

Crew Resource Management for Automated Teammates (CRM-A)

Robert J. Shively^{1*}, Joel Lachter¹, Robert Koteskey² and Summer L. Brandt²

¹ NASA Ames Research Center, Moffett Field CA 94040, USA

² San Jose State University, Moffett Field CA 94040, USA

{robert.j.shively, joel.lachter, robert.w.koteskey,
summer.l.brandt}@nasa.gov

Abstract. Crew Resource Management (CRM) is the application of human factors knowledge and skills to ensure that teams make effective use of all resources. This includes ensuring that pilots bring in opinions of other teammates and utilize their unique capabilities. CRM was originally developed 40 years ago in response to a number of airline accidents in which the crew was found to be at fault. The goal was to improve teamwork among airline cockpit crews. The notion of “team” was later expanded to include cabin crew and ground resources. CRM has also been adopted by other industries, most notably medicine. Automation research now finds itself faced with similar issues to those faced by aviation 40 years ago: how to create a more robust system by making full use of both the automation and its human operators. With advances in machine intelligence, processing speed and cheap and plentiful memory, automation has advanced to the point that it can and should be treated as a teammate to fully take advantage of its capabilities and contributions to the system. This area of research is known as Human-Autonomy Teaming (HAT). Research on HAT has identified reusable patterns that can be applied in a wide range of applications. These patterns include features such as bi-directional communication and working agreements. This paper will explore the synergies between CRM and HAT. We believe that HAT research has much to learn from CRM and that there are benefits to expanding CRM to cover automation.

Keywords: Crew Resource Management (CRM), Human-Autonomy Teaming (HAT), Automation.

1 Why Human-Autonomy Teaming (HAT)?

For centuries, automation has been making humanity more productive. Historically this automation has chiefly replaced the brawn of human labor, leaving people as the brains. Recently, however, machines have been used by humans for increasingly cognitive tasks. Computers can play chess better than humans. They can search the web to answer your questions. They can drive cars. But automation is not perfect. It is brittle, breaking, often catastrophically, when taken out of the “comfort zone” it was

designed for. It has crashed stock exchanges [1] and cars [2]. Self-driving cars still need to “phone home” when encountering a person directing traffic [3]. Aviation has not been immune to this increase in automation, from relatively simple autopilots to sophisticated flight management systems. These innovations have not always been smooth [4]; however, the development of automation as tools to improve flight safety and efficiency has continued.

As the nature of automation changes, so should the role of the human when interacting with automation. Automation is moving from the realm of simple “tools” into intelligent cognitive agents that can function as teammates, similar to human teammates. Innovations in artificial intelligence as well as increases in the speed and memory of the underlying hardware have spurred this shift. Cognitive agents can now suggest courses of action, monitor the operator actions and physiology, and monitor the environment. With all of these new capabilities, work has begun to investigate how the relationship of the human and the automation can be shifted toward that of teammates to best take advantage of this phenomenon. This idea of Human-Autonomy Teaming (HAT) goes beyond simply giving a person a computer. The idea is that humans and automation should work together the way (well-functioning) human teams do, bouncing ideas off of one another, backing each other up when they sense potential problems, and keeping each other informed of what they are doing. HAT is currently recognized as a promising solution to the problems of human operators managing increasingly complex work systems. A human-autonomy team has been defined as “the dynamic, interdependent coupling between one or more human operators and one or more automated systems requiring collaboration and coordination to achieve successful task completion;” [5] a definition that has been picked up by others [6, 7]. As such, it is being developed and pursued in many operational areas such as robotics [8], commercial aviation [9], and UAS operations [10].

Aviation, with its very systematic approach to safety, may prove to be both a source of inspiration for developing better human-autonomy teaming, and an industry where its benefits might be most profitably exploited. While computers have eliminated the need for the navigator and (together with an increase in reliability) the flight engineer, regulatory barriers have slowed the ascent of automation in the cockpit. Years of cultivating teamwork on the flight deck have contributed to aviation’s superb safety record. Here we explore potential synergies between the lessons the airline industry has learned from years of studying teamwork and this newer field of human-autonomy teaming. Can we make human-autonomy teaming more effective by looking at the aviation model? Can we introduce advanced automation to the flight deck more safely if we introduce it as a team member? In the cockpit, where the importance of teaming has long been understood, the skills necessary for good teaming and the training procedures for teaching those skills have been codified under the umbrella of Crew Resource Management (CRM). Here we explore the synergies between HAT and CRM that may allow these new computerized teammates to be incorporated into CRM and may improve HAT concepts by incorporating the lessons learned from the development of CRM.

