Bi-directional Communication in Human-Autonomy Teaming

Robert J Shively
NASA-Ames Research Center
Moffett Field, CA
USA

robert.j.shively@nasa.gov

Gilles Coppin
Telecom Bretagne
Breast
FRANCE

gilles.coppin@imt-atlantique.fr

Joel Lachter
NASA-Ames Research Center
Moffett Field, CA
USA

joel.lachter@nasa.gov

ABSTRACT

Communication between autonomous agents and humans is critical for effective Human Automation Teaming (HAT). This paper investigates the Human Factors issues, technological challenges and implementation issues that are associated with bi-directional communication to facilitate comprehension. In HAT, as well as human teams, good communication is paramount to coordinate task planning and execution. It is important that all necessary information be communicated without excessive chatter that can divert attentional resources from task performance. Human autonomy teams face increased challenges in this novel environment such as the definition of a mechanism to communicate with the autonomy (keyboard input or natural language), and defining a shared language that both humans and automation can understand. A paradigm shift is required to move from using autonomous agents as tools to treating them as teammates. Because communication is central to effective teaming, it would be useful to standardize methods for operators to communicate with automation that can be reused with autonomous systems of varying types. Such methods might be analogous to the system of windows, menus, and files that became the standard metaphors for interaction with computers on the 1980s. In this paper we identify issues that bi-directional communication may address and issues that need to be addressed for effective communication between the operator and autonomous agents. Bi-directional communication can help with understanding the capability of the automation, self-confidence in particular solutions, provide transparency to help calibrate trust and through those mechanism even help with agility to reduce brittleness. As mentioned, an effective channel of communication must be established and the content of information must be defined to provide shared situation awareness. Consideration must be given to both what needs to be communicated (context, options, authority), and the manner in which it is best communicated (how and when). A "bi-directional communications manager" is proposed, that is capable of tracking goals, plan progress, and user mental state, to determine when communication is necessary. Such a communications manager could be

customized for use with a wide variety of autonomous systems. Systems with such a manager would be far more interactive than current automation, encouraging iterative planning in which both automation and human operators contribute as well as joint execution and monitoring, in much the same way that crew resource management proscribes procedures that assure both pilots monitor the situation and vet plans in current commercial aviation. As autonomous systems develop, it is clear that humans need ways to understand and team with them effectively. At the same time, current automation has proven brittle, requiring human intervention in off-nominal situations. More comprehensive communication between humans and automation can help solve both of these problems. We believe that a standardized communication system for HAT will increase usability and adoption of advanced automation.

1.0 HUMAN AUTONOMY TEAMING (HAT)

For hundreds of years automation has been taking on new tasks that increase exponentially in complexity. Recently these advances appear to be closing in on a new threshold, one where the automation takes on tasks that require it to act autonomously in the world. Autonomous cars promise to save us the drudgery of our daily commutes and autonomous aircraft promise to move us and our stuff around without traffic constraints. However, there are still hurdles to overcome. Such systems are brittle, working properly within some bounded space for which they have been programmed, then failing when parameters fall outside that space. They lack transparency; human operators often do not know what the automation is doing or why. Operators often do not know when to trust automation, relying on it to handle conditions it cannot, or not taking advantage of it to handle conditions it can. When the automation performs a task, the operator is often less aware of the system state, causing confusion if the operator needs to step in. Compounding this, as automation does more of the work, operators become less practiced and then are unable to perform correctly when the automation fails.While each of these issues is troubling by itself, they often manifest together. An operator, over-trusting the system, does not realize that some parameter has gone out of bounds. Because it is out of bounds, the automation either quits or is no longer reliable. The out of practice operator must then try to regain situation awareness using opaquely presented information. This situation has been responsible for a number of accidents (e.g., Air France 447 [1], Korean Air 801 [2]). Similar concerns have been identified by others [3-7]. A number of authors have suggested that these concerns are ameliorated by developing interfaces and procedures analogous to guidelines for improving teamwork among humans [4, 7, 8, 9]. This area of research has been termed “Human-Autonomy Teaming” or “HAT.” [7,9] HAT is a simple concept: treat automation as a teammate, rather than a tool. This will help people get the most out of the automation and work with it more naturally. However, as with many simple concepts, the implementation can be complex. The next section outlines some of the key attributes that will allow users to move toward treating automation as a teammate.

1.1 Key Attributes of HAT

1.1.1 Pilot directed interface

Previous work has focused on task allocation between the humans and automation. However, this static relationship can lead to non-optimal performance. At the same time, if the automation re-allocates tasks, the operator can become disoriented. A well planned and understood allocation strategy coupled with an operator directed dynamic allocation of tasks based on workload, time pressure and other important factors allows a much more agile, flexible system. One potential enabler of such a dynamic system of task allocation is the concept of a play [10]. A play encapsulates goals, tasks, and a task allocation strategy appropriate for a particular situation. Operators can call a play to quickly adapt to a new situation. A pilot directed interface also facilitates understanding by explicit discussion of goals (as opposed to intent inferencing), as well as confidence, and rationale.

