Do We Need "Teaming" to Team with a Machine?

Craig Haimson, Celeste Lyn Paul, Sarah Joseph, Randall Rohrer, and Bohdan Nebesh

U.S. Department of Defense, USA
{crhaims,clpaul,skjosep,rmrohre,banebes}@tycho.ncsc.mil

Abstract. What does it mean for humans and machines to work together effectively on complex analytic tasks? Is human teaming the right analogue for this kind of human-machine interaction? In this paper, we consider behaviors that would allow next-generation machine analytic assistants (MAAs) to provide context-sensitive, proactive support for human analytic work – e.g., awareness and understanding of a user's current goals and activities, the ability to generate flexible responses to abstractly-formulated needs, and the capacity to learn from and adapt to changing circumstances. We suggest these behaviors will require processes of coordination and communication that are similar to but at least partially distinguishable from those observed in human teams. We also caution against over-reliance on human teaming constructs and instead advocate for research that clarifies the functions these processes serve in enabling joint activity and determines the best way to execute them in specific contexts.

Keywords: Human-Machine Symbiosis and Human-Machine Interface, Human-Machine Teaming

1 Introduction

The more sophisticated a system's ability to perform complex tasks in coordination with its users, the more it seems to function as something more akin to a human work partner than a mere tool [1]. Such collaboration between humans and technology is often referred to as human-machine teaming (HMT) (e.g., [2, 3]) because of its resemblance to human teamwork. HMT has been heralded as the key to transforming automation-enabled work practices across a number of domains, many of which are of critical importance to national defense [e.g., 4]. But just how important are human teamwork behaviors to HMT? Is a more human-like machine teammate necessarily a better machine teammate? This paper focuses on the potential role of HMT in one domain – intelligence analysis – and explores the extent to which human teamwork is an appropriate model for these HMT use cases.

Intelligence analysis comprises a set of interrelated activities that generate evidence-based information products from collected information, often with the goal of answering critical questions about adversaries' attributes, associations, beliefs, intentions, and actions. Many of these activities depend upon a human analyst's ability to find and fuse information acquired across multiple heterogeneous datasets. With the volume, velocity, and variety of data constantly growing, intelligence professionals require increasingly sophisticated tools to enable them to keep pace with fast moving events and the signatures these events generate. Advances in machine learning are driving the development of new analytic technologies capable of recognizing and responding to meaningful patterns across these datasets; however, to maximize their utility to analysts, these technologies must be managed by intelligent software agents that can deploy analytics in a coordinated fashion on behalf of human analysts who constrain, shape, and consume the consolidated results of their activities. We refer to these software agents as machine analytic assistants (MAAs) and envision that they will facilitate intelligence analysis by collaborating with their human partners on shared analytic projects.

HMT research often draws on concepts from human teaming to inform and ground HMT principles [e.g., 1, 3] since there is a wealth of research exploring effective human teaming processes. We will discuss several such processes and suggest how they may facilitate joint analytic work by human analysts and MAAs. However, we will also identify some ways in which human-machine teams may differ from their all-human counterparts, at least for the intelligence analysis use cases with which our research is concerned. Based on this assessment, we suggest that HMT research should not set itself the task of fully emulating human teamwork but should instead just focus on determining the functions and interaction designs that maximize the overall efficiency and effectiveness of joint human-machine work.

2 Envisioned Characteristics of MAAs

Intelligence analysis typically entails multiple iterative tasks involving searching, filtering, evaluating, and fusing information contained within a large number of sources, driven initially by broad exploratory goals that evolve into more focused, hypothesis-driven objectives [5]. Analysts may need to use a variety of tools and methods to find relevant information, expose associations between entities of interest, and preserve key results for further analysis or reporting. Moreover, they must often perform many of these tasks manually, which can severely limit the quantity of data they can consider and amount of information they can extract and synthesize.

