Three Myths about Roles

Myriam Abramson
Naval Research Laboratory
4555 Overlook Ave., SE
Washington, DC
myriam.abramson@nrl.navy.mil

Abstract

In multi-agent systems, roles provide an apriori decomposition of the task based on the interactions of the agents. Roles are instantiated by the agents' actions and/or messages. Role allocation has emerged as a key computational tool in the coordination of intelligent agents who are flexible enough to take on any role. It is in this context that some rethinking about roles and their representation is needed. Illustrations in the game of Go, a game of strategy, highlight the key issues.

Introduction

Roles have emerged as a powerful computational abstraction for the coordination of multi-agent systems. Roles are broadly defined here as encapsulated behaviors used to solve a task requiring no learning or planning from first principles and enabling the reusability of a solution in multi-agent systems. In this context, a strategy can be viewed as a composition of roles. Roles, as abstract entities, can be assumed and/or assigned and therefore lend themselves both to self-organization and mixed strategies where human-in-the-loop interactions require a higher level of granularity in the articulation of goals. In practice, the distinction between an agent and its role mitigates the usefulness of this conceptual abstraction.

It was the game of Go[1] that led John Nash to game theory. Indeed, the interaction between the stones in Go has a close resemblance to the interaction of intelligent agents due to the absence of constraints in the type of move, including suicide. Go is a two-player game played with stones of different colors. Starting from an empty board, stones are added in turn on the intersection of lines of a variable-size board, usually 19×19. Stones are captured if surrounded by stones of the opposite color. Winning the game means to acquire more territory than the opponent (Fig. 1). Territory is defined as free intersections surrounded by stones of your own color or the edges of the board. A mapping can be made between the territorial relationship of the stones on the board and their respective role in the game. Each move in the game can be explained away by its purpose. There are three phases in the game of Go which are not always temporally or visually distinct. The opening phase in the corners, called *josekis*, has been extensively studied and "dictionaries" of opening templates are available. The endgame, where the board can be decomposed into independent regions, has been practically solved with combinatorial game theory (Berlekamp & Wolfe 1994). The middle game, however, is still intractable and hasn't been reduced to any single approach. Go is a challenging platform for AI techniques (Bouzy & Cazenave 2001) and has replaced chess as the "drosophilia of AI." This paper claims that Go is also an excellent platform for multi-agent systems.

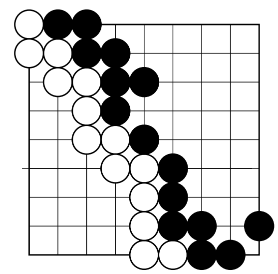

Figure 1: Territory: Black wins with a 32/30 free-intersection advantage over White.

This paper is organized as follows. We first introduce the practical contexts for multi-agent systems (MAS) in our research in which to frame common "myth conceptions" about roles, namely multiple roles, open environments, and mixed strategies. Three myths are then defined and possible approaches described.

Roles in Multi-Agent Systems

Some issues about roles arise in the following contexts in the development of multi-agent systems.

Multiple Roles

Role allocation can be viewed as a constraint satisfaction problem (CSP) similar to resource allocation problems

[1] A good introduction can be found at http://www.usgo.org.

where agents are variables $(x_1, ..., x_n)$ and roles are values that those variables can take. When a utility or preference is attached to an assignment, role allocation also entails a constraint optimization problem. There are typically N roles and N agents and the goal is to maximize the sum of utilities or Pareto efficiency in assigning agent x_i to role j subject to the constraint that $x_{i \in N} \neq x_{i' \in N-i}$. In dynamic environments with partial observability, this constraint must be relaxed and the optimization goal becomes elusive. The multi-agent system problem can be formally described as follows:

- $x_i = j$ ←assignment of an agent i to role j
- u_{ij} ←utility of role j to agent i, where $\sum_j u_{ij} = 1$
- w_j ←priority of role j in the global task, where $\sum_j w_j = 1$
- Maximize $\sum_{i,j} u_{ij} w_j$ or Pareto efficiency of a solution

When the global task consists of multiple coordination tasks and multiple teams coexist, an agent can take on multiple non-conflicting roles. The goal is then to maximize

$$\frac{\sum_t \sum_{i,j} u_{ij} w_j}{|N|} \tag{1}$$

where t is the number of teams. Over-constrained situations where a team cannot be formed because of an insufficient number of agents are not relevant here. Rather, an agent has to assume multiple roles in those situations. The capability to handle multiple roles is an important aspect of the coordination of intelligent agents to achieve the promised payoff of multi-agent systems. In order to extend role allocation algorithms to multiple roles, the myth that roles are "crisp", that is distinct and non-overlapping, must first be addressed.