2 What is Crew Resource Management (CRM)?

2.1 History of CRM in Human Teams

What can we learn from CRM as it has been developed for human teams? Although the discipline has expanded to other high-risk high-reliability industries, the multi-crew airline flight deck is where CRM has its roots.

Risk management has always been the core task for an airline pilot. In the early days of transport aviation, risks due to mechanical failure were more prominent and aircraft were less reliable. The majority of threats were clear, and external to the human who was tasked with operating the machine. Early airline captains were solo performers whose technical skills were sharpened by absolute necessity. Their selection, their environment, and their culture reinforced strong, independent personalities and isolated, individual decision making. As large aircraft became more complex, a single pilot could no longer operate the aircraft. The crew compliment grew but was comprised of people who were still focused on individual tasks. As the job quickly became a more team-oriented endeavor, flight deck culture and nature of the individual pilot had not necessarily changed to reflect this shift.

By the late 1960's and 1970's transport accidents due to mechanical failure had drastically decreased. Advances in aviation technology like jet engines, modern avionics, and increased automation, so lowered the accident rate that the majority of new occurrences were now found to be crew related. Thus, further safety improvement could most easily be found from within the human team [11]. Landmark accidents like UAL 173, caused by fuel exhaustion, and EAL 401, brought down by the distraction from a faulty landing gear indicator, highlighted the need for training on team leadership and decision making for multi-crew aircraft [12]. The rugged, isolated individual was perhaps not the ideal model for what was now clearly a team activity.

A breakthrough study, conducted by Ruffel Smith [13] suggested a correlation between the leadership and communication style of the captain with overall crew performance. Further research reinforced this hypothesis [14, 15], and during the 1980's, industry and government came together to form what was at the time, a new discipline of aviation team training called Cockpit Resource Management, or CRM. As it evolved and became recognized as applicable to the larger aviation community, it was renamed Crew Resource Management. CRM has been defined as “using all available resources-information, equipment, and people-to achieve safe and efficient flight operations” [16, pg. 20].

In the 1990's, CRM training programs were introduced at major airlines and in military aviation units. Many of the concepts originally trained were lifted from business school management training templates and were not particularly well received by pilots. The topics in these classroom seminars typically included assessments of personality and leadership style. While these assessments might have been good predictors of performance, they did not necessarily allow for actual behavioral change. The goal of the training was an attitudinal shift, but this was hard to measure and equally difficult to reliably move to the flight deck.

Eventually, over the course of two decades, emphasis has shifted to identifying observable flight deck team behaviors (good and bad) that could be trained and evaluated in actual line operations. Task analyses were conducted at major carriers and research organizations. Both technical and CRM skills were identified and then used to create training in high fidelity, full flight, line-oriented simulations. These Line Oriented Flight Training (LOFT) events are now a primary means of introducing, reinforcing, and evaluating CRM skills [17]. CRM training today still includes an indoctrination seminar for new hires, but now CRM skills are embedded in all aspects of multi-crew transport training. They are trained and evaluated alongside technical skills at every training opportunity. Current CRM doctrine uses the concept of Threat and Error Management (TEM). This paradigm seeks to acknowledge the universal existence of human error and outside threats. TEM seeks to engage the team in actively searching for those threats and inevitable errors, then through use of CRM skills, to avoid, minimize or correct them.

The goal of CRM is the optimization of the human team. Whether the task is performing a checklist, monitoring a standard operating procedure, leading the crew, or mindfully following the direction of the leader, the whole point is for the members of the team to form, and perform optimally, and in synergy. We will now look at recognized CRM concepts and tools that are trained to, and used by, human teams in aviation, with the intent of adapting concepts from this field to the optimization of the new human-automation team.

2.2 Generally Recognized CRM Training Concepts for Human Teams

Threat and Error Management. As noted above, Threat and Error Management is one recognized framework used to convey a common mental strategy for identifying potential threats to safe operation, for identifying potential errors, and for correcting them when they inevitably happen [17]. TEM defines a continuum of safety ranging from safe operations on one side to an undesired equipment state on the other (i.e. a series of errors has occurred and gone unrecognized and unmitigated so that an unsafe condition or accident has occurred). The goal of TEM is to always remain in, or return to, safe operations while avoiding an undesired equipment state.

Avoid. Teams are taught that all operations have certain inherent threats. A threat can be internal or external to the team. These are identified and called out during briefing with the intent that this affords a greater likelihood that these dangers will be consciously mitigated by all members of the team. These known threats are to be avoided and a strategy for doing so can be articulated even before the team begins its task. Common threats might include a new or inexperienced team member (internal), or a hazardous weather condition (external).

