

Received September 8, 2020, accepted September 23, 2020, date of publication October 5, 2020, date of current version October 16, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3028751

Metareasoning Structures, Problems, and Modes for Multiagent Systems: A Survey

SAMUEL T. LANGLOIS¹, OGHENETEKEVWE AKORODA¹, (Member, IEEE),
ESTEFANY CARRILLO², (Member, IEEE), JEFFREY W. HERRMANN³,
SHAPOUR AZARM³, HUAN XU², (Member, IEEE),
AND MICHAEL OTTE², (Member, IEEE)

¹Institute for Systems Research, University of Maryland, College Park, MD 20742, USA

²Department of Aerospace Engineering, University of Maryland, College Park, MD 20742, USA

³Department of Mechanical Engineering, University of Maryland, College Park, MD 20742, USA

Corresponding author: Samuel T. Langlois (slangloi@umd.edu)

This work was supported by the Air Force Research Laboratory (AFRL) under Grant FA8750-18-2-0114.

ABSTRACT Autonomous multiagent systems can be used in different domains such as agriculture, search and rescue, and fire protection because they can accomplish large missions more quickly and robustly by dividing them into separate tasks. Using multiple agents introduces additional complexity, which makes autonomous reasoning and decision making more challenging, however. Because agents such as ground robots, unmanned air vehicles, and autonomous underwater vehicles may have limited computational resources, they may need computationally efficient yet powerful reasoning algorithms (decision-making processes that perform deliberation and means-end reasoning). Metareasoning, which is reasoning about these reasoning algorithms, offers a way to tackle these challenges by monitoring and controlling reasoning algorithms to improve agent and system performance. Although metareasoning approaches for individual computational agents have been studied, no survey of metareasoning in multiagent systems (MAS) has yet appeared. This survey fills the existing gap by discussing the multiagent metareasoning approaches that have been studied in the literature. It identifies metareasoning structures, applications of metareasoning to reasoning problems, and the modes (techniques) used to control reasoning processes. This survey contributes to the study of MAS by providing a framework for discussing multiagent metareasoning, highlighting successful approaches, and indicating areas where future work may be fruitful.

INDEX TERMS Artificial intelligence, metacomputing, multiagent systems.

I. INTRODUCTION

Metareasoning describes reasoning about one's own decision-making process [15]. In dynamic, uncertain environments, metareasoning—a type of self-adaptation—seeks to improve an autonomous agent's performance by monitoring and controlling the agent's reasoning and decision-making processes. For example, consider a mobile ground robot that has a portfolio of planning algorithms that can determine a sequence of tasks for the robot; some algorithms run more quickly, but others generate better sequences that require less time to complete. When it needs to plan a sequence, the robot may run an algorithm selection procedure that selects the most appropriate algorithm based on relevant fac-

tors such as the number of tasks and the current computational workload. Selecting an algorithm to create the sequence is a metareasoning decision.

In an autonomous agent, metareasoning occurs in a meta-level that monitors and controls the reasoning algorithms in the agent's object level, which contains the reasoning algorithms that understand the environment and determine which, when, and how ground-level actions should be performed to achieve its goals [15]. In a multiagent system (MAS), the agents' object level may include additional reasoning algorithms such as coordination and teaming. Ground-level actions are actions that the agent takes that influence its environment and affect its state in the world. They include actions such as moving and sensing. In the example mentioned earlier, the meta-level action is selecting the planning algorithm, the object level includes the planning algorithm, and the

The associate editor coordinating the review of this manuscript and approving it for publication was M. Venkateshkumar¹.

ground level includes actions such as moving through the environment. Anderson and Oates [2], Cox [14], Cox and Raja [15], and Russell and Wefald [31] presented fundamental concepts in metareasoning. This paper describes how metareasoning has been applied in MAS.

MAS include swarms of unmanned air vehicles (UAVs, also known as drones), multi-robot systems, computer networks, intrusion detection systems, smart grids, and other applications [16], [33], [34]. For the purposes of this study, a MAS has multiple agents that may cooperate with each other by sharing information and coordinating their behavior to accomplish a mission or a set of tasks. Coordinating the agents' behaviors to achieve the system goal requires achieving consensus and synchronizing behaviors through appropriate task allocation [16], [21]. Despite previous work to develop collaboration algorithms (e.g., [12], [18], [19], [22], [32]), coordinating the agents' behaviors to achieve the system goal remains a challenge, for instance when communication between agents is unreliable [23].

In a multiagent system, agents must reason not only about their own actions but also about the actions of other agents, which make the environment more dynamic. Although the other agents can provide information, communication can be costly, and the other agents may have different objectives. This review considers communication as an object-level process because some reasoning algorithms (collaborative task allocation algorithms, for example) must communicate with other agents as part of their decision-making process.

Metareasoning can help the agents in a MAS adapt their reasoning algorithms and decision-making processes as the problem space or environment changes. Compared with metareasoning by a single agent (operating on its own), multiagent metareasoning may have different objectives (utility functions), problems, and structures. First, in a MAS, metareasoning may aim to maximize the utility of the MAS instead of the individual agent's utility. Second, in a MAS, an agent's object-level may include different types of reasoning algorithms such as communication, coordination, and team formation (which are not relevant when there is only one agent). Third, the agents in a MAS may use different structures to perform metareasoning (instead of each agent just metareasoning on its own), as described in Section II.

Although researchers have begun to study the potential for metareasoning to improve the performance of multiagent systems, we are not aware of any systematic survey of this topic. Existing surveys have focused either on multiagent systems or metareasoning but not both. For example, Anderson and Oates [2] reviewed research in metareasoning, but none of the work reviewed considered metareasoning for multiagent systems. Likewise, Conitzer and Sandholm [13] analyzed the computational complexity of some single-agent metareasoning problems.

