

Measuring Team Effectiveness

Henry Hexmoor

Computer Science & Computer Engineering Department
Engineering Hall, Room 313
Fayetteville, AR 72701
hexmoor@uark.edu

U.S.A.

Gordon Beavers

Computer Science & Computer Engineering Department
Engineering Hall, Room 313
Fayetteville, AR 72701
gordonb@uark.edu

U.S.A.

ABSTRACT

Accurately assessing team performance is required by team directors in order to improve the likelihood of achieving team goals. Unfortunately, there are no generally accepted and broadly useful methods for assessing team performance. This paper takes preliminary steps towards developing general criteria for such an assessment by suggesting how team performance of agents can be measured. These measurements can be used to assess the need to improve team effectiveness. The measurements might be obtained either by evaluations internal to a team (subjective) or by external means (objective).

Keywords

Agents, Teams

1. INTRODUCTION

In a previous paper [1] we considered fundamental issues about the nature of teams of agents and the mental states of the agents required for the formation of a team. That paper put forward some criteria for the existence of a team as a guide for the development of systems of artificial social agents. The conclusions of [1] are summarized in section 2 of this paper which contends that a "team" is an aggregation of agents with a common mental state that includes having a common intention, awareness of being a member of the team and the team's abilities, and a cooperative attitude. Furthermore team members must have autonomy in order to form intentions.

By measuring the effectiveness of teams we can begin to address methods that improve it. Environments such as Soccer have explicit agents (the coach) that specialize in monitoring team effectiveness and suggesting ways to improve it. Unlike pre-arranged teams, such as Soccer teams, we are seeking to develop methods that agents can utilize on their own in the organization of a team. It is our intention to provide methods to facilitate teamwork that are general and determined online without outside intervention.

Sycara's work ([7] and [9]) is closely related, as is [8]. Sycara intended to provide a general level of cohesion for heterogeneous agents and a toolkit for assembling teams. She is considering large open agent groups on the Internet.

Much like public services of a town, Sycara's middle agents mediate transactions between agents.

Beyond the model of teamwork such as Cohen and Levesque's famous Joint Intentions theory, there has been some work on monitoring team effectiveness and actually tools for composing teams; Karma-Teamcore is an example [10]. Tambe's work describes a methodology he calls Team Oriented Program whereby a human designer interacts with the system to determine roles and assign roles and plans to individuals. Our developments are similar but the intent is for the agents to determine when and how they work with a team and there is no notion of an external system designer.

The ideas, notation, and formalism found in this paper are constrained by the logic used to model autonomous agents, namely, multi-modal temporal BDI logic in the style of CTL*. In the remainder of this paper, section 2 outlines our approach to defining teams. In section 3 we will describe a few subjective team effectiveness quantification techniques and then offer concluding remarks.

2. TEAM CONDITIONS

In this section we review criteria set forth in [1] for the existence of a team. In order for a team to exist the agents forming the team must have a common goal as a part of a common mental state. The common mental state requirement should be understood to include the following four components:

1. The group must have a joint intention, in the sense that there is at least one intention ψ such that each individual agent of the team intends ψ and believes that the other team members also intend ψ . All intentional agents have some degree of autonomy.
2. Each individual team member must believe itself to have some ability that could contribute the achievement of the shared goal.
3. Each individual team member must be aware of being a team member.
4. Each individual team member must maintain a cooperative stance toward the recognized members of the team with respect to the shared goal.

Teams are individuated by their joint intention and membership. In order to develop the concept of intention we must first develop the concepts of “ability” and “choice”.

A requirement for being a team member is that an agent must possess an ability (in the sense of being able to perform an action), which contributes to the achievement of the team goal. Team members can be either actual, or virtual agents. There is no sharp distinction between virtual and actual agents. In the case of virtual agents, these actions may be limited to actions such as planning, strategy determination, communication, and computation. Using Wooldridge’s development of *Can* [11] as a model, the following definition of ability makes explicit that it is the action (or complex of actions) α , which is expected to bring about a temporal state in which ψ is true. We also make it explicit that ability depends upon the current temporal state in which the agent finds itself. The difference is that the definition of ability that follows requires the agent to be committed to a specific course of action. Consider an agent’s intention to move from location A to location D, where the agent may do so by first moving from A to B and then from B to D, or by moving from A to C and then from C to D, with the ability to follow either course. The action required by the following definition of ability for the agent to have the ability to move from A to D can either be stated as a disjunctive action or at a level that omits mention of the intermediate locations. The idea that needs to be captured is that at the current temporal state that agent i is in, there is a path on which i performs α and the action α leads to a temporal state in which ψ is true.