Repair. An important new implication of this model is that the team expects that it will make errors during the course of its action. In aviation training this was an innovation. Previous generations of aviation professionals were expected to perform procedures and tasks to near perfection. Individuals were likely to minimize or discount errors both in training and actual practice because there was a professional and social penalty for doing so. In the TEM paradigm, individual errors are expected,

so team members may be more likely to call out their own errors and those of others with less stigma attached. While individuals are expected to make errors, the team as a whole is expected to respond to those and correct them before they lead to an undesired equipment state.

Recover. Should the first two levels of defense fail to capture the inevitable errors and threats that a team is likely to encounter, an undesired equipment state may occur. The task then becomes recovering the operation to a safe state. It is expected that CRM skills be utilized to return to a safe operation.

Verbalize Verify Monitor. One CRM tool used to capture and mitigate error is known as Verbalize Verify Monitor or VVM. Using VVM, the team member planning on taking an action first verbally states the intention of the action. This allows both the acting member and a monitoring team member to verify that the action matches the intent of current team goals. If that test is met, the action is carried out and the results can then be effectively monitored.

Verbalize. A statement that focuses the team on the accuracy of the next action to be taken. This need not be a simple statement. Often, a question is appropriate. Especially if the initiator is checking the accuracy of their own perception. As an example: “I intend to turn right.” or “Is the next turn to the right?” may both be appropriate, depending on the verification goal of the person taking action.

Verify. When an intent is stated, another team member is needed to verify the accuracy of the statement. The verification will usually also be stated but may be passed by non-verbal means if appropriate to the workload and situation. The important aspect is that another member is now engaged and their input is received and acknowledged.

Monitor. This last step requires continued engagement of the team to ensure that the intent just stated and verified is actually carried out accurately. In the case of automation management, this is particularly useful during multi-step changes or complicated procedures. A simple aviation example might involve the pilot flying calling out an intended automation mode change, executing the change after ensuring the monitoring pilot is engaged, and then both pilots monitoring the results of the change for accurate results. This process is used not only in automation management, but in all aspects of team action. For instance, a flight deck team taxiing on an unfamiliar airport would use VVM to ensure the aircraft does not deviate from its cleared route (i.e. enter an undesired equipment state). Using VVM, the pilot steering would proactively call out each turn and holding point well before the action must be taken. This communicates intent, focuses the monitoring team member, and allows for checking of common situation awareness. VVM works to verify technical as well as non-technical tasks. It may even be used during briefing to verify and align team goals, or as a tool to engage low-response team members.

Standard Operating Procedures. A common observation made concerning airline crews is that it seems remarkable that persons who have never met each other could come together in a very short time, form a team, and then operate a complex piece of equipment in a challenging environment. This seemingly remarkable ability lies in the

training of those people to the same set of standard procedures. They have a reasonable expectation of what skill level they will encounter, and indeed, after the advent of CRM, even an expectation of how they will be treated as the team forms and acts.

Well known human factors tools like checklists, standardized callouts, and stabilized approach criteria all fall into the category of Standard Operating Procedures (SOPs). Some of these tools have been around as long as aircraft have existed and are used in many other high-risk high-reliability activities. SOPs ensure that a known, safe, efficient set of actions is used to navigate through complex procedures that require great accuracy. At an airline, they are developed by expert teams and vetted through internal and regulatory processes to verify their worth and efficacy.

SOPs are intended to be the familiar landmarks that provide reinforcement, guidance, and reassurance during the progression of the project that is the flight of a transport aircraft. Until the advent of CRM, there was no SOP for the operation of the human team. Even with an extensive set of well-defined procedures for operation of the equipment, it is possible for a group of humans to fail in the execution of those procedures if they cannot work together. CRM, then, can be thought of as set of SOPs for operating in a team with other humans. When used by high-functioning teams employing a full pallet of CRM skills, technical SOPs are landmarks that appear as familiar signposts that are the culmination and verification of various team actions. Conversely, to low-functioning teams without a good grasp of solid CRM skills, these landmarks may come as a surprise, presenting last-chance safety backstops, rather than the mile markers of a safe operation.

Systems Approach to Training CRM Skills (AQP and ISD). It is important to consider how the different needs and existing skills of a particular organization influence the CRM training product. How does an organization identify the particular CRM skills and values it would like to emphasize, and how does the organization evaluate whether or not the program is effective? For airline operations, this process is codified by the FAA in two Advisory Circulars. One, FAA Advisory Circular 120-51E, concerns CRM [18], the other FAA Advisory Circular 120-54A, concerns the Advanced Qualification Program or (AQP) [19]. The CRM publication outlines basic CRM training topics, and the AQP publication describes a systematic process for training and evaluating CRM skills alongside technical skills.