1.1.2 Transparency

Lyons [11] argues that the system and its operators should be able to recognize what the human/automated teammate is doing and why. He defines automation transparency as the enabler of such recognition. Transparency is more than simply making the information available to the human operator. Often the calculations done by automation do not correspond directly to those a human would do when performing the same task. To be transparent, the automation must bridge this gap by presenting information in ways that fit the operator's mental model. Transparency helps correctly calibrate trust in the sysem, which is a critical component, see below.

1.1.3 Bi-directional communication

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. We see bi-directional communication as key to solving a number of the issues typically found in highly automated systems. Bi-directional communication can make systems less brittle. History has shown that humans, generally, are better able to adapt to unfamiliar situations, poorly structured problems, or situations with incomplete information. With HAT, the human can provide the missing information and context, apply judgment from experience in similar situations that would not be recognized as relevant by the automation (Christofferson and Woods [4] have shown how automation does not generalize well from one domain to another) and if necessary, override the decisions made by automation in these situations. Bi-directional communication can also help to build shared awareness. Alberts and colleagues [19] define shared awareness as a cognitive capability that builds on shared information. Shared awareness of an event can be developed through four mechanisms: (1) directly, (2) though independent sensors, (3) through information that is passed from one agent to another, and (4) through information that is shared and the fused results presented to two (or more) agents. The last two, sharing and fusing of information require a communication channel through which information can pass in either direction. And, while the same information could be gathered independently (mechanisms 1 and 2), to assure shared awareness, this information needs to be cross-validated or eventually the information will be out of sync. Several of the key attributes of HAT are enabled by bi-directional communication. It is of paramount importance, therefore to ensure that this communication is effectively supported.

2.0 THE STRUCTURE OF BI-DIRECTIONAL COMMUNICATION

2.1 Definition

When bi-directional communication is discussed in the context of HAT, it is meant to convey the kind of communication that occurs between human team members. In such communication, each participant provides new information that the other might not have, each can offer hypothesis and pose "what if" questions, and each can particpate in exchanges before a shared view is arrived at. This kind of two-way, natural bi-directional communication is the key to effective teaming between humans and in-turn, effective teaming between humans and autonomy. A few key components are required for this communication and these are discussed below.

2.2 Channel

Some mechanism must be available to host the communication, that is, some kind of two-way channel. Ideally, this would be as natural as possible for the human operators. In some cases the most natural interface might be a natural language. However, this is not required. For some purposes machine speech recognition and generation are not up to the task. Further, speech is not always the best way of communicating, even among humans, who often point, draw diagrams, or physically guide each other. Applying pressure to a throttle or drawing a path on a sketch pad may be more natural ways of communicating particular ideas or intent. The critical aspect is that the human and the automation share a channel that they use to communicate. They need to able to input information, ask questions and otherwise interact through this channel.

2.3 Language

Bi-directional communication requires some ability to have a common cognitive understanding and for both the human and automation to communicate their intentions [12]. To have true bi-directional communication, a shared language will be required. Because computers and humans do not share an internal "language of thought", translation will be necessary. Ideally, most of the translation will fall to the automation; however, it may be necessary for the human to learn some formalization. The "round trip" nature of bi-directional communication provides a useful check that such translations have been successful. Figure 1 provides an example where the operator gives an initial command using a mix of natural language and gestures which the automation then interprets**.** The resulting interpretation is then displayed for the operator to validate. A shared communication channel allowing both implicit and explicit communication between the human and agent, is required.