MAAs will support the intelligence process by coordinating activities of data analytics that help human analysts find and combine information in ways that satisfy intelligence requirements, similar to software agents that support other forms of exploratory data analysis [e.g., 6]. Machine analytics operating on text [e.g., 7], images [e.g., 8], or other media can classify and cluster data, identify important concepts and relationships, identify anomalies, and reveal and quantify key trends. MAAs will serve as intelligent gateways to these powerful capabilities, assembling and executing multi-analytic workflows to generate summarized findings that meet analyst needs, both by responding to analysts' explicitly expressed requests and also supplying additional recommendations based on knowledge of analysts' mission goals and analytic history. In these ways, MAAs will help analysts find and organize data for efficient and effective review and assessment. MAAs will expand their repertoire of behaviors by learning new workflows and conditions of use, either through passive observation of users' actions or active participation in demonstration/training sessions with users [e.g., as in 9, 10].

Note that although MAAs will utilize logical reasoning to support data-driven inference and other logical functions in pursuit of these activities, they are unlikely to possess a level of knowledge or cognitive sophistication required for more than rudimentary analysis and reporting tasks; moreover, the nuances of the legal policies that govern the conduct of intelligence analysts [e.g., 11, 12, 13, 14] are too context-dependent, and the potentially disastrous consequences of automation failure are too severe to allow for anything beyond this in the foreseeable future, whatever the degree of artificial intelligence achieved. Thus, MAAs will not replace human analysts but will instead assume their more time-consuming and laborious information retrieval and manipulation tasks, freeing analysts to devote more time to interpretation (although MAAs could potentially also help structure the analysis process itself, as in [15]).

Achieving these modest yet still ambitious goals will require advances in software agent technologies that afford MAAs the ability to:

- Learn and understand users' evolving goals and maintain an awareness of available data and analytics that can satisfy them.
- Orchestrate and execute actions with flexibility, responding appropriately as requirements and results accumulate and change over time.
- Recognize when results reveal opportunities for useful follow-on analysis and then program and execute a new set of actions to exploit those opportunities in accordance with existing guidelines and constraints.
- Operate with a reasonable degree of independence to insulate users from an otherwise constant barrage of requests for input and validation while still making the most of available data and computational resources. This includes identifying conflicting goals and resolving lower level conflicts to avoid wasting time and computational resources on lower priority tasks.

Ideally, MAAs will have the ability to evaluate circumstances and consult policies that help them determine whether and when to request approval for a given task they will complete on their own Such an ability would minimize user involvement in routine tasks while ensuring a human is "in the loop" for more complicated, riskier decision making (e.g., about tasks that are resource intensive) and deeper, contextually-dependent analysis (e.g., about results that may have multiple interpretations or important national security implications). Johnson et al. [16] recommend a "combine and succeed" approach to allocating work between humans and automation, allowing for varying degrees of human involvement and multiple ways of achieving a given task based on circumstances. We believe this approach will be essential to the effective use of MAAs, which will vary in the level of autonomy with which they identify current information needs, decide which combination(s) of data and analytics are most likely to satisfy these needs, configure and execute analytics against appropriate data sources, and manipulate and interpret results (see [17] for a discussion of varying levels of autonomy across similar classes of tasks).

3 Human Teaming Behaviors and MAAs

The degree of interdependence envisioned between human and MAA tasks and the need for human-machine interactions that coordinate interdependent human-MAA work will create MAA HMT challenges. Research must address these challenges to ensure the successful application of MAAs to intelligence analysis. The complexity of the agents' behaviors will require that MAAs at times operate with different task sub-goals than their users currently hold while still working towards common overall objectives. In turn, the potential for MAAs to operate with different sub-goals than their users will create a need for users to ensure automation is aligned with their own situational understanding and mission priorities, in order to regulate use of computational resources, prevent adverse MAA activities, and maximize the fit and utility of machine contributions. Achieving these objectives will require the development of sophisticated teaming functions and human-machine interaction methods that allow both analysts and MAAs to coordinate their activities efficiently and effectively.