Definition 2.1 Can. $(\text{Can } i \alpha \phi \psi) \equiv (\text{Bel } i (\text{Agt } \alpha i) \wedge (\text{Achvs } \alpha \phi \psi)) \wedge (\text{Agt } \alpha i) \wedge (\text{Achvs } \alpha \phi \psi)$

that is, agent i in state ϕ has the ability to achieve a state in which ψ is true if and only if the action, or complex of actions, α for which i believes i is the agent and i believes that α will achieve a state in which ψ is true and in fact i is the agent and α does achieve ψ . Here $(\text{Agt } \alpha i)$ signifies that the agent i carries out the action α , and $(\text{Achvs } \alpha \phi \psi)$ signifies that α is the action that leads from the state ϕ to a state in which ψ is true. In the case that there are multiple actions which might lead to the same temporal state, an agent may be permitted to perform some of those actions, but not others. This leads to a notion of ability that depends on autonomy, which in turn is affected by permissions.

Generalizing to the multi-agent case yields the following definition.

Definition 2.2 Joint-Can. $(\text{J-Can } g \alpha \phi \psi) \equiv (\text{J-Bel } g (\text{Agt } \alpha g) \wedge (\text{Achvs } \alpha \phi \psi)) \wedge (\text{Agt } \alpha g) \wedge (\text{Achvs } \alpha \phi \psi)$

that is, the group g has the ability to achieve a state in which ψ is true if and only if the action, or complex of actions, α that g collectively believes g can collectively perform and g collectively believes that g will achieve ψ , and in fact g does collectively perform and α does achieve ψ .

In J-Can, group belief is reducible to individual belief, that is, each individual team member believes that it has some ability that could contribute to the achievement of the shared goal and collectively the team must have some non-trivial probability of achieving its intention. An agent that has the ability to achieve some state in which ψ is true may also desire to have this be the case. An agent may also have permission to perform an act that will bring the state about. Permissions are derived from social considerations. The combination of ability, desire, permission and autonomy enables the adoption of an intention.

The members of a team with the ability to achieve a goal must have a joint intention to achieve the goal in order to be a team. Adopting an intention is an act of self-determination, but even stimulus-response agents can be considered to be self-determined. Adopting an intention implies that the agent has some degree of autonomy. “Self-determination” does not distinguish between innate and “creative” behavior, but autonomy does. The notion of autonomy that we are trying to get at here has something to do with self-determination and complexity. Ants in a swarm are assumed to be too simple to be able to form intentions although they do individually make reactive decisions that collectively lead to coherent group behavior as judged by an external observer. Individual ants are unable to persuade other ants to pursue a goal other than the one reactively adopted. Ants are incapable of forming intentions. Intentional agents are more complex. Intentionality implies that agents who are team members have some degree of autonomy. Determining criteria for the existence of a team requires the consideration of autonomy and concomitant issues such as the importance of a cooperative attitude, abilities, shared intentions and shared worldview.

We also distinguish between a team, which is formed to solve a specific problem, and a coalition, which is a work group formed to increase the mutual benefit of the individual group members. Members of a team share a joint persistent goal while members of a coalition share a joint utility and strive for mutual benefit. Teams are work groups that attempt to achieve a goal even though in doing so they may engage in activities that adversely affect their individual condition. The team persists so long as the achievement goal persists. Coalitions on the other hand are work groups that pay constant attention to their individual benefit. Stability, viewed as the longevity of the work group, is desirable in coalition formation, however, the sense of stability comes from the expected longevity of benefit and there is no notion of a determinate persistent

intention as there is when a team is seeking a goal. Other desirable properties of a coalition are efficiency, low cost of forming a coalition, and distributed computational load over the agents joining the coalition. This makes coalitions suitable particularly in economically oriented settings. In comparison, teams seem more suitable when agents must reason about goal achievement and come to an agreement about working together.