This paper aims to address that gap by describing key aspects of previous research on metareasoning in multiagent systems, where there are challenges and opportunities for metareasoning that do not exist when a single agent is

operating independently. Researchers and system developers can benefit from a survey that not only describes basic definitions and previous work in multiple domains but also analyzes the progress of this stream of research and presents the next set of research challenges [25]. Such a survey will provide insights into how to use multiagent metareasoning effectively and help orient those beginning to work in this area.

This paper describes previous work on multiagent metareasoning from three perspectives: (1) the *metareasoning structure*, the relationship (if any) between the agents at the meta-level; (2) the *metareasoning problem*, the particular aspect of the agent's reasoning that the meta-level is controlling; and (3) the *metareasoning mode*, the means by which the meta-level modifies object-level reasoning. This paper focuses on these aspects because they describe the key relationships between metareasoning and the other aspects of the agents in the MAS.

This paper is not a comprehensive description of metareasoning or review of all research on metareasoning, which is beyond the scope of this work, and it does not consider system design problems that are solved off-line to determine aspects of the MAS (e.g., its structure or the control policy) before it begins operating. Internal details of metareasoning algorithms are also beyond the scope of this paper.

The remainder of this paper is organized as follows: Section II describes multiagent metareasoning structures. Section III describes multiagent metareasoning problems. Section IV describes multiagent metareasoning modes. Section V discusses our conclusions from this review. Section VI summarizes the paper and presents potential directions for future research.

II. METAREASONING STRUCTURES

There are multiple ways to implement metareasoning in a MAS. For example, in some approaches, each agent has its own meta-level that controls its own object-level reasoning. Other approaches use a centralized leader that controls the agents' reasoning. This section describes the different multiagent metareasoning structures that previous work has considered; the cited papers describe specific examples. (Note, because each metareasoning approach adopts a structure to govern reasoning about a problem through a mode, each study is discussed multiple times across Sections II, III, and IV.)

A. INDEPENDENT METAREASONING

In this type of structure, each agent in the MAS has its own meta-level, which performs metareasoning independently of the other agents, although the agents' object levels may communicate, as shown in Fig. 1. In the studies that we reviewed, this structure was very common in both non-cooperative MAS [28] and cooperative MAS [1], [5], [6], [9], [17], [20], [24], [26], [29], [30], [35], [36]. Because the meta-levels are independent, no additional communication or coordination between the agents is required, so implementing this structure is easier than implementing a decentralized structure. This structure adds the overhead of the meta-level to each

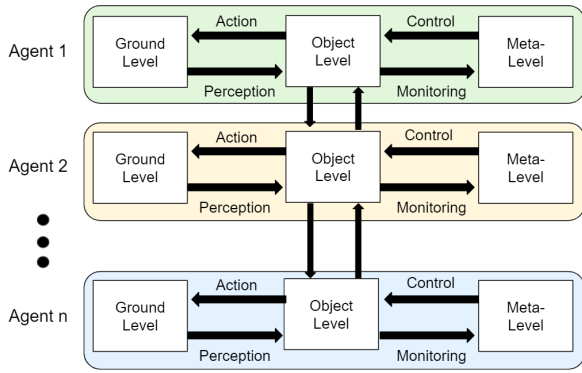


FIGURE 1. Multiagent independent metareasoning structure with three agents that communicate at the object level.

agent, which may affect the computational resources available at the object-level. In a cooperative system, moreover, the meta-level should consider how the agent’s reasoning affects not only agent-level performance but also system-level performance.

B. COUPLED METAREASONING

In this type of structure, each agent in the MAS has its own meta-level, which performs metareasoning independently of the other agents except for the following: when one meta-level has decided to halt its agent’s object-level reasoning, it communicates this decision to the other agents, and those meta-levels halt their object-level reasoning as well [10], [11]. As shown in Fig. 2, the meta-levels are coupled because one agent’s metareasoning decision depends upon another agent’s metareasoning decision (thus, the interaction goes one way; it is not bidirectional). In this structure, unlike the decentralized metareasoning structure, the meta-levels do not cooperate to make a coordinated decision; the meta-levels work independently, but they stop simultaneously. In a MAS in which the agents cannot act until they are finished reasoning, this coupling enables the agents to coordinate their behaviors. The communication cost in this structure is lower than the communication cost in the decentralized structure, but the metareasoning decision may be a poor one for agents that needed more time to find a better solution.

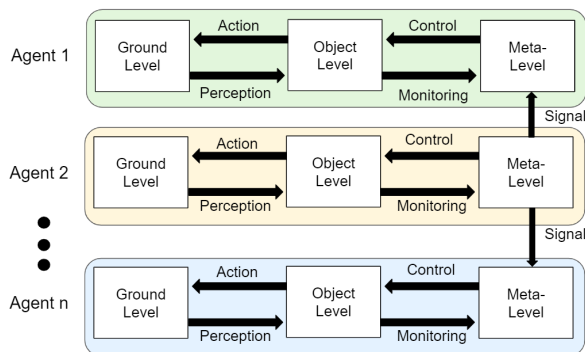


FIGURE 2. Instance of multiagent coupled metareasoning structure in which one agent sends a halt signal to the other agents.

C. DECENTRALIZED METAREASONING

In this type of structure, the agents in the MAS cooperate to determine how they are going to reason, as shown in Fig. 3. For example, Artikis [4] described a multiagent metareasoning approach in which the agents can propose and vote to enact changes to the rules that govern the MAS, even though the agents are competing against each other for external resources. This structure requires more communication and coordination, which may increase the overhead associated with metareasoning, but, by cooperating, the agents may be able to achieve better reasoning.

