

Diagnosing a Team of Agents: Scaling-Up

Meir Kalech Gal A. Kaminka
The MAVERICK Group
Computer Science Department
Bar Ilan University, Israel

Abstract

Agents in a team must be in agreement. Once a disagreement occurs we should detect the disagreement and diagnose it. Unfortunately, current diagnosis techniques do not scale well with the number of agents, as they have high communication and computation complexity. We suggest three techniques to reduce the complexity.

1. Introduction

Agents in a team must be in agreement as to their goals, plans and at least some of their beliefs [1, 3]. Unfortunately, they may come to disagree due to uncertainty in sensing, communication failures, etc.

Once a disagreement occurs the agents should diagnose it and provide a solution. The diagnosis process identifies which agents are in disagreement and about what they disagree, so that they can negotiate and argue, to resolve the disagreements [1]. We refer to this kind of diagnosis as *social diagnosis*, since it focuses on finding causes for *inter-agent* failures, i.e., failures to maintain relationships between agents in a team. Social diagnosis stands in contrast to *intra-agent* diagnosis, which focuses on determining the causes for components within agents.

In previous work [2] we focused on problems of disagreement between agents. We have shown that one can reduce the communication by centralizing the diagnosis, so all the agents may send their information to a single predefined agent who compares between these beliefs. However, in teams where the number of agents is scaled-up, the computation may be expensive. Moreover, we showed that further reductions in communications, based on using inference of other agents beliefs, is exponential in run time.

In this work we focus on tackling the complexity of communication and inference, to enable diagnosis of large-scale teams. We suggest new methods of social diagnosis, that reduce both the communication and computation.

2. Scaling Diagnosis Methods

In our previous work, we presented Querying algorithm to make a diagnosis [2]. Querying proceeds in three stages

(Figure 1). First, it observes its peers and uses a behavior recognition process to identify their possibly-selected behavior paths, based on their observed actions. Then, based on the hypothesized behavior paths it further hypothesizes the beliefs held by the teammates by belief recognition process (which led them to select these behavior paths, by enabling sets of pre-conditions and termination conditions). Finally, it queries the diagnosed agents as needed to disambiguate between these belief hypotheses. Once it knows about the relevant beliefs of each agent, it compares these beliefs to detect contradictory beliefs which explain the disagreement in behavior selection.




Figure 1. Querying process for a single agent

We suggest three methods that tackle the runtime and communication complexities of querying. Each method tackles the complexity stemming for a particular factor in the complexity of querying: the number of behaviors, the number of beliefs, and the number of agents. **(i) behavior querying** eliminates the behavior recognition process by querying about the selected behavior path; **(ii) shared beliefs** limits the belief recognition process by inferring only the propositions of the beliefs, not their value; and **(iii) grouping** abstracts the diagnosed agents by grouping together agents along disagreement lines, and selecting representative agents for diagnosis.

1. Behavior Querying: Generally a behavior is associated with several beliefs through its preconditions and termination conditions. Therefore, we expect that the size of belief hypotheses will be greater than the size of behavior path hy-

potheses, since each behavior path hypothesis may generate several belief hypotheses as previously described.

We can eliminate the uncertainty in the behavior recognition process by disambiguating the observed agent's behavior path using communication, instead of inferring all its behavior path hypotheses. This goal is achieved by querying the observed agent about its behavior path. Once the diagnosing agent knows the behavior path of the monitored agent, it continues to build the belief hypotheses that are associated only with that behavior path. The advantage of this method is that by a single query about the behavior path of the observed agent, it eliminates all the queries about the belief hypotheses associated with other (incorrect) behavior path hypotheses.

We predict an improvement relative to the behavior path hypotheses process in term of communication, since we now expect to see only one message from each observed agent independently of the number of behavior path hypotheses. Also, we predict an improvement in terms of runtime since the behavior querying method eliminates the belief hypotheses computation of all the behavior path hypotheses except for the correct one. So instead of the linear complexity of behavior recognition (in the number of behaviors in the behavior hierarchy), the number of behaviors has no effect at all, and the resulting complexity is $O(1)$.

2. Shared Beliefs: The main factor that causes a high runtime of the querying algorithm is the use of belief recognition process. This process grows **exponentially** in runtime with the number of beliefs associated with hypothesized recognized behavior paths. Even if the number of behavior path hypotheses is one, belief recognition will typically have multiple beliefs associated with it, and thus result in an exponential number of belief hypotheses.

We present a light-weight belief recognition technique whose complexity grows **linearly** with the number of beliefs. The key to this technique is to infer only the propositions associated with a belief, without hypothesizing about its value. In other words, the key is to infer that an agent has beliefs about p , without inferring what these beliefs are (p or $\neg p$). The diagnosing agent uses this technique to infer, for each agent, what propositions it holds. Then, for each pair of agents it queries for the values of propositions that are shared by the agent, and may thus be in conflict. For instance, if p is a proposition shared by agent A and agent B , a possible diagnosis is that agent A believes p while agent B believes $\neg p$. Thus the diagnosing agent should send a query to agent A and B about the value of p in order to determine if there is a contradiction.

Using this method, we expect that the communication will increase in the number of agents relative to the querying algorithm, since in teams we expect that most of the beliefs will be shared beliefs, so most of them are suspected. But, we expect to reduce the runtime complexity signifi-

cantly, since instead of inferring all the exponential number of belief hypotheses, we use a process that is linear in the number of beliefs.

3. Grouping: Regardless of how knowledge of the beliefs of teammates is inferred, the diagnosing agent must compare between the beliefs of the teammates after inferring those beliefs. This comparison is polynomial in the number of agents and in the number of beliefs. However, in a large-scale team, runtime may be too high in practice.

The grouping method abstracts the observed agents, grouping together agents that are in a similar state. It then uses a single agent from each group as a representative for all agents in its group. To determine the diagnosis, it only compares the beliefs of these representative agents, thus significantly reducing the number of comparisons.

The process is based on the assumption that two or more agents that have both the same role in the team and the same behavior path will have the same beliefs, at least with respect to their selection of role and behavior path. Based on this assumption only representative agents of each role and behavior path must be diagnosed.

To determine the different role/behavior path combinations, the diagnosing agent first disambiguates the behavior path of each monitored agent using *behavior querying* process. It then divides the team to groups based in their roles and behavior paths. This essentially divides the team along disagreement lines. It continues to do the diagnosis process only against representative agents of each group, either by querying algorithm or by shared belief methods. Finally, it uses the results of the diagnosis for the remaining members of the groups.

We predict that this process will reduce both the number of messages as well as the runtime, since the diagnosis process involves a significantly lower number of agents (only the representative agents of the groups), and likely this number is much smaller than the number of agents in the team (see the next section for an analysis of the maximum number of groups possible given a set of roles and behaviors). However, communications will still in the number of agents, though slowly, since the diagnosing agent has to disambiguate the behavior path of the agents by behavior querying in order to divide the team to groups.

References

- [1] B. J. Grosz and S. Kraus. Collaborative plans for complex group actions. *Journal of Artificial Intelligence Research*, 86:269–358, 1996.
- [2] M. Kalech and G. A. Kaminka. On the design of social diagnosis algorithms for multi-agent teams. in *International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 370–375, 2003.
- [3] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.