Human teaming seems a natural analogue for joint activity involving multiple autonomous-yet-interdependent actors. Decades of work in industrial and organizational psychology and management science have produced a rich literature on human team processes [e.g., 18, 19, 20], and there are many important lessons to be learned from this research about what makes human teams function effectively. We see a number of parallels between human-MAA teaming and human teaming, and we believe principles of human teaming should inspire and inform MAA HMT research and development; however, we do not feel that all principles of human teaming are equally relevant to MAA HMT and/or should necessarily be expressed in human-MAA interactions the same way they are in human teams.

Consider work by Salas and colleagues [18]. They conducted an extensive review and thematic analysis of two decades of human teaming research and identified five major factors that appear to affect the success of human teams: team leadership, mutual performance modeling, backup behavior, adaptability, and team orientation. These factors, along with the coordinating mechanisms of shared mental models, mutual trust, and closed-loop communication, all appear to be important for successful human teamwork involving either collaboration (team members work on a common task) or coordination (team members work on separate interdependent tasks contributing to a shared outcome). Tables 1 and 2 discuss Salas et al.'s human teamwork factors (Table 1) and coordinating mechanisms (Table 2), along with our observations regarding their applicability to MAA HMT.

As shown, many aspects of Salas et al.'s framework are highly applicable to MAA HMT use cases. There is solid evidence demonstrating the importance of these factors and coordinating mechanisms for human teaming, and it seems clear they will be important for MAA HMT as well. For example, functions that support mutual performance monitoring will allow analysts and MAAs to detect each other's errors, mitigating their impact. Similarly, capabilities that facilitate the development of shared mental models will enable analysts and MAAs to better understand each other's information needs, which should encourage proactive sharing and more efficient communication.

Table 1. Salas et al. [18] Teamwork Factors and Applicability to MAA HMT

Factor	Description	Applicability to MAA HMT
Team leadership	Planning, assigning, coordinating, and facilitating team activities in accordance with knowledge of evolving objectives and conditions. Also involves evaluating, developing, and encouraging team personnel.	User will direct MAA by providing goals, constraints, and feedback based on mission objectives and understanding of context. User will develop MAA by providing feedback on correctness/utility of its work, and teaching it new analytic procedures. MAA could plan/assign/coordinate some tasks and even instruct junior users. ****Not applicable: User will not need to motivate MAA.**
Mutual performance monitoring	Maintaining awareness of teammate performance to assess needs and identify errors.	User and MAA will monitor each other's performance to infer teammate's goals, plans, and needs; identify disagreements in priorities or interpretation of data; and detect errors in teammate's decisions or actions.
Backup behavior	Taking over some of a teammate's tasking to provide relief during periods of high workload. Also involves proactively offering information or support in anticipation of a teammate's needing it, or providing feedback when a teammate commits errors or has difficulty performing a task.	User and MAA will provide corrective feedback when they identify errors in each other's performance, and they will proactively offer information in anticipation of each other's needs. ****Not applicable: Workload rebalancing. User will never perform a task for MAA as long as MAA knows how to do it, and if MAA can execute a task for user, it will always do so.**
Adaptability	Modifying team work plans and processes based on changing needs and circumstances.	User and MAA will tailor analytic methods and MAA's level of autonomy according to changing complexity and uncertainty in requirements, data, and results.
Team orientation	Considering teammates' perspectives and needs, effectively leveraging teammates' contributions to achieve one's own tasks, and valuing team's success over one's own self-interest.	User will need to accept MAA's help and utilize it effectively, which could be complicated by fears that automation is "taking over" analysis process. ****Not applicable: MAA will not have interests or goals apart from those of user; thus, it is not clear either user or MAA would need to adopt the kind of collective outlook human teaming requires.**