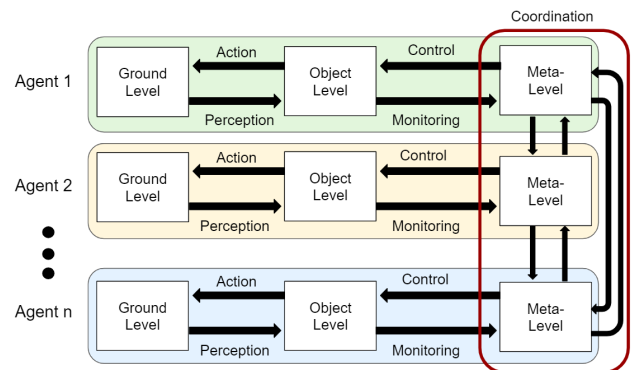


FIGURE 3. Multiagent decentralized metareasoning structure of three agents with meta-levels that communicate and coordinate.

D. MULTIPLE METAREASONING AGENTS

In this type of structure, the MAS includes, as shown in Fig. 4, additional specialized metareasoning agents that engage in metareasoning, and the reasoning of the other agents is influenced by the metareasoning agents. For example, a MAS in which some agents are trying to predict future events may have additional agents that determine which agents should cooperate [7], [8]. This structure adds resources (more agents) to perform metareasoning so that the other agents have no additional overhead. Keeping the metareasoning agents informed and communicating their metareasoning decisions will require more communication.




FIGURE 4. Multiple metareasoning agent structure of two agents for each meta-level agent.

E. CENTRALIZED METAREASONING

In this type of structure, a designated “leader” agent does the metareasoning and tells the other agents how to reason, as shown in Fig. 5. The leader’s objective is to maximize the performance of the entire system. For example, in a MAS in which agents are self-interested and will cooperate with some (but not all) of the other agents, the meta-level agent described by Pěchouček *et al.* [27] monitors the reasoning of the agents and shares that information with the other agents, which helps them make better decisions about forming coalitions (teams) of agents to accomplish tasks. Like the structure with multiple metareasoning agents, this structure adds additional resources and incurs the cost of additional communication. A single metareasoning agent that has information from the entire MAS should be able to make high quality metareasoning decisions.

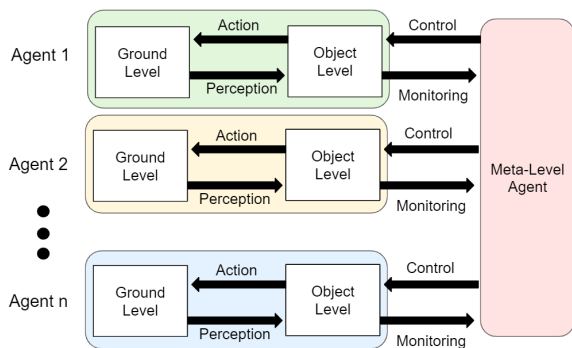


FIGURE 5. Centralized multiagent structure consisting of three agents with a central agent that monitors and controls all object level reasoning.

III. METAREASONING PROBLEMS

This section describes the metareasoning problems that previous work on multiagent metareasoning has considered. An agent’s object-level reasoning process may have many components, so there may be many options for which object-level reasoning activity the meta-level is controlling. There is no standard taxonomy of object-level reasoning for MAS, so the reasoning activities listed here are based on the papers in our survey.

As discussed in the following subsections, the papers that we reviewed discussed the following multiagent metareasoning problems: (1) multiagent coordination; (2) planning and scheduling; (3) communication; (4) resource allocation; (5) belief updating; (6) learning; (7) forecasting; (8) teaming; and (9) task delegation. In general, metareasoning about the reasoning processes that are solving these problems seeks to improve system performance while either reducing computational cost or satisfying constraints on computational resources.

A. MULTIAGENT COORDINATION

Research on multiagent systems has expanded the number of coordination methods available in multiagent settings.

Coordination methods present a trade-off between characteristics such as number of messages that should be sent and collaboration metrics, and their performance varies across different environments. Thus, it may be possible to improve performance by switching between coordination methods as the environment changes.

This problem can occur in a decentralized MAS in which the agents have little communication during runtime with their centralized controller and use a coordination policy to determine which tasks to perform. For a ship protection scenario, Herrmann [17] presented an approach in which each agent uses neural networks to estimate the performance of the candidate collaboration algorithms and selects the algorithm with the best estimated performance at that decision point. The best metareasoning policy studied reduced computational effort without degrading system performance.

When communication availability varies over time, the agents must determine which algorithm to use when communication availability changes. Carrillo *et al.* [9] determined a switching policy offline by evaluating the performance of collaboration algorithms at different levels of communication quality and showed that the metareasoning policy yielded better performance than using a single, fixed algorithm for decentralized task allocation in search, search and rescue, fire monitoring, and ship protection scenarios.

Negotiation is a type of coordination algorithm, and Raja and Lesser [29] described a MAS in which each agent’s meta-level must determine which negotiation algorithm to use. (This work also considered scheduling.)

B. PLANNING AND SCHEDULING

When an agent has many tasks to perform, its object level uses planning and scheduling algorithms to determine when it will do which tasks. Because planning and scheduling algorithms can be computationally expensive, controlling these algorithms is an important meta-level activity. For example, Raja and Lesser [29] described a meta-level that determines when to call the scheduling algorithm and selects among two scheduling algorithms that have different performance profiles so that urgent, high-value tasks can be scheduled quickly. Because computing the optimal policy using the real system state was intractable, they compared hand-generated heuristic policies and a policy derived from a Markov decision process that was based on a high-level (abstract) state. Using this abstract state reduced the overhead associated with the meta-level.

In the MAS considered by Rubinstein *et al.* [30], an agent has multiple scheduling algorithms (one that modifies its own schedule and another that creates hypothetical schedules in response to queries from other agents), and the meta-level controls these algorithms (and the resources that they use) by modifying their parameter values.