Open Environments

Open environments are characterized by dynamic team formation, multiple and changing goals, non-deterministic state transitions, variable resources, and unreliable communication (Abramson & Mittu 2004; Dastani, Dignum, & Dignum 2003). In this context, roles become fluid since the possibility of finding a better match for the role required exists. The cost in switching roles, such as increased communication, must be balanced against possible benefits. In the game of Go, as in open environments, stones can be added at any time if free intersections remain and there is no limit to the number of stones that can be played. In order to extend multi-agent systems to open environments, the myth that roles are "static" must first be addressed.

Mixed Strategies

Roles provide the high-level granularity for a human-in-the-loop type of strategy to direct the task while coordination can occur at the lower level of primitive action selection. In this context, it is not clear that only strategic level roles are those candidates for human intervention. The dichotomy between tactics and strategy and where priorities lie is not always clear. In the game of Go, strategic moves, called *fuseki*s, occur in the opening and affect the general outcome of the game. However, an evaluation function of the moves in Go has been so far elusive so that there is no guarantees of winning the game after the initial strategic phase. This provides a dilemma on what should be brought up to the human attention. In order to get some insights into this problem, the myth that roles are "flat", that is one-dimensional, either tactical or strategic, must first be addressed.

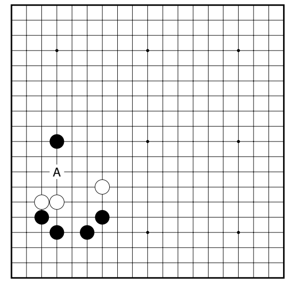

Figure 2: Myth 1 - In this situation, 'A' is the vital point for White for both attacking and defending. (Tisheng & Wen 1983)

Myth 1: Roles are "crisp"

Not only do we multitask out of necessity, but one highlight of intelligence is flexibility and adaptability in our interactions towards our peers and the environment. Premature role allocation leads to suboptimal situations and missed opportunities especially in an open environment. Maximum coordination efficiency is obtained when one action can fulfill multiple roles (Fig. 2). This suggests that coordination problems should be addressed at a lower level of granularity by decomposing roles into primitive actions (Abramson & Mittu 2004). Such an approach view roles at the spectrum of a continuum of primitive actions leading from large-scale coordination to local teamwork. Another approach, the token approach (Scerri *et al.* 2004), keeps non-conflicting roles intact and delays commitment to a specific role but does not directly address the problem of multiple role instantiations .

Myth 2: Roles are "static"

Sure, some roles are well-known apriori such as who drives the car and who is a passenger in the car. But new roles are constructed dynamically and created or destroyed all the time to respond to the need for coordination (Brooks 1995). That is the motivation given for the rise and fall of "organizations". This suggests that roles have to be synthetized on the fly to respond to dynamic situations affecting coordination (Fig. 3). In this context, roles are similar to goal-driven behaviors that can be incrementally learned or modified to perform a task. In addition, it was shown that the

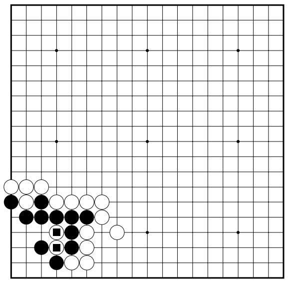

Figure 3: Myth 2 - Dead stones can suddenly acquire an important role. White can leverage from its (marked) dead stones to kill black.(Tisheng & Wen 1983)

fluidity of roles enables better coordination effectiveness. If all agents are running the same optimized algorithm, e.g. distributed constraint satisfaction, they should arrive at the same solution with a minimum of interference without the need for a consensus albeit with increased communication to share state information (Abramson, Chao, & Mittu 2005). How often to share state information when communication is unreliable while not contributing to congestion? Insect societies, such as the honeybees, demonstrate a plasticity in the division of labor of their members according to age and environmental conditions. Response thresholds (Bonabeau, Dorigo, & Theraulaz 1999) are a way of reacting efficiently to environment conditions in order to tune or switch roles in situations where the ownership of a role is not exclusive to an agent. How to use stigmergy to save on communication costs associated with role allocation in a dynamic environment?

Myth 3: Roles are “flat”

Roles structure interactions in relation to a goal. Roles can be represented as adjacency relations within a conceptual graph (Sowa 1984) relating agents in a task at the intentional level but their extensions into actions and/or messages might suggests different interpretations. CSP algorithms evaluate roles in relation to other roles but fail to take multiple goals into account leading to planning considerations. Getting the right local context in which to evaluate roles enabling a satisficing solution is necessary but difficult to do in a multi-agent system (Fig. 4). How much information should be relayed to the agents in order to obtain the correct context? How fine-grained should the interactions be? Similarly, in the game of Go, “whole-board” considerations (Yang & Straus 1996) arise in the interplay between strategic fusekis and tactical josekis in the opening (Fig. 5).