We take an agent's autonomy with respect to an action as a relative measure of its liberty to exercise individual preference with respect to that action. This individual preference is constrained by social considerations in that it includes consideration of other agents, how the other agents contribute to the agent's sense of freedom to choose and to the agent's performance and how it prefers to work with others. These issues may be combined to yield an assessment of an agent's *permission* to perform an action that will lead to another state. The consequence of an agent's autonomy consideration will determine whether the agent is self-directed, other directed (in the sense of delegating the task to another agent), shares responsibility with other agents (as in teaming), or partially self-directed [3]. Many parameters, both endogenous and exogenous may go into a cost/benefit calculation of these preferences [4]. The personal autonomy (with respect to a given goal) of an individual might vary from complete self-determination (internal, high individual autonomy) of every aspect of the agent's relationship towards the goal, to complete other-determination (external, low individual autonomy) of those aspects. The notions of group autonomy and cooperation can be developed in a similar manner. Autonomy is a complex and subtle concept currently under intensive study. Until the analysis of autonomy has been further developed, we shall treat it as an unexplained primitive. An agent i may prefer to share its autonomy.

Agents with group autonomy can join in group activity and enter a team if they further have a cooperative stance and come to have a joint intention. In addition to shared autonomy, we require agents to have a joint cooperative attitude. Agents must adopt the principle of social rationality in order to be cooperative [5]. Social rationality is used to develop a theory of joint responsibility. It is sufficient that team members adopt an obligation to cooperate that will be considered equivalent to having a cooperative attitude.

When an agent adopts an intention at a particular temporal state, the agent does so only after consideration of its desires, and its autonomy with respect to those desires. Desires are possible worlds that we assume the agent is able to partially order according to desirability, that is, we assume that there is a measure of individual benefit that partially orders the possible worlds. The result is that an action is chosen by the agent to cause a transition to the desired state. We use the notation $(\text{Choice } i \alpha)$ to indicate

that the application of the choice function (which has performed a cost/benefit analysis) by the agent i yields the action α . Likewise, $(\text{choice } g \alpha)$ indicates application of a similar choice function got group g . The state at which the choice function is implemented will affect the result, but the notation conceals this component. Note that it is immediate from the definition of *individual intention* that intentionality implies autonomy, which is in accordance with common sense.

Definition 2.3 Individual Intention

$$(\text{Int } i \alpha \phi \psi) = (\text{Autonomous } i \alpha \phi \psi) \wedge (\text{Choice } i \alpha)$$

Commitment to the intention (goal) is part of being a team member and this commitment will demonstrate itself in the persistence of the intention. Joint intention will be defined in the pattern of individual intention.

Definition 2.4 Joint Intention:

$$(\text{J-Int } g \alpha \phi \psi) \equiv (\text{J-Autonomous } g \alpha \phi \psi) \wedge (\text{Choice } g \alpha)$$

In addition to having a joint intention and joint ability the members of a team must be aware of their membership in the group and maintain a cooperative attitude towards the recognized members of the group with respect to the joint intention. Our notion of awareness captures an agent's cognizance of at least some of its teammates' actions. This cognizance may rely either on the agent's plan recognition abilities, communicated messages, or direct observations. In this paper, we will not expand on varieties of means for becoming aware nor do we give an account of ongoing awareness maintenance (or active monitoring) but rather refer the reader to [6]. Instead we capture this awareness with an agent's beliefs about other agents that are recognized members of the team.

As developed above, part of the requirement for being a team is that the agents forming the team must individually intend that the team achieve the goal. Teaming requires coordinated action of the team members and this coordination in turn requires that the individual agents be able to determine when to start or stop an activity, when to switch activities, and when to adjust their contribution. A potential team member must have the ability to accept or reject a goal, that is, revise its intentions. Part of this ability is the autonomy to make this decision, that is, without the autonomy to adopt an individual intention an agent cannot become a team member.