When the agents are working together to solve the planning problem, they must coordinate their reasoning to reach a good solution in a reasonable amount of time, so each agent’s metareasoning must consider not only its own computation

but also the progress of the other agents. Metareasoning must balance reducing the cost of computation and working longer to get a better solution. For a set of agents that are using any-time algorithms, Carlin and Zilberstein [10], [11] formulated this problem as a decentralized MDP so that the agents can decide when to stop computing based on their progress. They found that the metareasoning problem is computationally complex, but metareasoning policies can be determined by solving the decentralized MDP (if the number of agents is small) or using greedy algorithms for larger MAS.

Parker [26] studied metareasoning policies that update the values of control parameters that affect the algorithms that an agent uses to determine which tasks to perform and when to stop performing a task so that another agent can perform it. In particular, these parameters govern an agent's "impatience," which affects when it takes over a task that another agent is performing, and its "acquiescence," which affects when it will give up a task to another agent.

C. COMMUNICATION

An agent's object level will often communicate with other agents to share information that the other agents may use in their reasoning, but communication consumes resources. The meta-level may control communication to reduce computational cost. For example, in the MAS described by Xuan *et al.* [35], each agent decides, using a meta-level heuristic policy, whether to share its local information with the other agents, who cannot observe it directly; this communication has a cost but may help the other agents generate better solutions that lead to better system performance. Another option is to modify the communication policy during mission execution in response to the agents' needs [5].

D. RESOURCE ALLOCATION

Agents depend upon resources in order to reason and to perform actions. In some MAS, resources are allocated to agents by a resource allocation function. Changing resource allocation can affect the performance of the agents. (Here, resource allocation refers to providing agents with external resources, not the allocation or use of the agent's own resources.) Artikis [4] considered a resource allocation protocol that used rules to determine which agents had priority to get requested resources. The agents may propose changes to these rules and then vote to determine which rules were enacted.

E. BELIEF UPDATING

Each agent maintains beliefs about the environment and other agents, and it uses reasoning algorithms to update these beliefs when it obtains new, possibly inaccurate, information. Inaccurate information can be sent from agent to agent either voluntarily or involuntarily. For example, noise can distort a message involuntarily, but a competing agent may maliciously send false information. Belief updating is part of an agent's object-level reasoning, so it can be controlled by the meta-level if the agent realizes that some agents are

sending inaccurate information. Pinyol and Sabater-Mir [28] studied a metareasoning approach that adapted the way that an agent updates its beliefs in order to avoid two extremes (never trusting any agent, and always trusting every agent) that lead to poor system performance.

F. LEARNING

An agent learns by using the results of its actions to update the rules that it uses when reasoning. Metareasoning can monitor and control the learning algorithm so that the agent learns well with reasonable computational cost. For example, in Zhang and Lesser [36], the agents are communicating and coordinating with other agents during learning; correctly limiting the set of agents with which an agent coordinates reduces computational costs but has little impact on learning gains; system performance improved when agents communicated with smaller higher-quality groups of agents. In Noda and Ohta [24], each agent's meta-level modifies the agent's learning algorithm so that it is more likely to explore when it is choosing poorly.

G. FORECASTING

Agents may use forecasting to predict future events in the environment or actions by other agents. Metareasoning controls the forecasting algorithm in order to improve the predictions that are made. For example, grouping agents into clusters allows agents in the same cluster to share information and make better predictions about future events [10], [11]. For a MAS in which agents predict how other agents will behave, Borghetti and Gini [6] studied an algorithm selection approach that determined which prediction method to use by monitoring and then estimating each method's utility in the current context, which is an abstract, high-level state. Using this context reduces the overhead associated with estimating performance, which is done at the meta-level.

H. TEAMING

The agents in a MAS may decide to work together in small groups or teams to accomplish their tasks. Teaming is the reasoning process that determines which agents should work together. In this context, a team is a temporary arrangement that performs one or more tasks. For example, Pěchouček *et al.* [27] studied a centralized metareasoning approach that helps agents make better decisions about forming coalitions (teams) of agents to accomplish tasks when the agents lacked information about each other. In this approach, a meta-agent maintained a belief model of all agents in the system, continuously updated it based on events that occurred, and shared its conclusions with the agents; combining both inductive and deductive reasoning yielded more valuable information.

I. TASK DELEGATION

When an agent can assign tasks to another agent, task delegation is the reasoning process that determines which tasks to assign to which agents. (Here, task delegation refers to

one agent commanding another agent to do something, not collaborative task allocation, in which the agents are cooperating as peers.) Task delegation may be constrained by the nature of the relationships between the agents, however, so metareasoning can influence task delegation by redefining an agent's relationships with other agents [1], [20].

IV. METAREASONING MODES

This section arranges previous work on multiagent metareasoning by the mode that the meta-level uses to control object-level reasoning. (The mode does not describe how the meta-level decides. The mode describes what is specified or modified as the result of metareasoning; that is, the mode is the output, not the process.) This ranges from stopping a computational algorithm to changing the rules that govern agent interactions. The modes discussed herein occur in MAS and are determined as the agents perform their mission.

A. STOPPING AN ALGORITHM

In this mode, the meta-level determines when an object-level computation or algorithm should stop so that the object-level can begin other computations or act on the object-level decision. This mode is especially appropriate when the object-level algorithm is an anytime algorithm; the meta-level monitors the current state of the computation and tells the object-level when to stop the computation so that it can either start another computation or begin the selected action [10], [11].

B. MODIFYING A PARAMETER VALUE

The reasoning algorithms that agents use are governed by different parameters, so one opening for metareasoning is to modify the parameter value in response to changing conditions in order to improve performance.

For example, Pinyol and Sabater-Mir [28] considered a system in which a buyer agent receives information from multiple informant agents, who provide true and false information about multiple seller agents. The buyer agent modifies the function that it uses to update its beliefs about the sellers in response to the validity of the information that it has already received. A fixed updating function would lead to poor performance when conditions are changing or when an agent has incorrect prior beliefs about the other agents. This work used a simple update rule that required minimal computational resources.

