complex Dynamic Systems

Mair Allen-Williams
University of Southampton, UK

Nicholas R Jennings
University of Southampton, UK

ABSTRACT

Multi-agent systems draw together a number of significant trends in modern technology: ubiquity, decentralization, openness, dynamism and uncertainty. As work in these fields develops, such systems face increasing challenges. Two particular challenges are decision making in uncertain and partially-observable environments, and coordination with other agents in such environments. Although uncertainty and coordination have been tackled as separate problems, formal models for an integrated approach are typically restricted to simple classes of problem and are not scalable to problems with many agents and millions of states. We improve on these approaches by extending a principled Bayesian model into more challenging domains, using heuristics and exploiting domain knowledge in order to make approximate solutions tractable. We show the effectiveness of our approach applied to an ambulance coordination problem inspired by the Robocup Rescue system.

INTRODUCTION

As computing power and ubiquity increase, the use of multi-agent technology in complex distributed systems is becoming more widespread. Consequently, the scalability of such systems is becoming increasingly important. Furthermore, the inherent dynamism in many of these problems calls for timely online responses, rather than the offline computation of strategies. As an example, consider a street taxi business. Between fares, the taxis roam the streets looking for customers. If taxis are able to automatically share their locations using GPS and a dashboard display, they can attempt to spread out over different areas. Now, for multi-agent problems in which the complete state is not observed by any one agent, re-

cent work has advanced the state of the art for finding offline solutions in networks of stationary agents, with solutions in systems containing at least fifteen agents (Marecki et al., 2008). Building on this, in this chapter we describe a related online approach which is suitable to complex dynamic processes with mobile agents, also scalable into tens of agents. In order to provide a focus, and as a motivation for this work, we will consider the disaster response domain. Disaster scenarios form rich grounds for multi-agent distributed problem solving, allowing us to explore several features of complex multi-agent problems. While there are many characteristics which may be present in disaster scenarios, we will find that there are two common themes: uncertainty, and coordination.

The first of these, uncertainty, may concern the environment ("What's going on?") and the agent's position in the environment ("Where am I?"); it may be about any other agents which might exist in the environment ("Who else is around? Where are they?") and their behavior ("What are they going to do?"). In these uncertain situations, each agent must do some form of discovery to determine the essential characteristics of the situation, including the agent's collaborators, before and alongside directly working to achieve its goals. This discovery phase in a multi-agent system is tightly linked with the presence of other agents in the system. As well as determining which other agents are present, agents may be able to cooperate to search over different regions, sharing information with each other as appropriate.

In addition to explicitly sharing information, observing the behavior of the other agents allows an autonomous agent to make inferences about the system. For example, in a scenario involving a burning building, a rational agent will not enter the building. Out of the disaster domain, consider perhaps a car manufacturing company, receiving orders from a regular customer base of car dealerships, able to make judgements about local economics based on the orders. Beyond discovery, there will continue to be interaction between the agents in a multi-agent system, whether explicit via communications and negotiations, or implicit through activity. Achieving some subgoals may involve a collaboration between several agents, as in a rescue operation where two ambulance members are required to carry a stretcher, or a car manufacturer where different parts are manufactured in different local factories.

Now, this general problem of taking others into account, coordination, is the second key issue we have identified for multi-agent systems. In uncertain, changing or open systems, fixed protocols for coordination must function against a background in which agents are not fully aware of the situation; their environment, the resources available to them, or the behavior of the other agents. For example, a particular car part manufacturer may be manufacturing parts in assorted colours without necessarily knowing what orders are coming in or whether a particular colour is newly in vogue in one of the key towns supplied by the factories. The negotiation of coordinated behavior in such systems is intertwined with the discovery phase, as agents interact with one another, perhaps cooperating to determine properties of the situation. Another example might be of a team of milkmen, needing to prepare for adjustments to their regular standing orders: they will expect that in the summer customers are more likely to go on holiday; they may be able to make inferences about school holidays within a particular catchment area, and the teams may be able to compare notes when they meet at the depot.

In order to achieve such a comprehensive model for planning and acting, we have built on existing techniques for coordinated decision making under uncertainty in dynamic systems. We have developed an algorithm which explicitly models other agents, demonstrating a principled approach to coordinated behavior in uncertain and partially-observable multi-agent systems. Over the next sections, we provide the background to this work and motivate a coordinated approach using finite state machines to model agent behavior. We then go on to describe this approach in detail. In order to validate our model we test it on a problem taken from the disaster response domain, describing the problem and comparing the

35 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/bayesian-agent-adaptation-complex-dynamic/36569

Related Content

Decomposition Procedure for Solving NLP and QP Problems based on Lagrange and Sander's Method
H. K. Das (2016). *International Journal of Operations Research and Information Systems (pp. 67-93).*
www.irma-international.org/article/decomposition-procedure-for-solving-nlp-and-qp-problems-based-on-lagrange-and-sanders-method/163655

Towards a Model of Sustainability Capabilities and Competitive Advantage: Empirical Evidence from India
Santanu Mandal (2014). *International Journal of Applied Management Sciences and Engineering (pp. 41-57).*
www.irma-international.org/article/towards-a-model-of-sustainability-capabilities-and-competitive-advantage/106839

Ownership, Incentives and Regulation of CCP Risks
Louise Carter, Jennifer Hancockand Mark Manning (2018). *Global Business Expansion: Concepts, Methodologies, Tools, and Applications (pp. 489-521).*
www.irma-international.org/chapter/ownership-incentives-and-regulation-of-ccp-risks/202233

Reference Model Management
Oliver Thomas (2007). *Reference Modeling for Business Systems Analysis (pp. 288-309).*
www.irma-international.org/chapter/reference-model-management/28364

Developing Your Strategy
(2019). *Strategic Management of Business-Critical Information Assets (pp. 63-74).*
www.irma-international.org/chapter/developing-your-strategy/225444