Table 2. Salas et al. [18] Coordinating Mechanisms and Applicability to MAA HMT

Factor	Description	Applicability to MAA HMT
Shared mental models	Knowledge of a team's goals and tasks, as well as the dependencies that exist between them.	Both user and MAA will need to understand how their tasks affect each other's work and contribute to joint goals; will enable task coordination and anticipation of each other's information needs.
Mutual trust	Trusting that teammates will competently and conscientiously execute assigned tasks, accept and respect each other's contributions, and act in ways that benefit team.	User will need to have sufficient trust in MAA's competence to allow it to work independently. ****Not applicable: It is not clear that MAA's programmed acceptance of user's commands would constitute trust, and it is also not clear that user's trust in MAA would be the same as their trust in a human teammate. Trust in a human teammate includes believing the teammate will not act in ways that promote their self-interest at the team's expense; MAAs will not have self-interest, so user should not suspect MAA's "motives."**
Closed-loop communication	Communication in which communicants confirm they have received and understood each other's messages.	User and MAA will need to engage in closed-loop communication to ensure messages are received and interpreted correctly. MAA will also need to communicate its inferences about user's goals and needs (user should not need to infer MAA's goals and needs, as MAA will always communicate these explicitly). Will require mechanisms for user and MAA to identify and correct miscommunications or misunderstandings.

In contrast, the notion of workload-related backup seems less relevant to MAA HMT. This is partly because we assume MAAs and human analysts will generally perform different kinds of work, which would preclude their taking over each other's excess tasking (i.e., human-MAA teamwork will be more coordination-based than collaboration-based). However, we also expect MAAs and human analysts to be differentially affected by analytic workload. Machines do not have the same kinds of processing limitations as humans, so it is hard to imagine a situation in which a human would need to take over tasking from an MAA purely to lighten its workload, or in which an MAA would not perform a task for its user if it knows how to do it. It is possible that more collaborative HMT (e.g., humans and robots working together in urban search and rescue) would allow for workload-related backup, but it would seem to require that both human and machine teammates experience the same kinds of processing limitations (e.g., only being able to be in one physical location at a time). Thus, task rebalancing appears to expose one way in which human and machine teammates (and their associated teaming styles) may differ.

The more social aspects of team leadership, team orientation, and mutual trust highlight additional (and arguably more dramatic) differences that will exist between human and machine teammates. Software does not have an independent sense of its own interest; thus, factors like mutual trust, which helps ensure human teammates are willing to cooperate even when cooperation brings additional risks (e.g., the potential that teammates will act in ways that further their own interest at the expense of others'), seem inapplicable to machines. Moreover, while an analyst's willingness to work with and trust an MAA could be thought of as somewhat akin to team orientation and mutual trust, it is not clear the underlying constructs are the same. Although users may interact with machines in ways that resemble their interactions with humans, humans and machines are fundamentally different types of entities, and users fundamentally know this; thus, it not clear humans have the same underlying thoughts and feelings when they interact with machines as they do with humans. For example, humans need not concern themselves with the social costs of mistreating automation (other than ones they might experience if other humans observe them) and probably cannot truly empathize with or expect empathy from machines that do not experience life in the same way they do.

Our review of Salas et al.'s [18] framework suggests that MAA HMT will resemble human teaming but differ from it as well, especially when it comes to aspects of teaming that seem more dependent on teammates' core natures and the types of relationships they can form with each other. These kinds of differences need not prevent researchers from exploring the benefits of partnering humans with automation or from drawing inspiration from the human teaming literature in developing HMT capabilities; however, they do suggest the research community might think twice before assuming HMT need necessarily be the same as human teaming or that the interaction methods that support HMT need necessarily emulate human teammate interactions.

4 Discussion

We believe the future of intelligence analysis depends on the development of MAAs that can partner with analysts to improve the efficiency and effectiveness with which they exploit available data; without it, analysts will not realize the full benefits of analytics that have the potential to save them from extreme information overload. MAAs will be able to act with a great deal of independence and flexibility, thus eliminating much of the manual work currently associated with the use of analytic tools. However, commanding and controlling such technologies will pose challenges that must be addressed if intelligence analysis is to make the most of these technologies.

