12TH ICCRTS
"Adapting C2 to the 21st Century"

Introducing a Human-Automation Collaboration Taxonomy (HACT) in Command and Control Decision-Support Systems

C2 Metrics and Assessment;
C2 Concepts, Theory and Policy;
Network-Centric Experimentation and Applications

Sylvain Bruni*

Jessica J. Marquez*

Amy Brzezinski*

Carl Nehme*

Yves Boussemart*

* Student, Massachusetts Institute of Technology

This is a student paper submission.
Work supervised by Prof. Mary L. Cummings,
Massachusetts Institute of Technology

Point of Contact: Sylvain Bruni
617.417.0359 sbruni@mit.edu

Massachusetts Institute of Technology
Humans and Automation Laboratory
77 Massachusetts Avenue, Bldg 33-407
Cambridge, MA 02319

Introducing a Human-Automation Collaboration Taxonomy (HACT) in Command and Control Decision-Support Systems

Sylvain Bruni, Jessica J. Marquez, Amy Brzezinski, Carl Nehme, and Yves Boussemart
Massachusetts Institute of Technology, Cambridge, MA 02139

A significant trend in command and control human-systems integration is the design of decision-support systems that attempt to facilitate collaboration between the human operator and automation. Unfortunately, there currently is no framework that explicitly describes different aspects of human-computer collaboration. There is a large body of literature dedicated to enumerating both agents' respective roles in the decision-making process, yet none has explicitly encompassed higher level characteristics of collaboration other than the function allocations of humans and automated systems. We propose here a framework that more accurately portrays collaborative decision-support systems beyond simply role allocation, termed the Human-Automation Collaboration Taxonomy or HACT. HACT provides a structure to characterize and determine the degree to which a decision-support system is collaborative, for evaluation and comparison purposes. We identified three basic roles "Moderator", "Generator", and "Decider"; three primary system characteristics "Functional Transparency", "Information Transparency", and "Interactivity"; and one secondary system characteristic "Adaptability". This paper presents in detail these fundamental dimensions, and describes an application of this framework to a collaborative decision support system developed to assist human operators in the domain of mission planning.

I. Introduction

As technology has advanced, tasks that were previously performed solely by humans have evolved to include automation. Systems have become more complex and in turn, system designers have developed more sophisticated automation to either assist human operators or take over tasks that were assigned to human operators. This is especially true for decision support systems. Unfortunately, as new methods of human-computer collaboration emerge, no single framework effectively describes these as a whole. The lack of a human-automation collaboration taxonomy prevents decision support system designers from accurately comparing systems, including the advantages or disadvantages of particular designs. Currently, other frameworks do not consider human-automation systems with different types of collaboration in decision-making, which include not just functional role allocations of the humans and automation, but also the methods of communication and the adaptability of the system. Therefore, to better describe decision support systems beyond basic functional allocation, including those in the field of command and control, we introduce the Human-Automation Collaboration Taxonomy, or HACT.

The taxonomy presented in this paper is generalized to encompass human-computer collaboration within the context of decision support systems. Within command and control activities, we define collaboration as "the mutual engagement of agents in a coordinated and synchronous effort to solve a problem based on a shared conception of it" (adapted from (Dillenbourg, Baker et al. 1995)). An agent may be a human operator or an automated computer system, or "automation". HACT is only based on interactions between two agents (a human operator and automation). We also restrict the discussion of HACT to the analysis and decisions phases of the information processing model (Figure 1) and not the acquisition and action phases (Parasuraman, Sheridan et al. 2000). Typically, human-automation collaboration is an iterative process between the agents, and between the analysis and decision steps, which will be addressed in more detail in the next section.

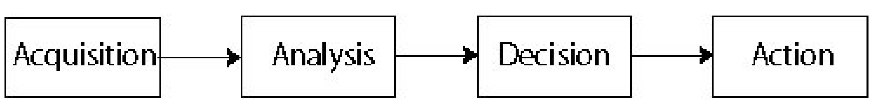

Figure 1. Four-step information processing flow (Parasuraman, Sheridan et al. 2000).

This paper first discusses existing automation taxonomies and how their scales do not fully represent human-computer collaboration elements. Next, HACT is presented, along with a supporting visualization. The taxonomy is then applied to a command and control decision support system example, in order to illustrate the usefulness in understanding and comparing different collaboration designs for a single decision support tool. Finally, we discuss how HACT addresses the shortcomings of previous automation taxonomies, and the usefulness of HACT as a potential design and cost analysis tool.

II. Background

Several taxonomies aimed at classifying and describing interactions between a human operator and a computer system exist, and they are generally based on the concept of "level of automation". Despite some variations, these levels of automation, or LOAs, refer to the role allocation between the automation and the human (Sheridan and Verplank 1978; Wickens, Mavor et al. 1998; Parasuraman, Sheridan et al. 2000). These LOAs emphasize particular attributes, such as authority in the decision making process, solution generation abilities, or scope of action. The relative importance of each attribute can vary tremendously across command and control systems, hence, several scales have emerged, each typically focusing around one or two specific attributes.

The widely used Sheridan and Verplank (SV) scale of levels of automation ranges from a fully manual system with no computer intervention, to a fully automated system where the human is kept out of the loop (Sheridan and Verplank 1978). Parasuraman et al. (2001) expanded the original SV LOA to include ten levels. At the lower levels, LOAs 1 through 4, the human is very active and involved in the decision-making and solution-generating processes, from finding a solution to sorting through possible alternatives suggested by the computer. Starting at level 5, the automation takes on a more active role in executing solutions, while still requiring consent from the operator before doing so (known as management-by-consent). The next level up, level 6, is typically referred to as management-by-exception and at this level, the automation commences a more active role in decisions, executing solutions unless vetoed by the human. Subsequently, levels 7 to 10, human operators are only allowed to accept or veto solutions presented to them. As levels increase, the human is progressively taken out of the decision-making loop, and the automation is given the authority to inform (or not) the human of its decisions. This scale addresses authority allocation, from human-controlled systems to automation controlled ones. However, although this LOA scale provides a range of possible distributions of authority in decisions making between the human and the automation, it provides a much smaller range for solution generation. The scale also falls short of describing other collaboration characteristics such as the different methods of communication between the human and the automation that include feedback or involve multiple solution iterations.