Tuple-space modeling (Gelernter 1985) is an attempt to establish an apriori context for coordination tasks involving communication while coordination graphs (Guestrin, Koller, & Parr 2002) provides an apriori context for tasks not necessarily requiring communication. Peer-to-peer discovery is a dynamic way of establishing the proper interaction context between agents.

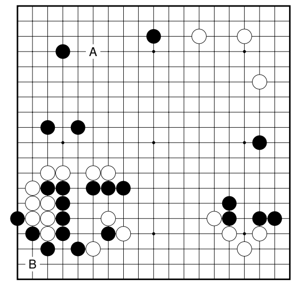

Figure 4: Myth 3 - Whole board considerations makes a “small” tactical move at B more “urgent” than the “big” move at A.(Tisheng & Wen 1983)

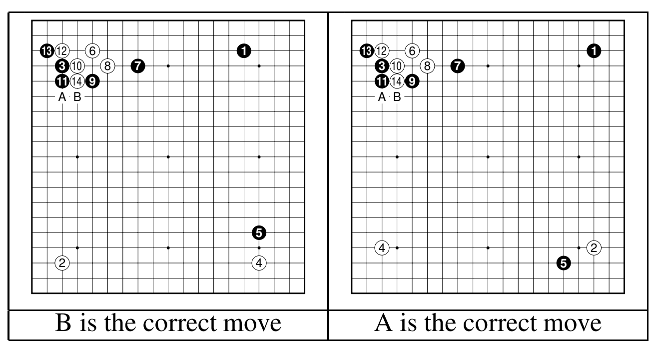

Figure 5: Interplay between fusekis (moves 1,2,3,4, and 5) and josekis (upper left corner)(Yang & Straus 1996)

Conclusion

Although roles are a powerful computational tool, a flexible granularity is needed to address the multiple-role problem in intelligent agents. Some of the usefulness of roles as apriori decomposition of tasks might be a limitation with new tasks requiring on-the-fly synthetization of roles. While behaviors are goal-oriented, roles are defined in relation to other roles for the only purpose of coordination and might only indirectly relate to the task. Coordination in multi-agent systems needs to interleave the creation of new roles with role allocation at the proper abstraction level.

The author wants to acknowledge useful discussions with Joe Collins and Mike Stein.

References

Abramson, M., and Mittu, R. 2004. Multi-agent systems in open environments. In *Proceedings of the Workshop on Large-Scale Coordination at the Third International Joint Conference on Autonomous Agents and Multiagent Systems*.

Abramson, M.; Chao, W.; and Mittu, R. 2005. Design and evaluation of distributed role allocation algorithms in open environments. In *International Conference on Artificial Intelligence*.

Berlekamp, E., and Wolfe, D. 1994. *Mathematical Go: Chilling Gets the Last Point*. A.K. Peters, Wellesley, Massachusetts.

Bonabeau, E.; Dorigo, M.; and Theraulaz, G. 1999. *Swarm Intelligence: from Natural to Artificial Systems*. Oxford University Press.

Bouzy, B., and Cazenave, T. 2001. Computer go: An ai-oriented survey. *Artificial Intelligence* 132(1):39–103.

Brooks, F. 1995. *The Mythical Man Month*. Addison-Wesley.

Dastani, M.; Dignum, V.; and Dignum, F. 2003. Role-assignment in open agent societies. In *Autonomous Agents and Multi-agent Systems*.

Gelernter, D. 1985. Generative communication in linda. *ACM Transations on Programming Languages and Systems* 7(1):80–112.

Guestrin, C.; Koller, D.; and Parr, R. 2002. Multiagent planning with factored mdps. In *Advances in Neural Information Processing Systems*, volume 14.

Scerri, P.; Farinelli, A.; Okamoto, S.; and Tambe, M. 2004. Token approach for role allocation in extreme teams: Analysis and experimental evaluation. In *Proceedings of 2nd IEEE International Workshop on Theory and Practice of Open Computational Systems*.

Sowa, J. F. 1984. *Conceptual Structures*. Addison-Wesley.

Tisheng, G., and Wen, L. 1983. *Strategic Fundamentals in Go*. Yutopian Enterprises.

Yang, Y., and Straus, P. 1996. *Whole Board Thinking in Joseki*. Fourth Line Press.