Many HMT researchers have based their recommendations on prior research in human teaming [e.g., 16, 21], and we are beginning to investigate ways in which lessons from the human teaming literature may be applied to HMT for MAAs. For example, we are exploring the kinds of dialogues analysts may need to engage in with automation to establish common ground and coordinate activity through negotiation. Note that these are among the top challenges Klein and colleagues [21] discuss with regard to HMT. We expect MAAs will participate in collaborative planning with their users as in [22] to establish joint high level goals and agree upon tasks to be performed separately in support of shared objectives. We also expect MAAs will communicate with their users about their interpretations of results, offering, defending, and/or challenging different explanations and hypotheses they or their users propose regarding the meaning or implications of a piece of data or analytic finding. Foundational work on dialogues between software agents [e.g., 23] will provide a basis for some of these dialogues about what actions to take and what results mean, but research from informal logic and human discourse analysis [e.g., 24] may also provide important background.

Although we have a strong interest in human teaming, we do not assume that all human teaming behaviors are relevant to HMT, or that the behaviors that are relevant need necessarily be expressed through styles of interaction that resemble human social interaction. On the contrary, we feel there is a critical need for research into understanding what form these dialogues and the processes that underlie them need to be to maximize efficiency and effectiveness of joint work in a particular domain. Assuming the research community should strive to fully emulate human teams imposes a daunting set of research requirements that may be unnecessary or even counterproductive if the end goal is optimizing human-machine task performance versus creating synthetic equivalents of machines' human counterparts (see discussion by [25]). It may seem reasonable to assume humans will be most comfortable and effective engaging with machines in a manner that mimics their interactions with other humans. However, treating tools like people is a fairly recent phenomenon in our species' natural history, and it seems just as reasonable to assume humans will be most comfortable and effective interacting with machines in ways that satisfy the computational requirements of joint work while maintaining distinct roles for users and the things they use. These interactions may not constitute "true" teaming in the human sense, but users may not need true teaming to work with automation successfully: They may only need something that accomplishes the teaming functions required for successful use of the technology.

We do not exclude the possibility that there will be circumstances in which effective human-machine coordination requires that humans are able to interact with machines in ways that closely resemble how they would interact with human teammates. However, rather than starting with the question of how to emulate human-like teaming behaviors in human-machine teams, we propose the research community instead begin by asking what functions are necessary for coordinating the tasks humans and machines will be performing in support of common goals. It can then explore the efficacy of different methods of instantiating these functions, clearly distinguishing computational goals from algorithms and means of implementation [26]. Some methods may have clear advantages, while others may be roughly equivalent in terms of efficiency and effectiveness measures, allowing choice of methods to be driven largely by development costs, user preference, and fit with operational systems and settings. We suspect the methods best suited to enabling HMT will vary with the types of tasks on which humans and machines are collaborating, and we suggest that understanding the features that cause different types of tasks to require different methods should be a research priority. Understanding these features will enable the community to generalize findings from one domain to another.

There may be scientific and practical benefits to more basic HMT research that seeks to replicate human teaming in human-machine teams in as direct and authentic a way as is possible, not the least of which would be gaining further insight into the nature of teamwork itself. However, at this time, we believe a more applied research agenda that treats HMT as a means to an end rather than an end in itself holds more promise for delivering solutions that best achieve successful coupling of humans and machines on a particular set of tasks. If a given HMT approach enables analysts to use MAAs to perform their work more efficiently and effectively, it will have achieved its purpose.