For airline operations using an AQP, the requirements attempt to ensure that CRM skills and the resulting training program are tailored to that airline's needs. A systematic approach is advocated called Instructional Systems Design (ISD). This model starts with needs assessment to see where the organizational baseline is. A training goal is then defined and objectives written to support that goal. Once training is written and delivered, data is gathered to facilitate a process of continual improvement.

2.3 Common Target Concepts for CRM Training

In the airline industry, major airlines generally still present a CRM indoctrination seminar that includes some aviation human factors and an introduction to CRM general concepts. This establishes that common vocabulary which is used for training and

evaluation of the concepts and tools throughout a pilot's career at a given carrier. This involves learning to operate the "human equipment" one is teamed with at the same level of proficiency as the mechanical equipment in use. Maximizing the team's efficiency and output is seen as being of equal importance to technical proficiency. How is this done quantitatively and what do those behaviors look like? An excellent summary of recognized CRM skills that might be generally trained to new hires and evaluated in experienced operators is found in the book Crew Resource Management [20]. Below, we briefly lay out the CRM skills they identified. We will then discuss the synergy these CRM skills have with HAT and the potential benefits of expanding CRM to cover automation.

Communication. Making sure that there is bottom up communication as well as top down is the core of all other CRM competencies. CRM training teaches techniques for clear communication. One example is the use of active listening, which is the mindful repetition of the sender's message back to them as a check for understanding. Another is teaching common, simple communication models to team members. An example of a model that might be trained during a CRM seminar is a three step process in which communication has not occurred until 1. A message is transmitted, 2. The message is received, and 3. Feedback is provided. Proactive communication utilized mindfully by all team member is essential for all other aspects of CRM to work.

In a high functioning human team, a person who has important information makes sure that information is communicated. This is true whether that team member is a decision maker or a subordinate. Alternate modes and channels are attempted until it is clear that the communication has occurred. As human teams form, they become more efficient in their communication. They begin to learn how to communicate with the individuals they are teamed with and then adapt their style to fit the person and the situation.

Briefing. Dedicated briefings are useful in organizing teams and maintaining a common plan so that actions are properly choreographed. A team leader may use the preflight briefing to set the tone of team interaction in addition to the simple passing of pertinent information. During flight, a pause to re-brief as conditions and goals change is also useful in quickly redirecting or re-focusing the team. Post flight debriefing is important in the continuing process of improvement for individuals, for teams that are likely to reform, and for the organization as a whole.

Backup Behavior. While briefings and SOPs set out roles and responsibilities for each crew member, in well-functioning teams, task allocation is not absolute. Particular circumstances may result in one crew member's assigned duties exceeding their workload capacity. Under such circumstances, workload should be shifted so that it is balanced across team members. This kind of transition can be seen during an off-nominal flight deck event requiring a shift from nominal SOP duties to alternate duties that compensates for the increased task complexity and better distributes the changed workload. For instance, the captain may assign the first officer both pilot

flying duty and the task of communicating with ATC, a shifting of roles as the flying pilot would normally rely on the pilot monitoring for this task. Because of the off-nominal (perhaps requiring someone to run a checklist and work on a change of destination), workload must be re-distributed.

Mutual Performance. “To err is human.” CRM attempts to prevent human errors from resulting in incorrect actions through mutual monitoring among crew members. Crew members are taught to give and take advice in an open and non-judgmental manner. An important aspect of this is the VVM technique discussed in 2.2 above. Monitoring is also an important aspect of Backup Behavior. Crew members are taught not only to watch for individual lapses in judgment, but also their partner’s overall workload and mental state, and to offer greater assistance when a partner becomes overloaded.

Team Leadership. While many of the CRM skills have the effect of making teams more egalitarian, there is still a recognized need for one person to be in charge. The desired goal of the leader is to ensure that all team members are used optimally, and that they are engaged both with the task and with the team. Good leaders organize the team in a way that makes appropriate use of each team member’s skill and ability and keeps them working together in a positive manner. For example, a good captain might have someone with more flight experience as a first officer, or might have a rookie. What is appropriate monitoring and mentoring in one case might be perceived as condescending and micromanaging in the other. It is the leader’s task to make this type of assessment and modify their interaction to obtain maximum results from the team.

Decision Making. A key goal of CRM is to improve decision making by encouraging the consideration of multiple possible courses of action and assessing each using as complete a collection of information as possible. To do this, crew members are taught to bring up potentially relevant information and alternative actions, considering the possible consequence of each, with an attitude of “what is right, not who is right.”

