Preface

This workshop focuses on models and algorithms for multiagent interaction without prior coordination. Interaction between agents is the defining attribute of multiagent systems, encompassing problems of planning in a decentralized setting, learning other agent models, composing teams with high task performance, and selected resource-bounded communication and coordination. There is significant variety in methodologies used to solve such problems, including symbolic reasoning about negotiation and argumentation, distributed optimization methods, machine learning methods such as multiagent reinforcement learning, and others. The majority of these well-studied methods depends on some form of prior coordination. Often, the coordination is at the level of problem definition. For example, learning algorithms may assume that all agents share a common learning method or prior beliefs, distributed optimization methods may assume specific structural constraints regarding the partition of state space or cost/rewards, and symbolic methods often make strong assumptions regarding norms and protocols. In realistic problems, these assumptions are easily violated — calling for new models and algorithms that specifically address the case of ad hoc interactions. Similar issues are also becoming increasingly more pertinent in human-machine interactions, where there is a need for intelligent adaptive behavior and assumptions regarding prior knowledge and communication are problematic.

Effective models and algorithms for multiagent interaction are most likely to be achieved as we bring together work from many different areas, including work on intelligent agents, machine learning, game theory, and operations research. For instance, game theorists have considered what happens to equilibria when common knowledge assumptions must be violated, agent designers are faced with mixed teams of humans and agents in open environments and developing variations on planning methods in response to this. The goal of this workshop is to bring together these diverse viewpoints in an attempt to consolidate the common ground and identify new lines of attack.

Authors submitted papers in the same paper format as was used in the AAAI-14 main track. We received a total of 10 submitted papers of which 8 papers were accepted, following a double-blind peer review process which resulted in two reviews for each paper. Each accepted paper is included in the workshop proceedings and the authors will be given time to present their work during the workshop. The workshop program will also include invited talks by three highly-distinguished researchers with background in models and algorithms for multiagent interaction: Subramanian Ramamoorthy (University of Edinburgh), Peter Stone (University of Texas at Austin), and Manuela Veloso (Carnegie Mellon University).

- Stefano Albrecht, Samuel Barrett, Jacob Crandall, and Somchaya Liemhetcharat