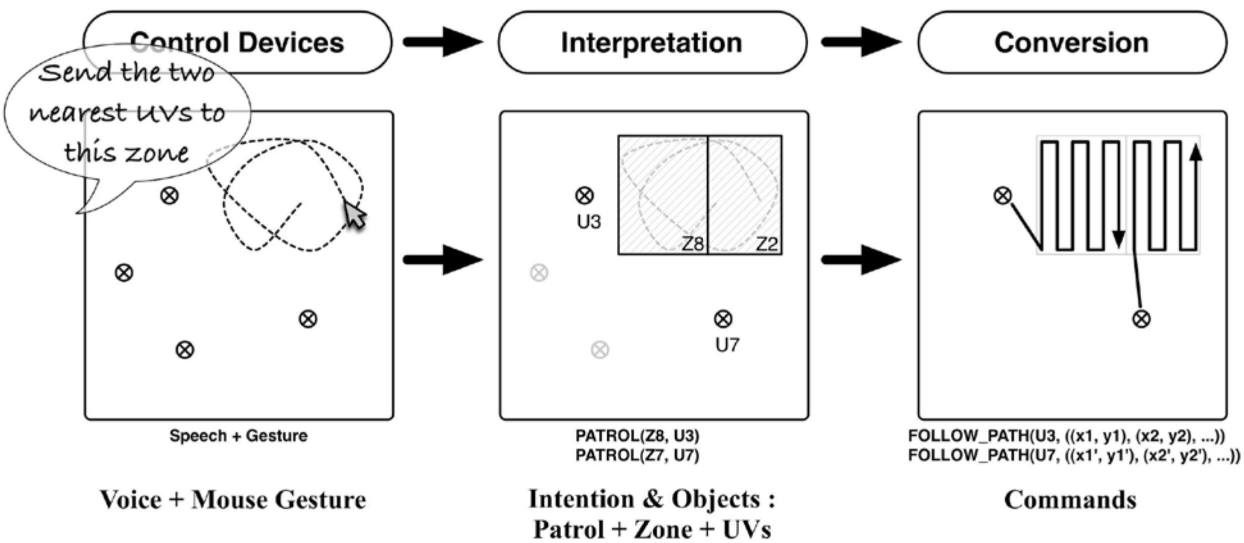

Figure 1: Communication Flow.

An example of such a language is the cockpit control language (CCL) developed by Riley [13]. CCL was based on the structured communications pilots use to talk to controllers and thus allowed pilots to easily translated controller instructions into commands for the flight management system. Such a codified, agreed upon language and syntax that would be understandable to humans as well as automation is critical to effective HAT.

2.4 Temporal Coordination

Most of the systems currently being discussed require real-time communication, or, at least, minimal delay. This can be critical to being able to effectively use teaming to make joint decisions. Systems operating in some environments such as under water or in deep space can have significant delays forcing the automation to largely operate independently.

2.5 Symmetry

The backbone of bi-directional communication is that both agents (human and computer) can input information, request information, and monitor the other agent; in short, the system needs to support symmetrical communication.

2.6 Examples

2.6.1 Everyday example: Navigation systems

We are probably all familiar with the navigation systems on our phone. They differ somewhat, but share many qualities. In general, these systems are extremely helpful and many of us don't know how we traveled without them. However, with HAT, they could be dramatically more useful, like a human navigator with knowledge of the area, traffic, weather, construction and all the factors that can affect your drive.

Currently, these systems will tell you that traffic is ahead and possibly the cause; an accident or a police action, for example. Then the system may (or may not) provide you with an alternate route (usually just one). That is about the extent of the (largely) one-way communication.

If you had a dialog with the system, you could ask several types of questions, bounce ideas or hypothesis off and interact more fully and naturally with the system. Perhaps you'd like to know where the accident is and ask how long taking an alternate route that you define would take. Or, you might have a prefered way to drive to work, and that doesn't show up as a recommendation, you can ask why not. This ability will allow you to treat this automation like you would treat a human teammate. Why are you saying that? What about this route?

The components of this bi-directional dialog are discussed below.

2.6.1.1 Channel

In this case, the channel may be verbal, a very natural channel for the human. This might be esspecially useful while driving or otherwise occipied. However, that requires automatic speech recognition and a natural language interface. While ideal for a human, the maturlity of this technology may limit its effectiveness. A keyboard and visual display interface may be serve the same function (althought not while driving!)

2.6.1.2 Language

The language may be normal spoken word, which again is optimal for humans. However, for the automation to parse and understand the semantics, it's useful to define syntax and even approved vocabulary.

2.6.1.3 Temporal coordination

This application would require real-time or near realtime porcessing and interface.

2.6.1.4 Symmetry

As indicated, this application would require the ability of both agents (human and automated) to be able to initiate, respond and build-on communications.

2.6.2 Aviation Example

Investigations as NASA recently looked at single pilot operations in commercial airlines. Questions included the workload of the single pilot and how automation may reduce that workload. In one instance, a flight was en route to San Francisco and a medical emergency arose which required a diversion to an alternate airport. The pilot was aided by an emergency landing planner system that when faced with a diversion would recommend an alternative airport. In this example, the system recommends diversion to Cheyenne, Wo. However, like many current systems, this is the end of the one-way communication.

This leaves the pilot without knowing why Cheyenne was chosen over Denver. Or how close other alternatives were - were alternatives just separated by a few numbers on a huge scale? What if the pilot knows that Denver has better medical facilities or has better company facitlities? How can he/she input that data to the system. Again, there is a need for a dialog between the agents in the system. The pilot should be able ask these questions, input data, ask what if questions and generally treat the automation like he/she would a teammate.

The components of this bi-directional dialog are discussed below.