Modifying a parameter can also affect the computational resources that reasoning algorithms consume. For example, Rubinstein *et al.* [30] used a metareasoning rule to modify parameters that affected the search algorithms used by different reasoning algorithms that were competing for resources; modifying the parameter values affected the computation resources used and the quality of the solution obtained by these algorithms (such as modifying its own scheduling and creating hypothetical schedules in response to queries from other agents).

Noda and Ohta [24] studied a MAS in which the agents are learning policies for choosing resources based on their own experience and the experience of other agents. Each agent's meta-level modified the values of parameters in the learning algorithm to encourage exploration or exploitation based on the agent's current performance.

In the multi-robot architecture proposed by Parker [26], each agent has a meta-level that monitors task performance and modifies the parameters of the algorithms that the agent's object-level uses to determine which tasks to perform and when to stop performing a task so that another agent can perform it. Parker studied different metareasoning policies that collected data about task performance and used that data to update the parameter values in different ways.

C. MODIFYING REASONING RULES

The agents in a MAS must follow certain rules that govern how they interact. Typically, these rules are determined a priori by the system designers. It is possible, however, to give the agents the ability to change these rules. For example, Artikis [4] considered a MAS in which the agents can propose and vote on changes to a resource-sharing protocol (the object-level reasoning). Although certain types of changes can be considered without harming the system, any change that is too far from the "desired" protocol cannot be approved (but the characteristics of the "desired" protocol change as aspects of the environment change). This constraint prevents the agents from enacting self-destructive protocols. Agents want to change the rules (in the resource-sharing protocol) to benefit themselves and evaluate possible changes from that perspective. When voting, each agent votes based on whether the change benefits itself, so that changes that benefit more agents are implemented.

D. SELECTING A REASONING ALGORITHM

Algorithm selection is a well-known metareasoning problem, and it occurs in a MAS when the agents must select algorithms that affect their reasoning about other agents or how they collaborate with other agents. Algorithm selection can be done by using a rule that determines which algorithm to use based on the current state or by estimating each algorithm's performance in the current state and selecting the best one. For example, in the MAS described by Raja and Lesser [29], an agent's meta-level determines which scheduling algorithm to use and which coordination (negotiation) algorithm to use.

In settings where communication availability varies over time, agents may want to switch collaboration algorithms (which depend upon communication) to use the best one. To keep online metareasoning effort low, Carrillo *et al.* [9] determined a switching policy offline by evaluating the performance of collaboration algorithms at high and low levels of communication and implemented a metareasoning policy that selected the collaboration algorithm based on the current communication quality.

In some cases, a simple collaboration algorithm that requires less time can generate high-quality solutions.

Herrmann [17] generated regression functions to estimate collaboration algorithm performance as a function of the current state. Although training the regression functions required much data, using them online requires little effort and allows one to add or remove algorithms without affecting the other regression functions. The results showed that a combination of one fast, simple algorithm and one more sophisticated, but time-consuming algorithm performed well.

For selecting the best method to predict the behavior of other agents, Borghetti and Gini [6] proposed a metareasoning approach that monitors the performance of different prediction methods and updates estimates of their performance. When the agent needs to make a prediction, it uses these estimates (which depend upon the current state) to select a prediction method.

E. AUTHORIZING COMMUNICATION

In a multiagent system, communication is an important reasoning activity when agents do not have complete information about the system state, but excessive communication is costly and brings no additional benefit. Determining whether to communicate is a metareasoning problem, and the general principle is to communicate if the expected gain is greater than its cost. Because finding an optimal communication policy is an intractable problem, Xuan *et al.* [35] considered different heuristics as meta-level control policies that determine when an agent should communicate. A hybrid policy that considered both the expected gain and the communication cost performed better than simpler heuristics that did not explicitly consider the communication cost. Becker *et al.* [5] studied an approach in which every agent follows the same policy (determined offline) and calculates the expected value of communication based on its beliefs about the other agents; the agent communicates if this value is positive. Having a computationally feasible procedure to determine the expected value of communication enabled the agents to make this decision based on the current state.

F. SHARING INFORMATION

In some MAS, an agent's reasoning depends upon its beliefs about other agents and how they reason. A meta-level can control agent reasoning by providing more or new information about other agents. For example, the meta-level agent described by Pěchouček *et al.* [27] controls the reasoning of the agents by sharing information that it has deduced about the other agents' reasoning.

G. DESIGNING COORDINATION

The ability of agents to coordinate generally depends upon their ability to communicate, and often agents try to coordinate with as many agents as possible. Because coordination requires communication, however, coordinating with every other agent may require excessive communication cost, and letting each agent determine its "coordination set" (the agents with which it will coordinate) may reduce communication cost. Zhang and Lesser [36] considered this problem

and developed a metareasoning approach in which each agent considers all possible coordination sets and selects the smallest one that still yields acceptable estimated performance. They found that tolerating just enough loss of performance can dramatically reduce computation cost.

Brueckner [7] used a different approach to determine which agents should work together (this work was also discussed by Brueckner and Parunak [8]). In order to decouple the metareasoning and reasoning functions, a set of metareasoning agents cluster the other reasoning agents into groups, and the agents in a group cooperate, which increases their effectiveness.

H. REDEFINING RELATIONSHIPS

In some MAS, the relationships between agents are not limited to peer-to-peer interactions, and these relationships affect how the agents reason. For instance, a superior agent may ask a subordinate agent to provide information or perform calculations for it. The corresponding metareasoning mode is to redefine (modify) these relationships in order to modify the agents' reasoning, which can affect system performance. (This does not change the existence of the relationship, as in designing coordination; it changes the nature of the relationship.) Kota *et al.* [20] considered a MAS in which an agent's meta-level determines when to redefine its relationships with other agents; an agent that is performing poorly will prefer to change its relationships despite the cost of reorganization. Ahmadi and Allan [1] studied a similar system and considered how limiting the number of relationships to redefine reduced the cost of reorganization.