Another ten-point scale was created by Endsley and Kaber (1999) where each level corresponds to a specific task behavior of the automation, going from "manual control" to "full automation", through intermediate levels such as "blended decision-making" or "supervisory control". This scale examines four functions (analogous to the Parasuraman et al. (2001) information processing model in Figure 1) present in command and control applications, specifically "monitoring systems status", "generating solutions", "selecting solutions", and "implementing solutions". Endsley and Kaber's LOA scale introduces the idea of breaking down the decision-making process into specific phases and how the introduction of automation differs in each of these phases. However, just like the SV scale, this scale also ignores the alternate communication methods that are possible between the human and automation.

Other similar LOA taxonomies (e.g., Endsley 1987; Riley 1989) describing how human operators and automation can be balanced within decisions support systems have been proposed, but unfortunately, they also only consider a subset of attributes. Endsley (1987) investigated the introduction of artificial intelligence into a five point LOA scale while Riley (1989) investigated the use of the level of information attribute in addition to the automation authority attribute creating a two dimensional scale. The limited set of attributes limits the applicability and scope of such one- or two-dimensional scales.

There are certain elements of human-computer collaboration that are not addressed in any of the LOAs reviewed. First, there is no mention of methods of making the automation more transparent to the operator, essential for maintaining mode awareness and detecting automation errors (Billings 1997). Second, the exchange of information between agents is important in any form of collaboration. Many systems claim to be collaborative but the manner in which information is exchanged cannot be described as "mutual engagement," which is a key attribute for collaboration. Finally, systems where the level of automation could change with time either through human actions (adjustable autonomy) or independently (adaptive autonomy) are not considered (Goodrich, Johansen et al. 2007). This unique characteristic of a potential decision support system should be considered as a step towards more elaborate forms of human-automation collaboration. HACT takes into account both the important attributes highlighted by previous LOA and these missing attributes.

III. The Human-Automation Collaboration Taxonomy (HACT)

A. A revised decision-making process framework

In order to better understand how human operators and automation collaborate, the information-processing flow diagram of Figure 1 was modified to focus on the specifics of collaborative decision-making. This new model, shown in Figure 2, features three steps: data acquisition, decision-making and action taking. The human-automation collaboration taxonomy presented in this paper applies directly to the decision-making phase of this model.

The data acquisition step is similar to that proposed by Parasuraman et al. (2000): sensors get the information from the outside world or environment, and transform it into working data. The collaborative decision-making process corresponds to the integration of the analysis and decision phases of Figure 1. First, the data from the previous step is analyzed, possibly in an iterative way where request for more data is sent to the sensors. The data analysis outputs some elements of a solution to the problem at hand. For example, in a mission planning situation, these elements of solutions may correspond to the current or projected status of some battlefield assets. The evaluation block will estimate the appropriateness of these elements of solutions for a potential solution. This block may initiate a recursive loop with the data analysis block. For instance, it may request more analysis of the domain space or part thereof to the data analysis block. At this level, sub-decisions are made to orient the search and analysis process. Once the evaluation step is validated, i.e., sub-decisions are made, the results are assembled to constitute feasible solutions to the problem. In order to generate feasible solutions, it is possible to loop back to the previous evaluation phase, or even to the data analysis step. At some point, one or more feasible solutions are presented to a second evaluation step which will select one solution (or none) out of the pool of feasible solutions. After this selection procedure, a veto step is added, since it is possible for one or more of the collaborating agents to veto the solution selected (like in management-by-exception). If it is vetoed, the output of the veto step is empty, and the decision-making process starts over again. If the selected solution is not vetoed, it is considered the "final solution" and is transferred to the action mechanism for implementation.

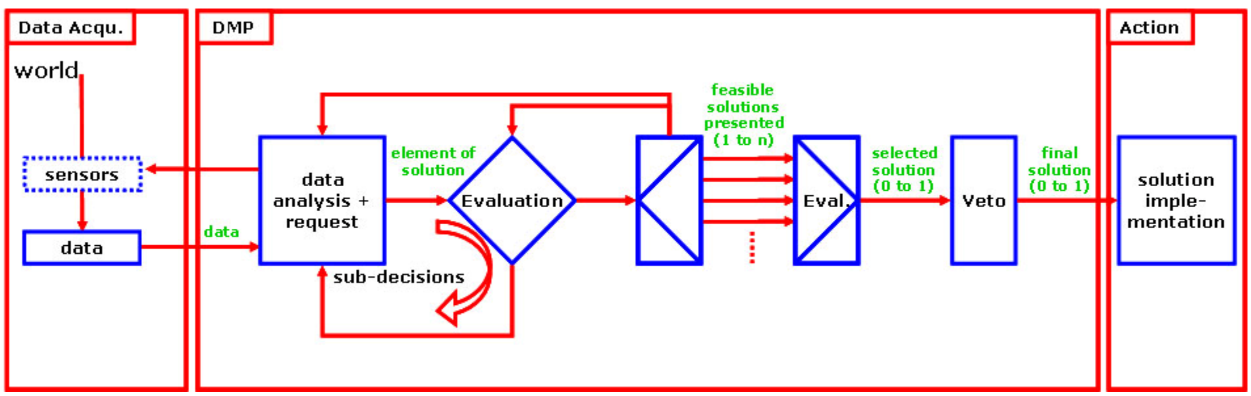

Figure 2. The collaborative information-processing model.