2.6.2.1 Channel

In this case, the channel may be verbal, a very natural channel for the human. However, the auditory channel is quite busy when flying, so a keyboard/visual display may be more compatible or perhaps a graphic editing tool coupled with the keyboard. The pilot might circle an area, then type in “what airports are avialalbe in this area?” This type of hybrid interface can be very effective.

2.6.2.2 Language

Aviation brings with it its own grammar and syntax, for example when a pilot speaks to air traffic control. Building on that would be compatible with the humand and much eaiser for the automation to understand.

2.6.2.3 Temporal coordination

This application would require real-time or near real-time processing and interface.

2.6.2.4 Symmetry

As indicated, this application would require the ability of both agents (human and automated) to be able to initiate, respond and build-on communications.

3.0 DISCUSSION

In order to communicate, the system must be first able to compute all the necessary elements such as the level of confidence (of its own proposals of decision or action), a course of action that can legitimate its communication and related decisions, its "own" expression of goals and decision consequences. This goes through the ability of learning from the user (classification of current functional state and cognitive state of the user) and from the situation. All these abilities mostly rely on Artificial Intelligence and Learning.

New capacities of autonomous systems lead to a paradigm shift as the system has now the possibility of building its own representations of world (the situation), self (system self evaluation) and the partners (learning from the user), from which it can propose or take decisions. A first step before decision concerns the communication of these representations towards the user as a teammate would do, while traditional systems are only mirroring the unique user's knowledge and representations.

The existence of these knowledge representations opens the perspective of using many of the features of human dialogs (etiquette, emotion, all performative - or perlocutionary - linguistic acts) and to develop efficient communication in the framework of HAT.

For the locutionary acts, we can refer to an existing framework proposed proposed by authors such as [Hoc, Chauvin, Flemisch] for human- human cooperation. Especially, teaming considered as cooperation needs to rely on a jointly structured ground (common ground or COFOR for Hoc) addressing views on external situation, hints on internal situation (functional states of team members) and elements of meta-knowledge (such as respective roles of team members), and to communicate bidirectionally on these elements.

As systems become more and more autonomous, HAT will become more and more critical and bi-directional dialog is key to impllmenting effective HAT.

4.0 REFERENCES

[1] Langewiesche, W. The human factor. *Vanity Fair*. Retrieved from http://www.vanityfair.com/business/2014/10/air-france-flight-447-crash on September 7 2018 (2014)

[2] National Transportation Safety Board. *Controlled Flight Into Terrain, Korean Air Flight 801, Boeing 747-300, HL7468, Nimitz Hill, Guam, August 6, 1997.* Aircraft Accident Report NTSB/AAR-00/01. Washington, DC. (2000)

[3] Parasuraman, R., Riley, V.: Humans and Automation: Use, Misuse, Disuse, Abuse. Hum. Factors. 39, 230–253 (1997)

[4] Christoffersen, K., Woods, D.D.: How to Make Automated Systems Team Players. In: Advances in Human Performance and Cognitive Engineering Research, vol. 2, 1–12. Emerald Group Publishing Limited (2002)

[5] Lee, D.D.: Review of a Pivotal Human Factors Article: "Humans and Automation: Use, Misuse, Disuse, Abuse. Hum. Factors. 50, 404–410 (2008)

[6] Parasuraman, R., Manzey, D.H.: Complacency and Bias in Human Use of Automation: An Attentional

Integration. Hum. Factors. 52, 381–410 (2010)

[7] Endsley, M.R.: From Here to Autonomy: Lessons Learned from Human-Automation Research. Hum. Factors. 59, 5–27 (2017)

[8] Chen, J.Y.C., Barnes, M.J.: Human-Agent Teaming for Multi-Robot Control: A Literature Review. Army Research Lab Technical report, ARL-TR-6328 (2013)

[9] Shively, R. J., Lachter, J., Brandt, S. L., Matessa, M., Battiste, V., & Johnson, W. W. Why Human-Autonomy Teaming? To be presented at the 8th International Conference on Applied Human Factors and Ergonomics (AHFE 2017), Los Angeles, CA. (2017)

[10] Miller, C.A., Parasuraman, R.: Designing for flexible interaction between humans and automation: delegation interfaces for supervisory control. Hum. Factors 49, 57–75 (2007)

[11] Lyons, J.B.: Being transparent about transparency: a model for human robot interaction. In: AIAA Spring Symposium Series (2013)

[12] Chen, J.Y.C., Barnes, M.J., Harper-Sciarini, M.: Supervisory control of multiple robots: human performance issues and user interface design. IEEE Trans. Syst. Man Cybern.–Part C: Appl. Rev. 41, 435–454 (2011)

[13] Riley, V., DeMers, B., Misiak, C., & Schmalz, B. The cockpit control language: A pilot-centered avionics interface. *Air & Space Europe*, *1*(1), 48–50 (1999)