V. DISCUSSION AND OPEN PROBLEMS

As shown in this survey, previous work on multiagent metareasoning has studied a wide variety of metareasoning structures, applied metareasoning to different problems, and used different metareasoning modes. This reflects the variety of MAS that have been developed for many different settings. Table 1 lists the papers that have been reviewed and summarizes the structures, problems, and modes that each one considered.

In principle, metareasoning should monitor and control an agent's object-level reasoning to optimize overall system performance, including the cost of computation and communication. Previous research has not yet developed a general formulation of this problem, which would require knowledge about the cost and performance of every aspect of an agent's reasoning process and how these combine to influence system performance. In the face of this complexity, researchers have reasonably adopted various and widely different approaches that address different parts of this problem. Although most studies have adopted (or are consistent with) the general framework of meta-level monitoring and control of object-level reasoning, researchers have investigated practical, system-specific approaches. (There is also the possibility of moving to a meta-meta-level that reasons about the metareasoning approach, with the goal of optimizing the

TABLE 1. Summary of metareasoning structures, problems, and modes in multiagent metareasoning papers reviewed by this survey.

Paper	Structure	Problem	Mode
Carrillo et al. [9]	Independent	Multiagent coordination	Select algorithm
Herrmann [17]	Independent	Multiagent coordination	Select algorithm
Raja and Lesser [29]	Independent	Multiagent coordination and Scheduling	Select algorithm
Becker et al. [5]	Independent	Communication	Authorize communication
Xuan et al. [35]	Independent	Communication	Authorize communication
Pinyol and Sabater-Mir [28]	Independent	Belief updating	Modify parameter value
Rubenstein et al. [30]	Independent	Scheduling	Modify parameter value
Borghetti and Gini [6]	Independent	Forecasting	Select algorithm
Noda and Ohta [24]	Independent	Learning	Modify parameter value
Zhang and Lesser [36]	Independent	Learning	Design coordination
Kota et al. [20]	Independent	Task allocation	Redefine relationships
Ahmadi and Allan [1]	Independent	Task allocation	Redefine relationships
Parker [26]	Independent	Task allocation	Modify parameter value
Carlin and Zilberstein [10], [11]	Coupled	Planning	Stop an algorithm
Artikis [4]	Decentralized	Resource allocation	Modify reasoning rules
Brueckner [7], [8]	Multiple metareasoning agents	Forecasting	Design coordination
Pěchouček et al. [27]	Centralized metareasoning	Teaming	Share information

metareasoning; for instance, a meta-meta-level can change the metareasoning policy that the meta-level uses to select the task allocation algorithm that the object-level uses.)

This variety makes directly comparing metareasoning approaches difficult; an approach that was developed for one type of MAS may be irrelevant to another MAS. Ultimately, however, the goal is to improve object-level reasoning, which often involves trade-offs between computational cost and solution quality (and other metrics); as a general rule, one can consider metareasoning whenever such a trade-off appears. In that case, the metareasoning approach should balance these competing objectives and adjust the reasoning as the situation changes. The work reviewed in this survey provides numerous options for doing this.

VI. SUMMARY AND CONCLUSION

Autonomous agents use reasoning algorithms to process sensor information, control their activities, and plan future tasks. Metareasoning monitors and controls an agent's reasoning so that it can make better decisions and achieve its goals more quickly while using fewer computational resources.

Based on a systematic review of the relevant literature, this survey identified and discussed multiple metareasoning structures, the application of metareasoning to different

problems, and the various modes by which metareasoning controls object-level reasoning in MAS. This paper focused on metareasoning that agents perform during MAS operation, and it provides both a narrative review that describes the multiagent metareasoning approaches that have been studied and a scoping review that identifies research gaps in multiagent metareasoning.

Increases in computational power, communication technologies, and collaboration algorithms are driving the development of MAS for more types of missions in multiple domains. Optimizing the reasoning of autonomous agents, which have limited resources, is an active area of research, and metareasoning approaches can help. But there is much more to do.

Based on our review of multiagent metareasoning approaches, we have identified gaps in the literature that can serve as a basis for further research. One gap is the paucity of work that has systematically tested and evaluated metareasoning approaches against each other. Evaluation can be done in simulation models, but evaluating metareasoning approaches implemented on real vehicles is also needed to build confidence and demonstrate their value.

Previous work has implemented metareasoning policies (e.g., which algorithm should be selected given the current state) that were determined by human developers. Future work should consider using machine learning techniques that can automatically develop high-quality metareasoning policies from data about algorithm performance.

Another gap in the literature is the lack of papers that address the general multiagent metareasoning problem of optimizing reasoning performance with limited time and resources. It is possible that a MAS provides benefits such as more information and additional resources, which should yield better metareasoning. Future research in multiagent metareasoning should focus on exploiting the advantages of MAS to perform metareasoning (such as including an agent that performs metareasoning for the other agents).

ACKNOWLEDGMENT

The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL or the U.S. Government.


REFERENCES

- [1] K. Ahmadi and V. H. Allan, "Efficient self adapting agent organizations," in *Proc. ICAART*, 2013, pp. 294–303.
- [2] M. L. Anderson and T. Oates, "A review of recent research in metareasoning and metalearning," *AI Mag.*, vol. 28, no. 1, p. 12, 2007.
- [3] K. Arkoudas and S. Bringsjord, "Metareasoning for multi-agent epistemic logics," in *Computational Logic in Multi-Agent Systems (Lecture Notes in Computer Science)*. Berlin, Germany: Springer, 2005, pp. 111–125, 2005.
- [4] A. Artikis, "Formalising dynamic protocols for open agent systems," in *Proc. 12th Int. Conf. Artif. Intell. Law (ICAIL)*, 2009, pp. 97–104.