B. Three basic roles

Within the decision-making process of Figure 2, three key roles have been identified: Moderator, Generator, and Decider. In the context of collaborative human-computer decision making, these three roles are fulfilled either by the human operator, by automation, or by a combination of both. Figure 3 displays how the three basic roles fit in the collaborative information-processing model, and the individual roles will be detailed in the next section. The Generator and the Decider roles involve parts of the model that are mutually exclusive: the domain of competency of the Generator (represented by the blue square to the left of Figure 3) does not overlap with that of the Decider (the green square to the right). However, the Moderator's role (represented by the red, dashed arrows in Figure 3) covers the whole decision-making process. Each role has its own scale, which lists the range of possible human-computer role allocations.

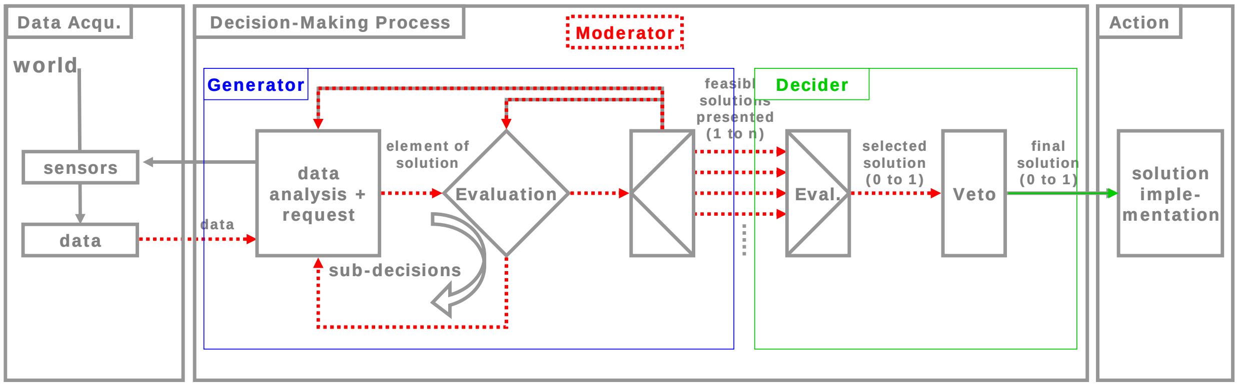

Figure 3. Visualization of the three collaborative decision-making process roles: Moderator, Generator and Decider.

The Moderator

Table 1. Moderator scale.

Moderator Level	Who assumes the role of Moderator?
1	Human
2	Mixed, but more human
3	Hybrid
4	Mixed, but more automation
5	Automation

The Moderator is the agent(s) that keeps the decision-making process moving forward (represented by the red, dashed arrows in Figure 3). The Moderator makes sure that the process goes from one block to another, and that the various phases are executed during collaboration. For instance, the Moderator may initiate the decision-making process and interaction between the human and automation. The Moderator may prompt or suggest that sub-decisions need to be made, or evaluations need to be considered. It could also be involved keeping the decision processing in pace when time deadlines must be met. The need for defining this role relates directly to ten-level SV LOA scale, where the difference between LOA 4 and 5 is who initiates generation of a solution (Parasuraman, Sheridan et al. 2000). However, we recognize that this moderation occurs in multiple portions of the decision making process and separate from the task of generating solutions and selecting them. The Moderator role has a five-point scale, displayed in Table 1, which describes the various types of role allocation possible. Scale levels 2 and 4 correspond to a "Mixed Moderator," meaning that both the human and automation are involved in keeping the decision-making process moving forward though the task is not equally shared. Scale level 3 is described as "hybrid", where both agents equally contribute to the Moderator role.

The Generator

Table 2. Generator scale.

Generator Level	Who assumes the role of Generator?
1	Human
2	Mixed, but more human
3	Hybrid
4	Mixed, but more automation
5	Automation

The Generator is the agent(s) that generates feasible solutions from the data. Typically, the Generator role involves searching, identifying, and creating solution(s) or parts thereof. The scale for this role is listed in Table 2, and is similar to that of the Moderator. Most of the previously discussed LOAs (e.g, (Endsley and Kaber 1999; Parasuraman, Sheridan et al. 2000)) address the role of a solution generator. However, instead of focusing on the actual solution (e.g., automation generating one or many solutions), we expand the notion of Generator to include other aspects of solution generation, i.e., all the other steps within the Generator box (Figure 3) such as the automation analyzing data to make the solution generation easier for the human operator. Additionally, it is acknowledged that the role allocation for Generator may not be equally shared between the human operator and the automation, thus, a five-point scale is used. An example of a scale level 3 Generator could involve a system where the human defines multiple constraints and the automation searches for a set of possible solutions bounded by these constraints. On the other hand, a Generator at a scale level 4 would be one where the automation proposes a set of possible solutions and then the human operator narrows down these solutions.

The Decider

The third essential role within our proposed collaborative decision-making process is the Decider. The Decider is the agent(s) that "makes the final decision", i.e. selects the potentially final solution out of the set of feasible solutions output by the Generator, and who has veto power over this selection decision. Veto power is a non-negotiable attribute: once an agent vetoes a decision, the other agent cannot supersede it. The concept of solution selection is discussed in Endsley and Kaber's LOA scale (Endsley and Kaber 1999). The veto power is also an important attribute that is described only in the Parasuraman et. al's (2000) LOA scale (upper levels). These aspects are embedded in existing LOAs but they are mixed and incomplete. We have improved upon previous frameworks by adding more resolution to the possible role allocations, listed in Table 3.

Table 3. Decider scale.