A basic requirement for the existence of a team is coordinated action, and it is impossible for an agent to coordinate when she does not perceive what or with whom she is to coordinate. Each individual team member must be aware of being a team member. The agents forming the team must individually be aware of the existence of the team and intend to coordinate their actions with other perceived team members. Since teams are individuated by

intention and membership, the agent must know what the joint intention is and that it is part of the team with that intention.

An agent in state ϕ , is aware of being a member of a social group g , which intends to achieve ψ if and only if the agent believes that it is a member of the team, that, together with the group, it has joint ability toward the desired state, and it has the intention to achieve that state, that is,

Definition 2.5 Aware

$$(Aware\ i\ g\ \phi\ \psi) \equiv (Bel\ i\ ((i\ \epsilon\ g) \wedge (J\text{-}Can\ g\ \phi\ \psi) \wedge (J\text{-}Int\ g\ \phi\ \psi)))$$

Note that awareness requires only a belief that the agent is a member of the group, and does not require that the agent know who the other members of the group are. The members of the team have to be jointly aware of the intentions of some other team member to contribute.

Definition 2.6 Joint Aware

A group g of agents in state ϕ is jointly aware with respect to state ψ if and only if each agent of the group g is individually aware of that state

$$(J\text{-}Aware\ g\ \phi\ \psi) \equiv \forall i \in g (Aware\ i\ g\ \phi\ \psi).$$

An analysis of the notion of cooperative attitude has not yet been performed, and thus we include it as a primitive component of teamness. We are now ready to compose the requirements we have outlined into one definition that captures the notion of a team.

Definition 2.7 Team:

A group of agents g in state ϕ is a team if and only if they have joint intention, group autonomy, joint benefit, joint awareness, and joint ability, i.e.,

$$(Team\ g\ \phi\ \psi) \equiv (J\text{-}Intention\ g\ \phi\ \psi) \wedge (G\text{-}Autonomy\ g\ \phi\ \psi) \wedge (J\text{-}Aware\ g\ \phi\ \psi) \wedge (J\text{-}Cooperative\text{-}Attitude\ g\ \phi\ \psi) \wedge (J\text{-}Can\ g\ \phi\ \psi).$$

3. MEASUREMENT METHODS

Effective teamwork requires constant monitoring of team conditions so that the member agents can adjust their activities with respect to one another and the intended goal. Effectiveness is a measure of the likelihood that the team will achieve its goal. This measure might be determined internally by the team members (subjective) or by an entity outside the team to which some degree of objectivity is assigned. Effectiveness is determined, in part, by the following properties for which suggestive measures are provided:

- A. Efficiency is a measure of the resources used by the team in its attempt to achieve the team goal where resources can be time, effort, etc.

- B. Influence is a measure of how agents affect the performance of other agents.
- C. Dependence is a measure of how much the ability of an agent to complete its task is affected by the performance of other agents.
- D. Redundancy is a measure of the duplication of effort by distinct agents.

These measures are not independent, but generally, an increase in the efficiency, and positive influence, or a decrease in dependence, redundancy, and negative influence will yield an increase in the effectiveness of a team. Effectiveness will depend on communication and resource management but these issues will not be addressed here.

Each team member may perceive its ability, or the abilities of other team members in relation to the overall team objective. If the agent perceives a significant shortfall, it may announce this impairment. If agents collectively perceive a problem, team conditions need to be examined. Consider a measure of efficiency. An agent’s assessment of its own efficiency takes a neutral value in this example. Consider each agent’s assessment of another agent’s efficiency as a percent deviation from the base level. In the following table we show three agents. Each cell in the table from row i to column j is agent i ’s assessment of agent j ’s efficiency. For instance, agent 2 believes agent 3 is 30% inefficient. There is no need for reciprocity, i.e, agents may not make the same assessment of one another’s efficiency. The first three entries in the last row are each agent’s averaged efficiency. For instance agent 2 is 90% effective. In the fourth column we show the sum of each agent’s perceptions about other agents. For example, agent 2 believes there is 40% inefficiency among agents 2 and 3. We can add the values in this column to see the overall perception about efficiency. In this example, there is perception of 20% inefficiency.