Task-Related Assertiveness. In order to maintain a collaborative decision making atmosphere it is important that team members develop an appropriate level of assertiveness. They must be able to communicate information and suggestions with persuasive logic while maintaining an ability to listen and be persuaded by other team members. This helps ensure that all available information is put on the table in a transparent manner when decisions are being made. In fact, it goes beyond this, to require team members to bring information forward even when it is not explicitly requested. Appropriately assertive input, particularly from subordinates, is key to sound communication. It is a cornerstone of good followership, the important but often neglected obverse to good leadership.

Team Adaptability. Humans show a strong natural tendency to maintain a particular course of action, even when, from a purely logical standpoint, it no longer makes

sense to do so (e.g, the sunk cost fallacy [21]). CRM attempts to counter this tendency, by encouraging continuous re-evaluation of the current course of action, recognition of possible threats to the current goals, and discussion of options.

Shared Situation Awareness. Many of the skills discussed in this section can be viewed through the lense of developing an accurate shared awareness of the situation. Communication and briefings serve to keep crew members on the same page, with shared goals and a shared understanding of the environment so that they do not act at cross purposes. CRM skills related to monitoring, assertiveness and adaptability serve to maintain the accuracy of this shared understanding.

3 CRM for Human-Autonomy Teaming

Automation research now finds itself faced with similar issues to those faced by aviation 40 years ago: how to create a more robust system by making full use of both the automation and its human operators. This section will examine two overarching HAT concepts that have been proposed and then looks at how the CRM skills identified above might be applied to a human-automation team.

3.1 HAT Concepts Supporting CRM-like Behavior

Bi-directional Communication. As with the VVM pattern in human teams, it is important to develop a style of communication that makes sure that information is communicated across all team members in human-autonomy teams. This pattern has been referred to as bi-directional communication [22]. For automation to participate as a teammate, it is critical to have a bi-directional communications channel. This will allow humans to team effectively with automation and allows the human (and automation) to question, share hypotheses, provide additional input, etc. just as human teammates would. This bi-directional communication is critical and enables many of the CRM elements to follow. It must be bi-directional to allow the pilot to input information into the system that the automation might not have access to via sensors or databases. For example, when deciding on an alternate airport due to a medical emergency, the pilot might know more about the medical facilities in and around particular airports than does the automation. For the automation to participate fully as a partner, it needs to share this information, therefore this channel needs to exist. Similarly, the automation needs to be able to alert, share hypotheses, level of confidence, etc. with the human teammate. This then, allows the human to better judge the value and understanding of the automation and trust appropriately. This is an example of transparency, being used to calibrate trust, enabled by bi-directional communications. Teammates often discuss options, brainstorm on solutions and openly discuss courses of action. For automation to be on a team, this bi-directional communication needs to exist. Bi-directional communication is key to solving a number of the issues typically found in highly automated systems. Bi-directional communication can make systems more transparent and less brittle and further can facilitate intent based interface design.

Working Agreements. In human-human teams, SOPs provide both a level of predictability for how team members will react in a variety of situations, along with the ability to plan for many situations offline. Work in HAT has developed a concept similar to SOPs for use with automation [23, 24]: working agreements. Working agreements encapsulate goals, procedures, and division of responsibility into a package that can be specified offline and instantiated quickly in real-time situations. Working agreements specify who (automation or human, and, in the case of humans, which human) is responsible for performing various acts associated with a particular situation. This responsibility can be conditional. For example, the automation might be given autonomous authority to follow a route unless potential hazards are detected, at which point it might alert the human operator for verification that it should proceed.

3.2 Developing CRM Skills with Autonomy

As discussed in section 2.2 above, human teams are taught certain CRM skills that provide standardized mechanisms for using the situation specific SOPs to improve performance on certain measures thought to improve operational outcomes. Here we discuss how automation designed to work with bi-directional communication and working agreements can mirror those same skills.

Communication. CRM in the context of human teams emphasizes the need for communication to flow both up and down the chain of command. Our bi-directional communication pattern is designed to enable something similar between humans and automation. Part of the reason for CRM was that superiors do not always want to hear what their subordinates are saying and subordinates are often scared to speak-up. Analogous problems may occur between human and automation. While we can assume automation will not fear speaking up, how can we make sure that the human listens appropriately? Some research indicates that manner in which information is communicated influences the degree to which operators accept and rely on automation [e.g., 25]; however more work is needed in this area.

Further, just as humans must adapt their communication styles to their teammates, it may be appropriate to build similar adaptability into the automated team members. Just as humans must be trained to operate synergistically within their teams, perhaps automation should also be capable enough to recognize individual human style in order to maximize HAT performance. How will the automation provide a metaphorical “touch on the shoulder” when its human partner is deemed not to be listening?