Decider Level	Who assumes the role of Decider?
1	Human makes final decision, automation cannot veto
2	Human or automation can make final decision, human can veto, automation cannot veto
3	Human or automation can make final decision, human can veto, automation can veto
4	Human or automation can make final decision, human cannot veto, automation can veto
5	Automation makes final decision, human cannot veto

The three roles, Moderator, Generator and Decider, focus on the tasks or actions that are undertaken by the human operator, the automation, or the combination of both within the collaborative decision-making process. They constitute the three main pillars of our Human-Automation Collaboration Taxonomy (HACT). These pillars form the basic building blocks of HACT, however, collaboration can be further characterized beyond basic role allocation. Other attributes pertaining to automation, such as information management and inter-agent communication, are crucial in further describing human-automation collaboration and will be discussed.

C. Three primary characteristics

We define three additional attributes aimed specifically at characterizing the automation and the system: Functional Transparency, Information Transparency, and Interactivity. The first two characteristics relate to the transparency of the collaboration, at the functional and informational levels. The third characteristic describes inter-agent communication through the system. These attributes describe intrinsic properties of the system, hence they are considered orthogonal to the basic roles of Moderator, Generator and Decider.

Functional Transparency

Functional Transparency denotes how the system provides feedback to the human operator about the way the automation works. In particular, this characteristic pertains to features that were purposefully included by the system designer to convey this information. As such, Functional Transparency applies to all levels of the Moderator, Generator and Decider scales that involve automation (which correspond to scale levels 2 and above). The Functional Transparency characteristic assumes that the automated system performs some sort of processing beyond basic input capture and display. We define three levels for Functional Transparency, as described in Table 4.

Table 4. Functional Transparency scale.

Functional Transparency Level	Description
Black	The system is opaque ("black box"): the human operator has no means to understand how automation works.
Grey	The system has features allowing the human operator to obtain a partial representation of the automation's internal process. The system may provide limited abstract information or use metaphors (e.g., trees or other high-level structures) to represent the process.
White	The system has features allowing the human operator to obtain a complete representation of the automation's internal process ("white box"). This is mostly a theoretical level, as the system's complexity may prevent complete automation representation.

Information Transparency

The Information Transparency attribute refers to the type of information that is presented to and used by the agents during the collaborative decision-making process. This characteristic is defined at three levels (Table 5), which represent the levels of abstraction of the information. The Information Transparency characteristic only applies to the Decider and Generator roles as they are the two roles dealing with information exchange between the agents.

Table 5. Information Transparency scale.

Information Transparency Level	Description
Raw	The agents collaborate using unprocessed low-level information, such as instantaneous sensor readings or measurements.
Mixed	The agents collaborate using both raw and aggregate information.
Aggregate	The agents collaborate using processed data, such as consolidated sensor measurements into abstract structures like a trend graph or post-imaging processing.

Interactivity

Interactivity defines how agents communicate with one another. It is characterized by a bimodal scale (Table 6) and is only applicable to the Moderator and Generator roles at the levels where both agents are involved (i.e. mixed or hybrids, scale levels 2, 3 and 4 in Table 1 and Table 2).

Table 6. Interactivity scale.

Interactivity Level	Description
Command	Agents assign orders to the other. The recipient may provide confirmation and/or feedback regarding the outcome of the command. At this level, agents unilaterally solicit actions
Dialogue	Both agents are engaged in a back and forth discussion, while working together towards reaching a goal.

Table 7 summarizes where each of the primary characteristics are applicable within each of the basic roles. Each line contains a different role while the columns contain the primary characteristics. Each cell represents a different level (1 through 5) for a different role. A shaded cell means that a primary characteristic does not apply to that particular level. For instance, a Moderator at level 2 or higher can have a certain level of Functional Transparency, whereas Information Transparency is not applicable to this basic role.

Table 7. Basic roles vs. primary characteristics

	Functional Transparency					Information Transparency					Interactivity				
Moderator	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*
Generator	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*
Decider	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*

D. A secondary meta-characteristic

HACT describes human-automation collaboration based on three basic roles (Moderator, Generator, and Decider), which are further characterized by a set of primary attributes (Functional Transparency, Information Transparency, and Interactivity). A secondary attribute completes the taxonomy by defining the Adaptability of each of the basic roles and of their primary attributes. Both systems with adaptive autonomy (where the level of automation changes without human intervention) and systems with adjustable autonomy (where the level of automation is changed by the human) are examples where the type of collaboration between the human and the automation changes. In order to take this attribute into account in HACT, Adaptability of a system is defined as: (a) the default scale level of the role or of the primary attribute, (b) all other possible levels that the system can adopt, and (c) whether the automation, the human operator, or both can trigger a change between these levels. By applying the attribute of Adaptability to each of described roles and primary characteristics, we are able to specify what part of the decision-making process has Adaptability. For example, a command and control system can have Adaptability in its Decider role such that the system is by default at level 1 (human makes final decision, automation cannot veto) but can be changed by the human operator to level 2 (human or automation can make final decision, human can veto, automation cannot veto). In this case, the human operator gives the possibility to the automation to make the final decision (decision offload), but retains veto power over it.

E. Visualization

In order to visually understand the relationships of the basic roles and attributes previously described, an innovative visualization technique for HACT was developed. If a decision support system is described using HACT, its corresponding HACT visualization would illustrate how much collaboration is present in the system. An implementation of the HACT visualization is demonstrated in this section.

The visualization begins with three nodes that represent the three basic roles of human-computer collaborative decision-making: Moderator (represented with an “M”), Generator (“G”), and Decider (“D”). These are placed in the center of a series of concentric circles (Figure 4). Each collaboration circle corresponds to an increase in collaboration. A sub-tree extends from each node; this sub-tree depends on the type of collaboration present in the decision support system for each HACT role. Types of collaboration relate directly to the basic role’s scale level and the levels of each of the pertaining primary characteristics. The longer the sub-tree, the more collaborative the basic role within the decision support system described. The more branches the sub-tree has, the more Adaptability the system has. Figure 4 depicts all the possible types of collaboration that can be described with HACT, i.e., the complete collaboration space. This current representation shows that some types of collaboration are more or less collaborative than other types. A simple linear relationship between levels was employed. For instance, within Interactivity, “Dialogue” is deemed more collaborative than “Command”. Thus, “Dialogue” is given a weight of one, while “Command” is zero. This simple heuristic represents ordinal rankings as “objective” weightings between levels have not been determined either experimentally or operationally. In order to illustrate one decision support system with the HACT visualization, only the possible branches that exist in the system are depicted within the corresponding sub-tree. If the system has Adaptability, the other branches that are possible are also drawn (see examples in Figure 7).