	1	2	3		
1	0.0	+0.1	+0.2	+0.3	0.88
2	-0.1	0.0	-0.3	-0.4	0.77
3	+0.1	-0.2	0.0	-0.1	0.91
	1.0	0.9	0.9	-0.2	0.84

Table 1 Efficiency

Once agents become aware of their average efficiency (first three values in the last row), they may use that in computing a subjective measure of efficiency. This is done by computing a weighted average of the other agents’ assessments of the agent’s efficiency. Weights are the agent’s overall score (first three values in the last row). The idea is to differentially weight assessments of agents toward one another. The vote of a more efficient agent will count more heavily than the vote of a less efficient agent.

Weighted averages are shown in the rightmost column. For instance, agent 3's subjective efficiency is at 0.91. By averaging the values in the last column, we arrive at an overall measure of efficiency. For the group of agents in our example, this is 84%.

The notions of independence and influence are related, but distinct. When one agent depends on another agent to accomplish a task before it begins its own task, the failure of the first agent interferes with the task of the second agent. There are also cases where influence has nothing to do with dependence. Therefore, we shall include separate measures for influence and independence. We address influence first. Influence can be either beneficial or harmful. If agents collectively perceive a problem with influence, team conditions need to be examined. We do not consider self-influence. Consider each agent's perception of the influence of another agent as a percent deviation from neutral. In the following table, we show three agents. Each cell in the table from row *i* to column *j* is agent *i*'s perception of the influence of agent *j*. For instance agent 2 believes agent 3 has a 20% positive influence on agent 2. There is no need for reciprocity, i.e., agents may not perceive the same amount of influence between them. The last row is the average influence that agent has produced. For instance agent 3 has produced 40% synergy (positive influence). Agent 2 has produced a 10% negative overall influence. In the last column we show, for each agent, an averaged interference it has experienced from others. In this example, all three agents have experienced 10% synergy. If we add the amount of interference produced and subtract it from amounts experienced, there is a net +0.2 influence. We can think of this group as being mutually supportive.

	1	2	3	
1	-	-0.1	+0.2	+0.1
2	-0.1	-	+0.2	+0.1
3	+0.3	-0.2	-	+0.1
	+0.2	-0.1	+0.4	+0.2

Table 2 Influence

Robustness is the ability to maintain team effectiveness in light of possible failures. Dependence of one agent on other agents detracts from robustness. Each team member may estimate its dependence on another member which is represented as a deviation from independence. The following table again shows three agents. Each agent is considered fully independent at level 0.0. Each cell in the table from row *i* to column *j* is agent *i*'s estimate of its dependence on agent *j*. For instance, agent 2 depends on agent 3 by 20%, which can be interpreted as agent 2 cannot finish 20% of its task without agent 3 completing its task. The first three entries in the last row are each agent's total

dependence on others. For instance, agent 1 depends upon agent 2 by 10%. These values reflect the team's dependence on the agent. In this example, agent 2 is the most dispensable whereas agent 1 is the most needed. In column 4, we show the sum of each agent's dependences on other agents. For example, agent 1 depends on agents 2 and agent 3 by 10% and 30% respectively. Since agent 2 depends on agent 3 by 20%, agent 1 indirectly depends on agent 3 by another 20%, which adds up to 60%. In this example, agent 2 is the most dependent on other members of the team ($0.6 + 0.2 + 0.1 = 1.0$), whereas agent 3 is least dependent ($0.1 + 0.1 = 0.2$). The values in the bottom row indicate the amount that other agents depend on the agent in that column.

	1	2	3	
1	-	0.1	0.3	0.6
2	0.7	-	0.2	1.0
3	0.1	0.0	-	0.2
	0.8	0.1	0.6	

Table 3 Dependence

Initial team organization activities are expensive. A plan must be discovered and agreed upon, and this activity can require extensive negotiation and communication. Cooperation requires monitoring the physical and intentional states of other team members. This monitoring might be effected through agent-to-agent communication. These activities often involve duplication of efforts. In addition, different agents may be better suited to particular activities than other agents and an agent's abilities may wax and wane over time. These considerations could cause reassignment of duties. In a sense/plan/act model, centralizing planning could greatly reduce the total amount of calculation, negotiation, and communication to determine a joint plan for the team.