- [5] R. Becker, A. Carlin, V. Lesser, and S. Zilberstein, "Analyzing myopic approaches for multi-agent communication," *Comput. Intell.*, vol. 25, no. 1, pp. 31–50, Feb. 2009.
- [6] B. J. Borghetti and M. Gini, "Weighted prediction divergence for metareasoning," in *Proc. Metareasoning*, 2011, pp. 249–264.
- [7] S. A. Brueckner, "Swarming geographic event profiling, link analysis, and prediction," in *Proc. 3rd IEEE Int. Conf. Self-Adapt. Self-Organizing Syst.*, Sep. 2009, pp. 71–81.
- [8] S. A. Brueckner and H. V. D. Parunak, "Applying metareasoning to swarming systems: Approaches and architectures," in *Proc. 4th IEEE Int. Conf. Self-Adapt. Self-Organizing Syst. Workshop*, Sep. 2010, pp. 317–322.
- [9] E. Carrillo, S. Yeotikar, S. Nayak, M. K. M. Jaffar, S. Azarm, M. Otte, and H. Xu, "Communication-aware multi-agent metareasoning for decentralized task allocation," Univ. Maryland, College Park, MD, USA, Tech. Rep., 2020.
- [10] A. S. Carlin, "Decision-theoretic meta-reasoning in partially observable and decentralized settings," Ph.D. dissertation, Univ. Massachusetts Amherst, Amherst, MA, USA, 2012.
- [11] A. Carlin and S. Zilberstein, "Bounded rationality in multiagent systems using decentralized metareasoning," in *Decision Making with Imperfect Decision Makers* (Intelligent Systems Reference Library). Berlin, Germany: Springer, 2012, pp. 1–28.
- [12] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Trans. Robot.*, vol. 25, no. 4, pp. 912–926, Aug. 2009.
- [13] V. Conitzer and T. Sandholm, "Definition and complexity of some basic metareasoning problems," 2003, *arXiv:cs/0307017*. [Online]. Available: <https://arxiv.org/abs/cs/0307017>
- [14] M. T. Cox, "Metacognition in computation: A selected research review," *Artif. Intell.*, vol. 169, no. 2, pp. 104–141, Dec. 2005.
- [15] M. T. Cox and A. Raja, "Metareasoning: An introduction," in *Proc. Metareasoning*, 2011, pp. 3–14.
- [16] A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey," *IEEE Access*, vol. 6, pp. 28573–28593, 2018.
- [17] J. W. Herrmann, "Data-driven metareasoning for collaborative autonomous systems," ISR, UMD, College Park, MD, USA, Tech. Rep. 2020-01, 2020.
- [18] A. Khan, B. Rinner, and A. Cavallaro, "Cooperative robots to observe moving targets," *IEEE Trans. Cybern.*, vol. 48, no. 1, pp. 187–198, Jan. 2018.
- [19] S. Koenig, P. Keskinocak, and C. A. Tovey, "Progress on agent coordination with cooperative auctions," in *Proc. AAAI*, vol. 10, 2010, pp. 1713–1717.
- [20] R. Kota, N. Gibbins, and N. R. Jennings, "Decentralized approaches for self-adaptation in agent organizations," *ACM Trans. Auto. Adapt. Syst.*, vol. 7, no. 1, pp. 1–28, Apr. 2012.
- [21] N. K. Krothapalli and A. V. Deshmukh, "Distributed task allocation in multi-agent systems," in *Proc. Inst. Ind. Eng. Annu. Conf.*, 2002, pp. 1–6.
- [22] L. Ma, H. Min, S. Wang, Y. Liu, and S. Liao, "An overview of research in distributed attitude coordination control," *IEEE/CAA J. Automatica Sinica*, vol. 2, no. 2, pp. 121–133, Apr. 2015.
- [23] S. Nayak, S. Yeotikar, E. Carrillo, E. Rudnick-Cohen, M. K. M. Jaffar, R. Patel, S. Azarm, J. W. Herrmann, H. Xu, and M. Otte, "Experimental comparison of decentralized task allocation algorithms under imperfect communication," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 572–579, Apr. 2020.
- [24] I. Noda and M. Ohta, "Meta-level control of multiagent learning in dynamic repeated resource sharing problems," in *PRICAI 2008: Trends in Artificial Intelligence*. Berlin, Germany: Springer, 2008. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-540-89197-0_29
- [25] G. Paré, M.-C. Trudel, M. Jaana, and S. Kitsiou, "Synthesizing information systems knowledge: A typology of literature reviews," *Inf. Manage.*, vol. 52, no. 2, pp. 183–199, Mar. 2015.
- [26] L. E. Parker, "L-ALLIANCE: Task-oriented multi-robot learning in behavior-based systems," *Adv. Robot.*, vol. 11, no. 4, pp. 305–322, Jan. 1996.
- [27] M. Pechoucek, J. Tozicka, and V. Marík, "Meta-reasoning methods for agent's intention modelling," in *Autonomous Intelligent Systems: Agents and Data Mining*. Berlin, Germany: Springer, 2005. [Online]. Available: https://link.springer.com/chapter/10.1007/11492870_11
- [28] I. Pinyol and J. Sabater-Mir, "Metareasoning and social evaluations in cognitive agents," in *Autonomic Computing and Communications Systems* (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering). Berlin, Germany: Springer, 2010, pp. 220–235.
- [29] A. Raja and V. Lesser, "A framework for meta-level control in multi-agent systems," *Auto. Agents Multi-Agent Syst.*, vol. 15, no. 2, pp. 147–196, Aug. 2007.
- [30] Z. B. Rubinstein, S. F. Smith, and T. L. Zimmerman, "The role of metareasoning in achieving effective multiagent coordination," in *Proc. Metareasoning*, 2011, pp. 217–232.
- [31] S. J. Russell and E. Wefald, *Do the Right Thing: Studies in Limited Rationality*. Cambridge, MA, USA: MIT Press, 1991.
- [32] M. Senanayake, I. Senthoooran, J. C. Barca, H. Chung, J. Kamruzzaman, and M. Murshed, "Search and tracking algorithms for swarms of robots: A survey," *Robot. Auto. Syst.*, vol. 75, pp. 422–434, Jan. 2016.
- [33] G. Weiss, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Cambridge, MA, USA: MIT Press, 1999.
- [34] M. J. Wooldridge, *An Introduction to MultiAgent Systems*. Chichester, U.K.: Wiley, 2009.
- [35] P. Xuan, V. Lesser, and S. Zilberstein, "Communication decisions in multi-agent cooperation: Model and experiments," in *Proc. 5th Int. Conf. Auto. Agents (AGENTS)*, 2001, pp. 616–623.
- [36] C. Zhang and V. Lesser, "Coordinating multi-agent reinforcement learning with limited communication," in *Proc. Int. Conf. Auton. Agents Multi-Agent Syst.*, 2013, pp. 1101–1108.