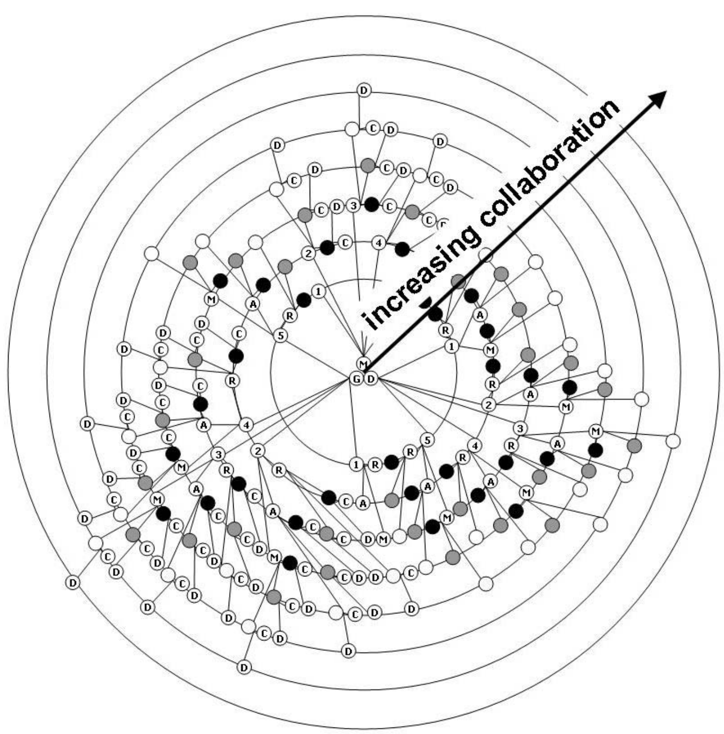

Figure 4. HACT visualization (complete collaboration space).

IV. Application to a Command and Control Decision-Support System

A. StrikeView: a prototype decision support system for Tomahawk mission planning

StrikeView was developed as a collaborative decision support system to support a Tactical Tomahawk missile strike coordinator in the task of mission planning. The human operator is in charge of matching a set of military missions aimed at destroying specific targets, with a set of available missiles aboard availabl launch platforms. Several system interfaces were created, at various levels of automation, to investigate how human operators collaborate with automation in mission planning tasks (Bruni and Cummings 2005; Bruni and Cummings 2006). Two possible interfaces of StrikeView are shown in Figure 5 ("Interface 1") and Figure 6 ("Interface 2").

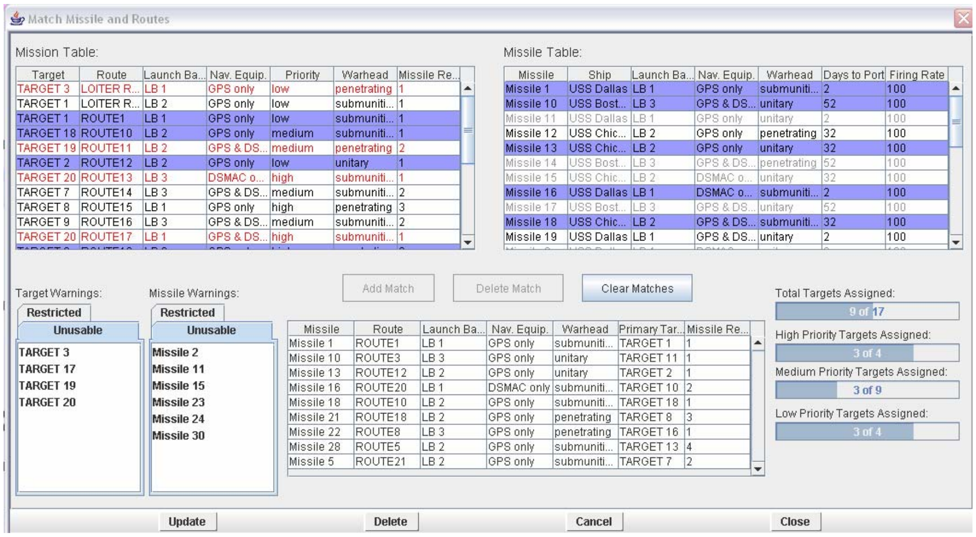

Figure 5. StrikeView Interface 1.