Briefing. As with human-human teams, human-autonomy teams must share a common plan to assure that actions are properly choreographed. This goes beyond the flight plan to include alternate airports (depending on flight progress), weather, aircraft status, any potential areas of concern, and any issues that would normally be discussed with the crew. Digital representation of the flight plan exists and so is straightforward to transfer to the automation, however, other aspects may be harder to transmit. It may be necessary to build a “briefing interface” on top of the bidirectional communication channel to allow the crew to easily and fully provide this information

to the automation. Such an interface would be a logical place to define working agreements between the crew and the automation.

Here may be an ideal opportunity to build the human-automation team in much the same way as human-human teams are built. A briefing for a human team provides not only information transfer, but, importantly, affords the team an opportunity for each member to adjust team dynamics and to begin forming communication strategies.

It is also important to consider that briefings do not only occur at the beginning of a task. High functioning crews use them to re-focus the team when the situation has changed and at the end of the mission to review lessons learned, both good and bad. These other modes of briefing may have utility in helping humans and automated team members align goals and adapt to each other.

In addition, human also use briefings to assess one another and to modify their behavior to better conform to each teammate's preferred style of communication or operation. Could the briefing opportunity for the automation to gather information on the human team member's preferences? Perhaps the human team member might proactively push this preference information to the automation as part of the briefing. One could perhaps envision an opportunity for carrying this information in a profile of some kind that could follow a human team member from station to station.

Backup Behavior. SOPs help human team members anticipate the needs of others by providing clarity about each other's responsibilities. CRM training teaches pilots formal ways to modify these responsibilities by shifting workload between members to create balance during periods of high workload or pressure while maintaining a clear understanding of who is responsible for what. Similar flexibility can be built into working agreements between humans and automation. As noted above, a key feature of working agreements is the ability to specify the conditions under which each party is responsible for taking certain actions. Stress and workload levels can be among these conditions. However, doing so requires that these levels can be sensed or conveyed to the automation without adding to the overall workload. Sensing raises its own issues; how does the operator know that the automation has sensed high workload and changed the task allocation? An alert or annunciation system would have to be very sophisticated to avoid distracting the operator in a high workload situation. Operators could initiate such changes vocally or with a simple interface such as a button or dial. This would allow the operator to control the timing of any task reallocation and assure that his or her situation awareness is maintained. Again, this is an opportunity for further research.

Mutual Performance Monitoring. An important aspect of CRM is crew members monitoring the performance of other crew members. This is also true for automation. The automation needs to be able to monitor the crew; are tasks being performed in a timely manner, is the pilot planning and staying "ahead of the aircraft"?

As in the CRM construct of TEM, errors should be expected, recognized, and mitigated as necessary. Humans find a variety of ways to begin monitoring each other's actions. Indeed, they start to monitor the quality of team interaction as well. In high-functioning human teams, members quickly learn each other's style. Even on a

short mission, with formerly unfamiliar teammates, this information about how the other operates may provide enhanced monitoring ability. HAT monitoring of error and compliance may benefit from perception of, and adjustment to, these individual human characteristics and preferences.

An extreme case of required monitoring would be for nefarious behaviors. In human teams, much of this kind of screening is accomplished as the team forms and briefs. Under HAT, we would expect that under most circumstances, the human would lead the team. However, if the pilot has significantly deviated from the flight plan without a plausible explanation, the automation may have the authority to take certain actions; contact air traffic control, company dispatchers or in extreme cases, take control of the aircraft. In addition, it may be necessary for the automation to monitor pilot physiological state; heart rate, blood pressure, eye gaze, etc. to verify that the pilot is fit and operating at acceptable levels of physiological states. This could become critical in the case of incapacitated pilots; the automation may have the authority to take control in certain situations. For either of these types of monitoring to occur, the automation needs insight into the actions on the cockpit and physiological sensors on the pilot/crew.

Monitoring needs to be bi-directional in nature. That is, the pilot/crew must be able to monitor the automation just as the automation monitors the pilot. To do so, the pilot needs insight into the automation. Transparency into the processing and decision making provided by bi-directional communication is critical to this monitoring. When the automation alerts the pilot or offers suggestions (e.g., alternate airports), the logic of the processing needs to be available for examination by the pilot. But, the need goes even further, the pilot needs to have indications that the automation is monitoring and performing as intended. The bi-directional communication interface should be designed to provide this information.

Team Leadership. For the foreseeable future, the pilot (human) will be the pilot in command and therefore the leader of this team, and thus ultimately responsible for the performance of the team. However, automation can still fully participate as a teammate, as long as several of the attributes discussed thus far are in place; e.g., communications, monitoring, transparency. These will provide the leader with the mechanism and the information required to direct and coordinate the activities of team members, encourage team members to work together; assess performance; assign tasks; develop team knowledge, skills, and abilities; motivate; plan and organize; and establish a positive team atmosphere.