Agents might duplicate tasks in part or in whole. For example, assembly inspection of a widget on an assembly line by all agents on the line is redundant and should be done only once. In the following table each cell from row *i* to column *j* contains agent *i*'s overlap with agent *j*. For instance, agent 2 overlaps with agent 3 by 20%. Overlap is symmetric. In the fourth column we add the overlap with other agents and see the agent's total overlap. For example, agent 3 has 50% overlap with the agents 1 and 2.

	1	2	3	
1	-	0.1	0.3	0.4
2	0.1	-	0.2	0.3
3	0.3	0.2	-	0.5

Table 4 Overlap and Repetitions

In proceedings of 20th IASTED International Multi-conference: Applied Informatics (AI 2002), p. 338-343, Innsbruck, Austria.

Consider the largest overlap in table 4, which is between agent 1 and 3. A service agent might be useful in eliminating this overlap. This agent can either perform the duplicated action or keep track of which of the two agents 1 and 3 performs the overlapped task and alert the other agent of the completion of this action so it can be skipped. The service agent may also pay special attention to agent 3 since it has the most overlap with the other two agents. Consider a case where an agent has an overall negative influence on other agents and the group depends least on the agent. Such an agent might be the one to be evicted from the group. In the examples in this section, agent 2 comes closest to this case since the group depends on the agent by only 10% (table 3) and it has a 10% negative overall impact on the group (table 2).

4. CONCLUSION

We have developed a formal account of teaming that relies on some joint understanding of a set of mental notions: intentions, autonomy, awareness, and benefit in [1]. In particular, nontrivial abilities of team members, inter-agent awareness, and benefit were introduced as additional criteria for a team. This paper has extended the treatment of teams to include the measurement of team effectiveness along the dimensions of efficiency, robustness and cohesion.

5. ACKNOWLEDGEMENTS

We thank an anonymous reviewer for helpful comments. This work was partially supported by AFOSR grant F49620-00-1-0302.

6. REFERENCES

- [1] G. Beavers and H. Hexmoor, Teams of Agents, in Proceedings of the *IEEE International Conference on Systems, Man, and Cybernetics* (SMC'2001), Tuscon, AZ, 2001.
- [2] H.J. Bürckert, K. Fischer, Holonic Transport Scheduling with TELETRUCK. *Applied Artificial Intelligence* 14(7), 697-725, 2000.
- [3] H. Hexmoor, A Cognitive Model of Situated Autonomy, In *Proceedings of PRICAI-2000 Workshop on Teams with Adjustable Autonomy*, Australia, 2000.
- [4] H. Hexmoor, Case Studies of Autonomy, In *Proceedings of FLAIRS 2000*, J. Etherage and B. Manaris (eds), p. 246- 249, Orlando, FL, 2000.

- [5] N. R. Jennings and J. R. Campos, Towards a Social Level Characterisation of Socially Responsible Agents, in: *IEEE Proceedings on Software Engineering*, 144 (1), 1997, 11-25.
- [6] G. Kaminka, D. Pyandeth, M. Tambe, Monitoring Deployed Agent Teams, *International Conference on Autonomous Agents* (Agents '01), 2001.
- [7] T. Lenox, R. Payne, S. Hahn, M. Lewis, K. Sycara, MokSAF, How should we support teamwork in human-agent teams? *CMU-RI-TR-99-31, The Robotics Institute, Carnegie Mellon University, 1999.*
- [8] E. Salas, C. Bowers, and E. Edens (Eds), *Improving Teamwork in Organizations* Lawrence Earlbaum Pubs, 2001.
- [9] K. Sycara, Brokering and Matchmaking for Coordination of Agent Societies: A Survey. In *Coordination of Internet Agents*, A. Omicini et al. (eds.), Springer., 2001.
- [10] M. Tambe, D. Pynadeth, and N. Chauvat, Building Dynamic Agent Organizations in Cyberspace, *IEEE Internet Computing*, March-April issue, 2000.
- [11] M. Wooldridge, *Reasoning about Rational Agents* The MIT Press, 2000.