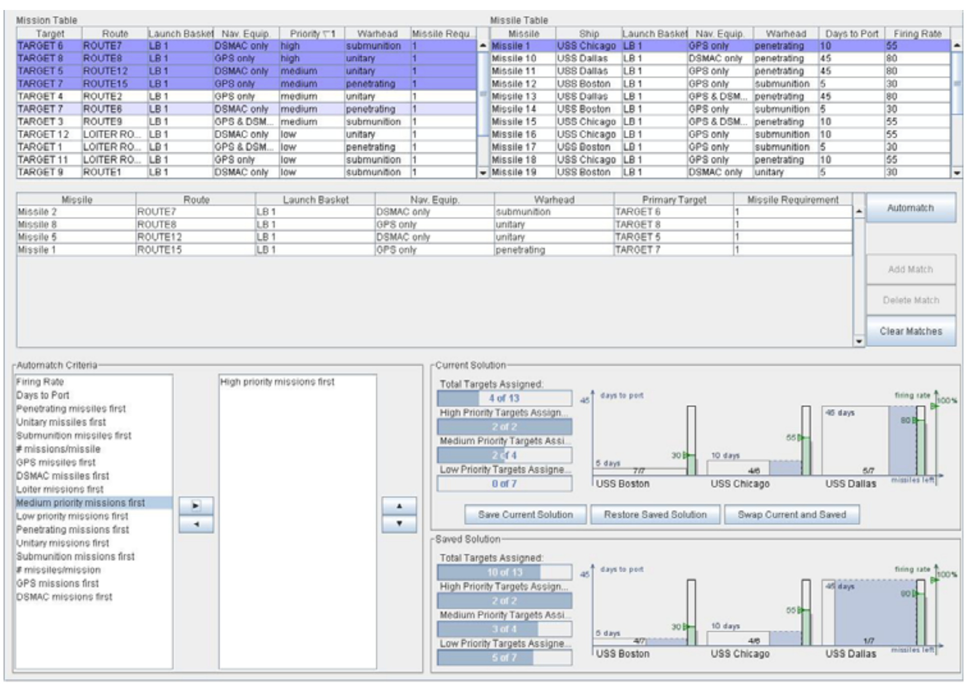

Figure 6. StrikeView Interface 2.

Interface 1 was designed to support manual matching of the operational assets (the missiles) to the operational goals (the missions). This interface provides raw data tables with all the characteristics of both missions and missiles that must be matched. However, Interface 1 only provides very limited automated support, such as basic data sorting, mission/missile assignment summaries by categories, and colored feedback on mission/missile incompatibility and current assignment status. Therefore, this interface mostly involves manual problem-solving.

Interface 2 was designed to offer the human operator the choice to either solve the mission/missile assignment task manually as in Interface 1 (note on Figure 6 that the top part of Interface 2 is a replica of Interface 1 shown in Figure 5), or to leverage the automation and collaborate with the computer to generate solutions. In the latter instance, the human operator could steer the search of solution in the domain space, by selecting and prioritizing search criteria, and giving them to the computer. Then, the automation's fast computing capabilities can be used to perform a heuristic search based on the criteria defined by the human. The operator can then either keep the solution output or modify it directly and manually, or modify the search criteria to get a new solution.

B. Applying HACT to StrikeView Interface 1

Table 8 summarizes the HACT categorization for StrikeView Interface 1. StrikeView is at level 1 for Moderator because the human operator fully controls the process. The automation does not intervene at that level. Since Interface 1 only features basic automation support, the generation of solutions is considered manual, hence the Generator role is at level 2. With respect to the Decider, StrikeView is at level 1 since only the human operator can validate a solution for further implementation, with no possible veto the automation. The functional transparency of this interface is categorized as "Black". The information used in this interface is only "Raw" data (Information Transparency), and the Interactivity between the human operator and the little automation present in the interface is defined at the "Command" level. This interface is not adaptable: the human cannot change how much automation is involved (for the three basic roles), nor change the type of information or interaction with the automation.

Table 8. HACT applied to StrikeView Interface 1.

	Moderator (M)	Generator (G)	Decider (D)
Scale level of role	1	2	1
Primary characteristics			
Functional Transparency	n/a	black	black
Information Transparency	n/a	raw	raw
Interactivity	n/a	command	n/a
Secondary characteristic			
Adaptability	no	no	no

C. Applying HACT to StrikeView Interface 2

We similarly applied HACT to StrikeView Interface 2 (Table 9). The Moderator remains at level 1 because the human operator is still in full control of the process, including which tasks are completed, at what pace, and in which order. The automation does not intervene. This system corresponds to levels 2, 3, and 4 for Generator because both the human and the automation participate in the generation of solutions, with more or less implication of each agent. Solutions can be generated manually by the human without intervention of the computer, but since the human can ask the computer to search and build solutions based on specific criteria, the computer can be involved in the generation of solutions. With respect to the Decider, StrikeView is at level 1 since only the human operator can validate a solution for further implementation. There is no veto possible from the automation.

Table 9. HACT applied to StrikeView Interface 2.

	Moderator (M)	Generator (G)	Decider (D)
Scale level of role	1	2 - 3 - 4	1
Primary characteristics			
Functional Transparency	n/a	grey	black
Information Transparency	n/a	mixed	mixed
Interactivity	n/a	dialogue	n/a
Secondary characteristic			
Adaptability	no	yes	no

The primary characteristics of each basic role are also presented in Table 9. There is no Functional Transparency attribute for the Moderator role, because it is at level 1, where the automation is not involved in moderation. The Functional Transparency of the Generator is categorized as "Grey" because the possibility for the human operator to select and order the automated search criteria gives insight as to how the automation conducts the search. However, the Decider is described as a “Black” box, as this StrikeView interface does not provide a specific means to the operator to understand how the automation decides on a solution.

There is no Information Transparency attribute for the Moderator role (per construction of the taxonomy). Within Generator and Decider, both raw (e.g., names and characteristics of the missions and missiles) and aggregated (e.g., group search criteria) information is available to the human and the automation. Hence, the Information Transparency attribute chosen is “mixed” for both Generator and Decider.

There is no Interactivity attribute for the Moderator role, because at level 1 the automation is not involved in process moderation, nor for Decider (per construction of the taxonomy). For the Generator role, the Interactivity between the human operator and the computer is labeled as “dialogue”. When human operators decide to use the automation to solve the problem, they first provide the computer with search criteria, and the automation returns a possible solution to the problem, which the operator can decide to refine, either by modifying the search criteria or by modifying the solution directly and manually.

In StrikeView Interface 2, the system is adaptable for the Generator role: the human can decide to do manual matching, or to combine it with automated matching, or to let the computer perform an automated search. This corresponds to a default Generator level of 2, with a possible human-initiated switch to levels 3 or 4 respectively.