References

1. Lyons, J.B., Mahoney, S., Wynne, K.T., Roebke, M.A. Viewing machines a teammates: A qualitative study. AAAI Spring Symposium Series, pp. 166-170, Palo Alto, CA (2018).
2. McDermott, P.L., Walker, K.E., Dominguez, C.O., Nelson, A., Kasdaglis, N. Quenching the thirst for human-machine teaming guidance: Helping military systems acquisition leverage cognitive engineering research. In: Gore, J., Ward, P. (eds.) 13th International Conference on Naturalistic Decision Making, pp. 236-240, Bath, UK (2017).
3. Parasuraman, R., Barnes, M., Cosenzo, K. Adaptive automation for human-robot teaming in future command and control systems. The International C2 Journal 1(2), 43-68 (2007).
4. Ryan, M. Human-machine teaming for future ground forces. Center for Strategic and Budgetary Assessments (2018).
5. Pirolli, P., Card, S.K. The sensemaking process and leverage points for analyst technology. Proceedings of the 2005 International Conference on Intelligence Analysis (2005).
6. St. Amant, R., Cohen, P. Intelligent support for exploratory data analysis. Journal of Computational and Graphical Statistics 7(4), 545-558 (1998).
7. Jurasfky, D., Martin, J.H. Speech and Language Processing. 2nd edn. Prentice Hall, Upper Saddle River, NJ (2008).
8. Li, S.Z., Jain, A. Handbook of Face Recognition 2nd edn. Spring-Verlag, London, UK (2011).
9. Allen, J., Chambers, N., Ferguson, G., Galescu, L., Jung, H., Swift, M., Taysom, W. PLOW: A collaborative task learning agent. Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07), pp. 1514-1519 (2007).
10. Azaria, A., Krishnamurthy, J., Mitchell, T.M. Instructable intelligent personal agent. Proceedings of the Thirtieth Conference on Artificial Intelligence (AAAI-16), pp. 2681-2689 (2016).
11. Reagan, R.W. Executive Order 12333: United States Intelligence Activities, U.S. Federal Register (1981).
12. Bush, G.W., Executive Order 13284: Amendment of Executive Orders, and Other Actions, in Connection With the Establishment of the Department of Homeland Security, Federal Register (2003.

13. Bush, G.W. Executive Order 13355: Strengthened Management of the Intelligence Community, Federal Register (2004).
14. Bush, G.W. Executive Order 13470: Further Amendments to Executive Order 12333, United States Intelligence Activities, Federal Register (2008).
15. Tecuci, G., Marcu, D., Boicu, M., Schum, D. COGENT: Cognitive agent for cognitive analysis. AAAI Fall Symposium Series, pp. 58-65, Arlington, VA (2015).
16. Johnson, M., Bradshaw, J.M., Hoffman, R.R., Feltovich, P.J., & Woods, D.D. (2014). Seven cardinal virtues of human-machine teamwork: Examples from the DARPA Robotic Challenge. IEEE Intelligent Systems, 74-70 (2014).
17. Parasuraman, R., Sheridan, T.B., Wickens, C.D. A model for types and levels of human interaction with automation. Proceedings of IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans 30(3), 286-297 (2000).
18. Salas, E., Sims, D.E., Burke, C.S. (2005). Is there a "big five" in teamwork? Small Group Research 36(5), 555-599 (2005).
19. Marks, M.A., Mathieu, J.E., Zaccaro, S.J. A temporally based framework and taxonomy of team processes. Academy of Management Review 26, 356-376 (2000).
20. Tuckman, B.W. Developmental sequence in small groups. Psychological Bulletin 63, 384-399.
21. Klein, G., Woods, D.D., Bradshaw, J.M., Hoffman, R.R., Feltovich, P.J. Ten challenges for making automation a "team player" in joint human-agent activity. IEEE Intelligent Systems, 91-95 (2004).
22. Ferguson, G., Allen, J. A cognitive model for collaborative agents. AAAI Fall Symposium Series, pp. 112-120, Arlington, VA (2011).
23. Parsons, S., Wooldridge, M., Amgoud, L. An analysis of formal interagent dialogues. In: Castelfranchi, C., Johnson, W.L. (eds.) Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems, pp. 394-401, New York, NY (2002).
24. Walton, D., Macagno, F. Profiles of dialogue for relevance. Informal Logic 36(4), 523-562 (2016).
25. Groom, V., Nass, C. Can robots be teammates? Benchmarks in human-robot teams. Interaction Studies 8, 483-500 (2007).
26. Marr. D. Vision: A Computational Approach. Freeman & Co., San Francisco, CA (1982).