SAMUEL T. LANGLOIS was born in Pitman, NJ, USA, in 1996. He received the B.S. degree in biomedical engineering from Rutgers University–New Brunswick, New Brunswick, NJ, USA, in 2019. He is currently pursuing the M.S. degree in systems engineering with the University of Maryland, College Park, MD, USA. He is also a Research Assistant with the Institute of Systems Research, University of Maryland, under the supervision of Dr. J. Herrmann. His research interest includes data-driven applications of metareasoning in multiagent systems. He is also a member of INCOSE.



OGHENETEKEVWE AKORODA (Member, IEEE) received the B.S. degree in mechanical engineering from Howard University, Washington, DC, USA, in 2019. He is currently pursuing the M.S. degree in systems engineering with the University of Maryland, College Park, MD, USA. He is also a Research Assistant with the Institute of Systems Research, University of Maryland, under the supervision of Dr. J. Herrmann. His research interest includes applications of metareasoning in multiagent systems. He is also a member of INCOSE, NSBE, and ASME.



ESTEFANY CARRILLO (Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from the University of Maryland, College Park, MD, USA, in 2012 and 2017, respectively, where she is currently pursuing the Ph.D. degree in aerospace engineering.

She is also a Research Assistant with the Department of Aerospace Engineering, UMD, under the supervision of Dr. H. Xu. Her research interests include formal methods and hybrid systems theory in the design of verifiable controllers for complex high-level tasks and control of multi-agent systems.



JEFFREY W. HERRMANN received the B.S. degree in applied mathematics from the Georgia Institute of Technology, Atlanta, GA, USA, in 1990, and the Ph.D. degree in industrial and systems engineering from the University of Florida, Gainesville, FL, USA, in 1993.

He is currently a Professor with the University of Maryland, College Park, MD, USA, where he holds a Joint Appointment in the Department of Mechanical Engineering and the Institute for Systems Research. He is also the author of the textbook *Engineering Decision Making and Risk Management* (Wiley, 2015). His current research interests include collaborative search and tracking and engineering design decision making. He is also a member of ASEE, IISE, ASME, the Design Society, and INFORMS.



HUAN XU (Member, IEEE) received the B.S. degree in mechanical engineering and material science from Harvard University, Cambridge, MA, USA, in 2007, and the M.S. and Ph.D. degrees in mechanical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2008 and 2013, respectively.

She is currently an Assistant Professor with the Department of Aerospace Engineering, University of Maryland, College Park, MD, USA, and the Institute for Systems Research, University of Maryland. Her research interests include formal methods and controls in the design and analysis of unmanned autonomous systems. She is also a member of AIAA, AUVSI, and INCOSE. She is also a member of the Maryland Robotics Center and a Subject Matter Expert for the UMD UAS Test Site.



SHAPOUR AZARM received the B.S. degree in mechanical engineering from the University of Tehran, Tehran, Iran, in 1977, the M.S. degree in mechanical engineering from George Washington University, Washington, DC, USA, in 1979, and the Ph.D. degree in mechanical engineering from the University of Michigan, Ann Arbor, MI, USA, in 1984.

He joined the University of Maryland, College Park, in 1984, where he is currently a Professor of mechanical engineering, and a Professor of applied mathematics & statistics, and scientific computation program. His research interests include engineering optimization and decision-making models and methods. He is also a Review Editor and a Senior Advisor of (journal of) *Structural and Multidisciplinary Optimization*. He was the Editor-in-Chief of the *Journal of Mechanical Design* (ASME Transactions) and an Associate Editor of (journal of) *Mechanics-Based Design of Structures and Machines*.



MICHAEL OTTE (Member, IEEE) received the B.S. degree in aeronautical engineering and computer science from Clarkson University, Potsdam, NY, USA, in 2005, and the M.S. and Ph.D. degrees in computer science from the University of Colorado Boulder, Boulder, CO, USA, in 2007 and 2011, respectively.

From 2011 to 2014, he was a Postdoctoral Associate with the Massachusetts Institute of Technology. From 2014 to 2015, he was a Visiting Scholar with the U.S. Air Force Research Laboratory. From 2016 to 2018, he was a National Research Council RAP Postdoctoral Associate with the U.S. Naval Research Laboratory. He has been with the Department of Aerospace Engineering, University of Maryland, College Park, MD, USA, since 2018. He is the author of over 30 articles. His research interests include autonomous robotics, motion planning, and multi-agent systems. He is also a member of the American Institute of Aeronautics and Astronautics, and the IEEE Robotics and Automation Society. He has served as an Associate Editor for IEEE conferences, including ICRA and IROS.

...