D. Visualization

Figure 7 displays the HACT visualizations for StrikeView Interfaces 1 (left) and 2 (right), using the simple heuristic of types of collaboration as described in section III E. A comparison of both pictures quickly demonstrates that Interface 2 (right) is a much more collaborative and adaptable system than Interface 1 (left). Indeed, both Interface 2’s Generator and Decider sub-trees extend further out compared to those of Interface 1. In addition, Interface 2’s Generator sub-tree features several different branches, which accounts for that interface’s adaptability factor.

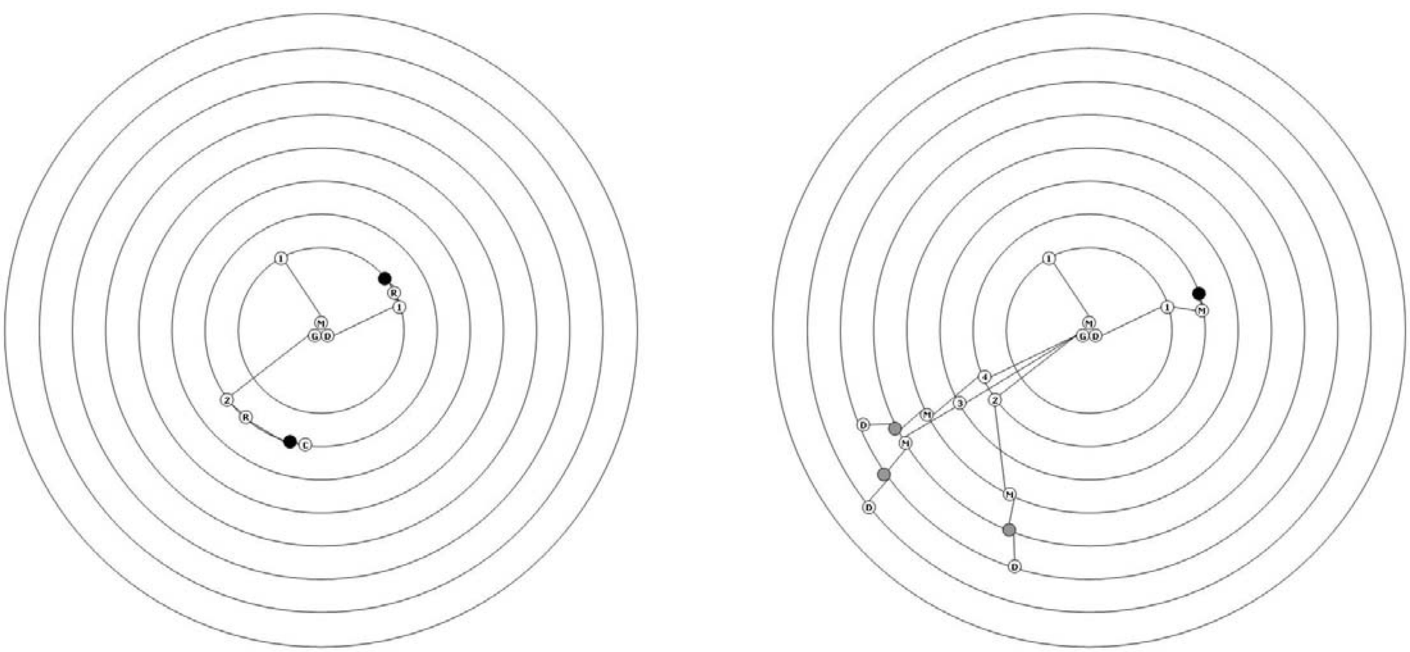

Figure 7. HACT visualizations for StrikeView Interfaces 1 (left) and 2 (right).

V. Discussion

The Human-Automation Collaboration Taxonomy was designed to address the gaps pertaining to previous frameworks. One main issue with previous human-automation frameworks was that their scales were too coarse in that they did not explicitly separate the solution generation from the solution selection phases. This distinction can be important in systems where, for example, the computer may generate automatically a single solution that the user can veto. In this case, the Generator is fully automated, and the Decider is fully human. In contrast, one could envision a system where the human leverages automation in order to jointly generate a solution. Such a system would be at a mixed level on the Generator scale, and, assuming the user chooses the final solution, the Decider would remain completely human. In the SV LOA scale (Sheridan and Verplank 1978), both systems would be at a level four (i.e. "the automation suggests one alternative"). The two systems are however very different with respect to the human involvement in the solution generation. This lack of granularity was addressed in our taxonomy by decoupling the roles of humans and automation in the decision process into our three basic roles: Moderator, Generator, and Decider.

The Endsley and Kaber scale (1999) addresses this issue in a similar way: the first system in the previous example would then be described at a "rigid system" level, whereas the second one would be at the "decision support" level. The Endsley and Kaber scale, however, lacks in granularity as it does not qualify the mode of interaction between the human and the automated agent. Taking the same previous example, there is a subtle, but important difference between the human operator ordering the automation to find an optimal solution as opposed to the automation suggesting how to improve a human generated solution. The latter implies a more sophisticated system, which is more collaborative. In both cases the system would be rated at the "decision support" level in the Endsley and Kaber scale, but would score differently in HACT on the Interactivity and Moderator attributes. The first example would be at the "Command" level for Interactivity and level 1 for the Moderator; the second example would be at the "Dialogue" level for Interactivity and level 3 for the Moderator.

