Mutual State Capability-Based Role Assignment Model (Extended Abstract)

Somchaya Liemhetcharat
The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
som@ri.cmu.edu

Manuela Veloso
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
veloso@cs.cmu.edu

ABSTRACT

We formally present the Mutual State Capability-Based Role Assignment (MuSCRA) model, as we introduce that an agent, acting in a team, has capabilities that depend not only on its own individual skills, but also on its teammates and their mutual state. The MuSCRA model includes a description of roles in terms of its association value with states and actions. Role assignment policies are evaluated with a utility accounting for the match between the new mutual state capabilities and the desired roles, weighted by a risk factor.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms

Keywords

Capability, Role Assignment, Heterogeneous Teams, Multi-Agent

1. INTRODUCTION

In heterogeneous teams, agents have different capabilities with regards to different roles. Common approaches to this problem use market-based techniques, where agents bid over the roles or tasks [1], and do not explicitly model the capabilities of agents. Other approaches model single-agent capabilities [3], or model uncertainty in capabilities [2], but do not incorporate the state of the agents.

We focus on role assignment in a heterogeneous team, where an agent's capability depends on its teammate and their mutual state, i.e., the agent's state and its teammate's state. The capabilities of an agent are represented by a mean and variance of the utility attained by performing the action, which captures information about both the innate abilities of the agent, as well as how effective a particular pairing of agents would be when the action is taken in their mutual states. For example, an agent may perform an action well with a particular teammate in some mutual states, but not in other mutual states, or with any other agent.

We present a formal framework (named MuSCRA) for representing these situations, and formally describe how to compute the value of a role assignment policy, using a notion of risk. There are several real scenarios in dynamic environments where MuSCRA can be applied. In robot soccer, robots may have different capabilities in kicking the ball accurately and passing the ball to one another, which depends on the role assignments and mutual states of robots. Similarly, in urban search and rescue, different robots have different capabilities, for example the ability to cross rough terrain, the speed of movement, and the ability to detect trapped humans.

Cite as: Mutual State Capability-Based Role Assignment Model (Extended Abstract), Somchaya Liemhetcharat and Manuela Veloso, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

2. MODELLING AGENT CAPABILITIES

A heterogeneous team of agents consists of agents with different capabilities, and the goal is to find an assignment of roles for the agents such that the best team configuration is achieved, in terms of the utility attained. The capability of an agent to successfully perform an action depends on the teammate and their mutual state, i.e., the agent's state as well as the teammate's.

2.1 The MuSCRA Model

Definition 1. *A Mutual State Capability-Based Role Assignment (MuSCRA) model is a tuple* $\{\mathcal{X}, \mathcal{A}, \mathbf{a}, \mathcal{R}, \mathrm{S}, \mathrm{E}, \mathrm{C}, \rho\}$

- $\mathcal{X}$ *is the set of states*
- $\mathcal{A}$ *is the set of actions*
- $\mathbf{a}$ *is the set of agents*
- $\mathcal{R}$ *is the set of roles*
- $\mathrm{S} : \mathcal{R} \times \mathcal{X} \rightarrow \mathbb{R}$ *is the association between roles and states*
- $\mathrm{E} : \mathcal{R} \times \mathcal{A} \rightarrow \mathbb{R}$ *is the emphasis of actions in roles*
- $\mathrm{C} : \mathbf{a} \times \mathcal{X} \times \mathcal{A} \times \mathbf{a} \times \mathcal{X} \rightarrow (\mu_{\mathrm{C}}, \sigma^2_{\mathrm{C}})$ *is the function of capabilities, where*
$$\mathrm{C}(a_1, x_1, A, a_2, x_2)$$
returns the mean and variance of the utility obtained when agent a_1 *in state* x_1 *performs action* A *while agent* a_2 *is in state* x_2
- $\rho \in (0, 1)$ *is the amount of risk to take in assigning roles*

States and Actions

The set of states $\mathcal{X}$ is the set of all possible states of the agents. $\mathcal{X}$ is not the joint state-space of the team — each $x \in \mathcal{X}$ represents a state that a single agent can be in. Similarly, $\mathcal{A}$ is the set of all possible single-agent actions, not the joint action-space.

Agents and Roles

The set of agents, $\mathbf{a}$, represent the team of cooperative agents whose roles are being assigned. The roles of the team, $\mathcal{R}$, represent associations with certain states of the world, as well as an emphasis in certain actions. Each role is assigned to a single agent, and so roles can be viewed as the smallest element of a team.