The coordination with the automation may be through working agreements. As with human teams, it is incumbent upon the leader to know the abilities of individual team members, and understand how to communicate effectively with them. In HAT, this may look like extensive training for the human team leader on the automated team member's capabilities and limitations. In addition, it may be desirable for the automated team member to practice the analog of good followership as described for human teams. This might involve the automation retaining some capability to recall communication or interface preferences of the human team leader.

Decision Making. Good decision making in human teams involves gathering and integrating information, identifying alternatives, and considering the consequences of each alternative. CRM encourages options developed by one team member to be evaluated and refined by other team members. This suggests that for automation to have good CRM skills it should be able to do three things that current automation typically cannot: evaluate options proffered by a human operator, give reasoning behind options it proffers, and compare options.

Evaluation of options proffered by the operator might be facilitated with a “course of action scratch pad,” that would allow the operator to input a proposed course of action (e.g., commands, reconfiguration, or routing) and have the automation evaluate it, presenting predicted outcomes (e.g., risk assessment, estimated fuel usage, ETA, etc). Similarly, when proffering a course of action, automation should be able to give similar evaluations along with an indication of what options were considered in developing this course of action.

Task-Related Assertiveness. A related measure of good CRM in human teams is the ability of members to communicate their ideas, opinions, and persuasively while remaining open to being convinced by the facts that other options are better. Working agreements can be used to implement such task related assertiveness in an automated system. A working agreement can specify conditions under which automation should “speak up” by alerting human operators to problems with the current course of action and/or offering alternatives. In addition, several channels for gaining the human’s attention may be desired. In human teams, a physical touch, a change in verbal tone or cadence, or even specific standardized phrases are all methods currently in use to gain the attention of a crew member who is not attending to a particular message. Some of these may translate well to automation while others clearly will not.

Team Adaptability. In human teams, good CRM requires that the team be able to alter a course of action or adjust strategies when new information becomes available, rather than push forward with a suboptimal or even infeasible plan. Properly designed automation can help with this. While people are often reluctant to give up on a course of action once it is embarked upon, automation has no such limits as long as it is running open loop. Automation can detect when the current course has become suboptimal and propose deviations. To prevent the automation from overwhelming the human operator with new options (e.g., modifying an aircraft’s trajectory every time the wind shifts slightly), working agreements can be developed that limit proposals such proposals to cases where the risk or cost difference meets a certain threshold. Alternatively, a working agreement could be developed that gives the automation authority to make small deviations from the current course of action autonomously.

Shared Situation Awareness. In human teams, it is important to maintain a common understanding of the task and team environment to keep everyone working toward a common goal. When new information becomes available it must be communicated or team members may find themselves working at cross purposes (e.g., if one pilot hears

a controller say descend to FL270 and the other hears FL260). This issue is even more important in dealing with human-automation teams because the human and the automation do not innately have the same information available to them. The automation takes in data from various sensor feeds at a level of detail that, even if it is available to the human, the human cannot process. Similarly, the automation can make calculations much faster and more precisely than the human, allowing it to quickly recognize and react to changes in the environment. The automation may not, however have the full range of senses that the human has. This is particularly important when it comes to understanding other people who, in most cases, set the objectives of the system. The bi-directional communication channel discussed in 3.1 above will be very important in allowing both the human and the automation to integrate this information and assuring that these representations are compatible. In some cases, it may be appropriate to “re-brief” as mentioned above in order to check for common goals and understanding of the mission. One could envision that this action could be called by either the automated component (if it senses that the human is taking inconsistent action) or by the human member of the team.

4 Conclusion

CRM has become deeply integrated into airline crew training. As automation rises the level of a teammate, it is imperative that this new status be reflected in CRM curriculum. It is recommended that airlines review their CRM training and incorporate this new more powerful automation paradigm as a critical component. It is perhaps equally imperative that research into CRM be incorporated into the design of these new non-human team members. The human-automation team should be developed in such a way that the human team member may eventually trust and interact with the automated team member in many of the same ways as they would with another human.

References

1. Steiner, C.: Automate this: How algorithms took over our markets, our jobs, and the world. Penguin Group, New York (2012).
2. Stewart, J.: Why Tesla’s autopilot can’t see a stopped truck. Wired, <https://www.wired.com/story/tesla-autopilot-why-crash-radar/> (2018).
3. Lawler, R.: Nissan’s SAM uses humans as a backup for self-driving tech. Engadget, <https://www.engadget.com/2017/01/05/nissans-sam-uses-humans-as-a-backup-for-self-driving-tech/> (2017).
4. Wiener, E.L.: Cockpit automation. In: Wiener, E. L., Nagel, D. C. (eds.) Human Factors in Aviation, pp. 433-461. Academic Press, Inc., San Diego (1988).
5. Cuevas, H. M., Fiore, S.M., Caldwell, B.S., Strater, L.: Augmenting team cognition in human-automation teams performing in complex operational environments. Aviation, Space, and Environmental Medicine, 78, B63-B70. (2007).