A third issue was the lack of definition of the type of information exchanged between the human and the automation. Similar to Riley's approach, we solve this issue by defining an Information Transparency scale. The simplest of decision support systems would provide "Raw" level of information, which may require a lot of effort on the part of the operator to process and assess the important variables. A more collaborative system would aid the operator in processing this information, presenting it at an "Aggregate" level. Similarly, a collaborative system would ideally be transparent to the operator. The concept of Functional Transparency was non-existent in previous frameworks, but it is crucial in assessing the relationship between the human and the automation. Indeed, a lack of understanding of the underlying processes performed by the computer could lead to an erroneous mental model, and be detrimental to the trust in the automation if the system does not behave as expected (Billings 1997; Parasuraman and Riley 1997). Finally, the Adaptability of the system is also a criterion often ignored by other frameworks, but it is an important one since systems with adjustable autonomy are critical in workload management and system adaptation, especially for supervisory control systems. We address this issue in our framework by providing a systematic way to note the different levels or configurations in which the system may operate, the union of which essentially embodies the overall adaptability of the system.

While the HACT framework describes human-automation collaboration in decision support systems, the HACT visualization can be used to analyze the collaboration of existing systems, assist in designing a collaborative system for a specific application, and compare the collaboration between different applications. Various types of existing collaboration systems can be represented through the HACT framework and visualization. The collaboration for each system can be then correlated to the system's decision-making ability and quality. Such an analysis would provide insight as to what combination of collaborative elements is most common and most successful for decision support, particularly in command and control. In turn, these results can be used to suggest improvements for designing human-automation collaboration for future decision support systems. Future work using the HACT framework and its visualization will focus on determining if certain collaboration designs tend to facilitate good decision-making.

Within a specific application, the HACT visualization can be used to graphically examine trade-offs among different decision support system designs. Cost functions (with metrics such as money, time, performance, safety, etc.) can be applied to HACT-defined individual collaboration components, so as to facilitate a cost-benefit analysis. This would provide a comprehensive description of the benefits and drawbacks of different collaboration designs, as well as the costs involved in implementing those designs. By using HACT visualizations to compare different design trade-offs, a design could be chosen that best satisfies the requirements and constraints of the decision-support system. Future work will also exploit the HACT visualization as an interactive tool, allowing decision-support system designers to change proposed designs and see the resultant affects on the human-automation collaboration and associated design costs.

While HACT and its visualization have been presented within the context of command and control decision-support systems, the framework and visualization are universal and broad enough to allow for graphical comparisons of decision support systems from different applications. Because HACT divides collaboration into specific parameters with defined levels, human-automation collaboration between different systems can be graphically compared using the HACT visualization. This may enable researchers to determine why certain types of collaboration facilitate better decision making than others. It is possible that some collaboration designs, regardless of application, can hinder the decision-making process, either because they promote too much or not enough collaboration, or because role or characteristic levels within the collaboration conflict with one another.

VI. Conclusion

The Human-Automation Collaboration Taxonomy, or HACT, and its visualization are concepts and tools aimed at providing a common description and analysis framework for the specific domain of human-computer collaborative decision-making. By specifying a clear set of parameters and attributes of collaboration, HACT manages to overcome some of the setbacks of previous scales and frameworks. This framework thus allows for comparisons of human-automation collaboration of past, present, and future decision-support systems. The HACT visualization is potentially a powerful aid that may allow for a better understanding of the benefits and drawbacks of human-automation collaboration. HACT takes the abstract idea of human-automation collaboration and turns it into a physical representation that can be analyzed, compared, and contrasted.

Acknowledgments

This research effort is sponsored by the Office of Naval Research, the Air Force Office of Scientific Research, the NASA Jenkins Predoctoral Fellowship Program, Charles River Analytics, Boeing, and the American Association of University Women Fellowship. The views and conclusions presented in this paper are those of the authors and do not represent an official opinion, expressed or implied, of any of the agencies listed above. Special thanks to Prof. Mary L. Cummings from Massachusetts Institute of Technology, USA, Prof. Gilles Coppin from the Ecole Nationale Supérieure des Télécommunications de Bretagne, France, and Prof. Duncan Campbell from Queensland University of Technology, Australia, for their contributions, supervision, and careful insight.

References

Billings, C. E. (1997). Aviation Automation: The Search for a Human-Centered Approach. Mahwah, NJ, Lawrence Erlbaum Associates.

Bruni, S. and M. Cummings (2005). Human Interaction with Mission Planning Search Algorithms. Human Systems Integration Symposium, Arlington, VA.

Bruni, S. and M. L. Cummings (2006). Tracking Resource Allocation Cognitive Strategies for Strike Planning. COGIS 2006 - Cognitive Systems with Interactive Sensors, Paris.

Dillenbourg, P., M. Baker, et al. (1995). The Evolution of Research on Collaborative Learning. Learning in humans and machines. Towards an interdisciplinary learning science. P. Reimann and H. Spada. London, Pergamon**:** 189-211.

Endsley, M. (1987). The application of human factors to the development of expert systems for advanced cockpits. Human Factors Society 31st Annual Meeting, Santa Monica, CA.

Endsley, M. R. and D. B. Kaber (1999). "Level of automation effects on performance, situation awareness and workload in a dynamic control task." Ergonomics **42**(3): 462-492.

Goodrich, M. A., J. Johansen, et al. (2007). Managing Autonomy in Robot Teams: Observations from Four Experiments. Human Robot Interaction. Washington D.C.

Parasuraman, R. and V. Riley (1997). "Humans and Automation: Use, Misuse, Disuse, Abuse." Human Factors **39**(2): 230-253.

Parasuraman, R., T. B. Sheridan, et al. (2000). "A Model for Types and Levels of Human Interaction with Automation." IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans **30**(3): 286-297.

Riley, V. (1989). A General Model of Mixed-Initiative Human-Machine Systems. Human Factors Society 33rd Annual Meeting, Denver, CO.

Sheridan, T. B. and W. Verplank (1978). Human and Computer Control of Undersea Teleoperators. Cambridge, MA, Man-Machine Systems Laboratory, Department of Mechanical Engineering, MIT.

Wickens, C. D., A. Mavor, et al. (1998). The future of air traffic control: Human operators and automation. Washington DC, National Academy Press.