Role-State Association and Role-Action Emphasis

Roles are associated with states of the world, and is represented by the function $\mathrm{S} : \mathcal{R} \times \mathcal{X} \rightarrow \mathbb{R}$. S indicates how strongly associated a state and role are, where a higher value indicates a stronger association. Similarly, different roles emphasize different actions, and is represented by the function $\mathrm{E} : \mathcal{R} \times \mathcal{A} \rightarrow \mathbb{R}$. S and E obey the following:

$$\forall R \in \mathcal{R}, 0 \leq \mathrm{S}(R, x) \leq 1, 0 \leq \mathrm{E}(R, A) \leq 1 \tag{1}$$

$$\forall R \in \mathcal{R}, \sum_{x \in \mathcal{X}} \mathrm{S}(R, x) = 1, \sum_{A \in \mathcal{A}} \mathrm{E}(R, A) = 1 \tag{2}$$

Eqn. 1 states that all associations and emphases of roles are between 0 and 1. Eqn. 2 normalizes across roles to ensure that the weighting of every role is equal.

Mutual State Capabilities

The capability function $\mathrm{C}(a_1,\ x_1,\ A,\ a_2,\ x_2)$ returns the mean and variance of the utility obtained when agent a_1 performs action A involving a_2 when a_1 and a_2 are in their mutual state (x_1, x_2). C takes into account that an agent's ability to perform an action and achieve the desired outcome depends on the teammate, and their mutual state, i.e., the agent's state and the teammate's state, and that there is uncertainty in the utility obtained by performing the action.

Risk

The ρ term in the MuSCRA model represents how much risk to take while assigning roles to the team. The utility U_r of a role assignment is normally distributed, and given a certain value of $\rho \in (0, 1)$, u is a value such that $P(U_r \leq u) = \rho$.

2.2 Evaluating Policies in MuSCRA

In order to evaluate different assignments of roles in the MuSCRA model, we define that a role assignment policy $\pi : \mathcal{R} \rightarrow \mathbf{a}$ is an assignment of roles to agents such that every agent has at most 1 role, i.e., $\pi(R) = \pi(R') \Rightarrow R = R'$.

Utility of a Policy

Given a role assignment policy π, we determine the utility of the team thus assigned, taking into account the capabilities of each agent and its assigned role. We define the utility of a role assignment policy as: $U : \pi \rightarrow (\mu_\pi, \sigma_\pi^2)$, where μ_π and σ_π^2 represent the mean and variance of the policy's utility. We can compute U as shown below:

$$U(\pi) = \sum_{\substack{R, R' \in \mathcal{R} : R \neq R' \\ x, y \in \mathcal{X} \\ A \in \mathcal{A}}} \phi(\cdot) C(\pi(R), x, A, \pi(R'), y)$$

where ϕ is a weight function:

$$\phi(R, x, A, R', y) = \mathrm{E}(R, A)\mathrm{S}(R, x)\mathrm{S}(R', y)$$

Using the action emphasis function E and role-state association function S, ϕ determines how much weight to place on the utility of an action taken by a role. Thus, actions with more emphasis in the role will reflect a higher weight in ϕ. Similarly, highly associated states of the agent and its teammate will have higher weights during U's calculation.

Incorporating Risk into Utility

Given the risk parameter ρ, we define the value of a policy:

Definition 2. *The value of a policy is given by the function $V : \pi \rightarrow \mathbb{R}$, where:*

$$V(\pi) = \mu_\pi + \sqrt{\sigma_\pi^2}\Phi^{-1}(\rho)$$

where Φ^{-1} is the inverse of the standard normal cumulative distribution function, and μ_π and σ_π^2 are the mean and variance returned by $U(\pi)$ respectively.

3. CONCLUSION

We formally defined the Mutual State Capability-Based Role Assignment (MuSCRA) model, and described each of its components. Capabilities of agents in MuSCRA are defined not only as a pairing between an agent and action, but also incorporates the teammate, and their mutual state, i.e., the state of the agent and its teammate. This allows a generalization of capabilities to include the fact that the utility of an action in a team depends on the composition of the team, as well as the state of the world. In addition, capability is represented by a mean and variance, to signify the uncertainty in the world, as well as the reliability of the data collected from observations. We defined how to determine the utility of such a policy in terms of a mean and variance, and how to incorporate the risk factor to retrieve the value of a policy, which adjusts the mean-to-variance trade-off in the role assignment policy.

Acknowledgments

This work was partially supported by the Lockheed Martin, Inc. under subcontract 8100001629/1041062. The views and conclusions contained in this document are those of the authors only.

4. REFERENCES

[1] V. Frias-Martinez, E. Sklar, and S. Parsons. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proc. RoboCup Symp.*, pages 532–539, 2004.

[2] C. Guttmann. Making allocations collectively: Iterative group decision making under uncertainty. In *Proc. 6th German Conf. Multiagent System Technologies*, pages 73–85, 2008.

[3] L. He and T. R. Ioerger. A quantitative model of capabilities in multi-agent systems. In *Proc. Int. Conf. Artificial Intelligence*, pages 730–736, 2003.