6. Langan-Fox, J., Canty, J. M., Sankey, M. J.: Human-automation teams and adaptable control for future air traffic management. *International Journal of Industrial Ergonomics*, 39, 894–903 (2009).
7. Strybel, T. Z., Keeler, J., Mattoon, N., Alvarez, A., Brakezyan, V., Barraza, E., Park, J., Vu, K-P. L., Battiste, V.: Measuring the effectiveness of human autonomy teaming. In: Baldwin C. (eds.) *Advances in Neuroergonomics and Cognitive Engineering*. AHFE 2017. *Advances in Intelligent Systems and Computing*, vol 586, 23-33. Springer, Cham (2017).
8. Chen, J. Y. C., Barnes, M. J.: Human-agent teaming for multi-robot control: A literature review (ARL-TR-6328). Aberdeen Proving Grounds, MD: Human Research and Engineering Directorate (2013)
9. Brandt, S., Lachter, J. B., Russel, R., Shively, R. J.: A human-autonomy teaming approach for a flight-following task. In: Baldwin C. (eds.) *Advances in Neuroergonomics and Cognitive Engineering*. AHFE 2017. *Advances in Intelligent Systems and Computing*, vol 586, 12–22. Springer, Cham (2017).
10. Fern, L., Shively, R. J.: A comparison of varying levels of automation on the supervisory control of multiple UASs. In: *Proceedings of AUVSI's Unmanned Systems North America*. Washington, DC (2009).
11. Kanki, B. G., Helmreich, R. L., Anca, J.: *Crew resource management*. 2nd edn. Academic Press, San Diego (2010).
12. Helmreich, R. L., Foushee, H. C. : Why CRM? Empirical and theoretical bases of human factors training. In: Kanki, B. G., Helmreich, R. L., Anca, J. (eds.) *Crew Recourse Management*, 2nd, pp. 3–57. Academic Press, San Diego, pages (2010).
13. Ruffel Smith, H. P.: A simulator study of the interaction of pilot workload with errors, vigilance, and decisions. NASA-TM-78482. NASA Ames Research Center, Moffett Field (1979).
14. Chidester, T. R., Helmreich, R., Gregorich, S., Geis, C.: Pilot personality and crew coordination: Implications for training and selection. *International Journal of Aviation Psychology* 1, 23–42 (1991).
15. Chidester, T. R., Kanki, B. G., Foushee, H. C., Dickinson, C. L., Bowles, S. V.: *Personality factors in flight operations. Volume 1: Leader characteristics and crew performance in a full-mission air transport simulation*. Moffett Field, California: NASA Ames Research Center. (1990).
16. Lauber, J.: Resource management in the cockpit. *Air Line Pilot* 53, 20–23 (1984).
17. Federal Aviation Administration: Flightcrew member line operational simulations: Line-oriented flight training, special purpose operational training, line operational evaluation. Advisory Circular 120-35D (2015).
18. Federal Aviation Administration: Crew resource management training. Advisory Circular 120-51E (2004).
19. Federal Aviation Administration: Advanced qualification program. Advisory Circular 120-54A (2006).
20. Shuffler, M. L., Salas, E., Luiz, X. F.: The design, delivery and evaluation of crew resource management training. In: Kanki, B. G., Helmreich, R. L., Anca, J. (eds.) *Crew Recourse Management*, 2nd, pp. 205–232. Academic Press, San Diego, pages (2010).
21. Kahneman, D.: *Thinking, fast and slow*. Farrar, Straus and Giroux, New York (2011).
22. Shively, R. J., Lachter, J., Brandt, S. L., Strybel, T. Z.: Human-autonomy teaming in a flight following task. In: *NATO HFM 247 Technical Activity Description* (in press).
23. Miller, C. A., Parasuraman, R.: Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors*, 49, 57–75 (2007).

24. Gutzwiller, R. S., Espinosa, S. H., Kenny, C., Lange, D. S.: A design pattern for working agreements in human-autonomy teaming. In: International Conference on Applied Human Factors and Ergonomics, pp. 12–24. Springer, Cham (2017).
25. de Visser, E J., Krueger, F., McKnight, P., Scheid, S., Smith, M., Chalk, S., Parasuraman, R.: The world is not enough: Trust in cognitive agents. In: Proceedings of the Human Factors and Ergonomics Society 56th Annual Meeting, pp. 263-267 (2012).