Self-Confidence of Autonomous Systems in a Military Environment

William G. Kennedy[1], Ciara Sibley[2], and Joseph Coyne[2]
[1]George Mason University, 4400 University Drive, Fairfax, VA, wkennedy@gmu.edu
[2]Naval Research Laboratory, 4555 Overlook Avenue SW, Washington, DC, first.last@nrl.navy.mil

Abstract

The topic of the self-confidence of autonomous systems is discussed from the perspective of its use in a military environment. The concepts of autonomy and self-confidence are quite different in a military environment from the civilian environment. The military's recruit indoctrination provided a basis for the concept, the factors affecting the concept, and its measurement and communication. These and other aspects of the topic self-confidence in autonomous systems are discussed along with examples based on current research on the interface between human operators and such systems.

Introduction

The succinct directions provided by a car's navigation system offer little insight into the system's underlying processes. This leaves the driver wondering why the automation makes some of its recommendations (e.g., does it have access to traffic information or historical data). In this situation the driver is directly controlling the vehicle and can choose to follow or ignore the navigation system's directions. However it becomes increasingly important to understand the system's decisions, and confidence in those decisions, when the system becomes more than just advisory and actually drives the car. In a normal day-to-day environment, an autonomous machine's use is our voluntary choice and the machine's self-confidence could be anywhere on such a scale and be tolerated. However, autonomous machines, their use, and their self-confidence in a military environment where lives could be affected is a very different story.

The use of some autonomy is a necessary in many situations such as those where there are significant communication delays before human direction can be provided (e.g., space missions). NASA began to use more autonomy in its Mars Rover Curiosity (for navigation) only recently, in 2013, due to the long communication delays (20 minutes on average) (NASA, 2013). The military's use of automation could be necessitated by requirements for immediate action, such as evading enemy fire; or requirements for electronic silence, such as covert data collection; or simply to reduce the workload for operators responding to all the chatter from multiple autonomous systems.

Command and control are vital to military operations, not just to maintain military discipline, but because lives could be on the line. Therefore, military expectations for command and control of autonomous systems are vital and are not necessarily the same as civilian expectations. In this paper we describe first what autonomy means in a military environment, then we discuss how the military environment affects the concept of self-confidence using the questions proposed by the symposium organizers. Finally, we discuss how such self-confidence applies to members of a human-autonomous system team.

Autonomy in a Military Environment

The current Department of Defense (DoD) roadmap for unmanned systems (DoD, 2013), points out that the terms "autonomy" and "autonomous systems", are often used when "automatic systems" would be more appropriate. As the Defense Science Board noted in its 2012 report, autonomy is a capability "that enables a particular action of a system to be automatic or, within boundaries, 'self-governing'". Furthermore, they contend that "all autonomous systems are *supervised* by human operators at some level". They discuss that human oversight should and will always be possible since autonomous machines are built by humans with software that "embodies the designed limits on the actions and decisions delegated to the computer" (Defense Science Board, 2012, p 1).

In this paper, we will use ***autonomous machines*** to refer to the unmanned aerial vehicles (UAVs) or similar hardware and software artifacts and ***autonomous systems*** to refer to the combination of autonomous machine(s) and at least one human operator. Further, using autonomous to describe systems can easily bring to mind the idea of fully independent, rogue, or killer robots from popular movies. In stark contrast, the Defense Science Board emphatically noted (Defense Science Board, 2012, p. 24), “there are no fully autonomous systems just as there are no fully autonomous soldiers, sailors, airmen or Marines.” The same year, the official DoD directive on the topic, DoD Directive 3000.09, was issued on the matter particularly related to weapons. The directive made it clear that weapon systems, either autonomous or semi-autonomous, must be designed and tested so that the commanders and operators can “exercise appropriate levels of human judgment over the use of force” (para 4.a). It also provides a definition for the concept of a human-supervised autonomous weapon system:

> An autonomous weapon system that is designed to provide human operators with the ability to intervene and terminate engagements, including in the event of a weapon system failure, before unacceptable levels of danger occur.

This directive is intended to make the U.S. policy clear to all Department of Defense components. It does not apply to cyberspace systems, unguided or operator guided munitions, mines, or “unexploded explosive ordnance”. It should be noted, however, that automatic weapons systems can be put in a fully automatic mode when defense of ships or installations requires response times for which humans couldn’t defend themselves. This is referred to as supervised autonomy by the directive.

Parasuraman and Riley (1997) define automation as the transference of actions formerly performed manually by a human to a machine, and not part of the original purpose of the machine, using examples like starters for cars and automatic elevators. They also discuss how there continues to be a “much greater collective emphasis on the technological than on the human aspects of automation” which results in designers automating everything that *could* be automated, rather than considering what *should* be automated for effective teamwork. They note that this can lead to unhealthy trust in the automation and inadequate supervision of the automation.

“Over automating” a system fundamentally changes the human’s engagement, undermining the fact that the human’s role is still to provide direction and monitor the performance of that system. Although through automation, humans have transferred some authority for the execution of a task, they cannot divest themselves of the ultimate responsibility for the results (Army R600.20, 2-1.b). This principle also applies to an autonomous system: someone is still personally responsible.

For the rest of this paper, autonomy and self-confidence of autonomous machines in a military environment will be discussed with respect to their immediate human operator (at the bottom of the human military hierarchy) who currently may operate a single UAV, but who the DoD envisions will one day control/supervise multiple UAVs. This future capability will require an increase in responsibility of both the operator and the vehicle’s automation. Successfully bringing about these new roles requires not only advances in automation, but also understanding the interaction of the human decision maker and the machine.

The Naval Research Laboratory (NRL) developed the Supervisory Control Operations User Testbed (SCOUT™) (Sibley, Coyne, Morrison 2015) as a tool to begin investigating some of the human-automation interaction issues involved with implementing this new supervisory control paradigm within military operations. SCOUT was developed to represent the primary tasks that a future operator would likely need to control, or manage, multiple UAVs (including prioritization of targets, route planning and airspace management, communication, and monitoring vehicle and payload status). It also captures a range of mission/task performance metrics, which are synchronized with detailed data on the user’s behavior and physiological data (including eye gaze, pupil and eye lid opening size, heart rate variability, and key stroke data). Some of the examples in the following sections will reference examples from tasks within SCOUT.

Figure 1. Participant interacting with the SCOUT™ simulation environment with integrated eye tracking system.

The Self-Confidence of Military Machines

We now turn to the questions posed for this symposium. Each will be discussed with respect to an autonomous machine in the military environment. We consider the autonomous machine as part of the military hierarchical chain of command and discuss self-confidence as it relates to the supervision of multiple autonomous machines.

1. What does "self-confidence" mean in the context of autonomous systems?

The military, through initial training or "boot camp", instills some basic concepts such as history, mottos, codes of conduct, values, warrior ethos, duty, honor, integrity, responsibility, etc. in all new members. Part of this training is to build a military basis and reference for self-confidence. The Army's Initial Entry Training Soldier's Handbook (TRADOC Pam 600-4) provides this background for all new recruits. It includes several examples of noble behavior and several quotes including this from Eleanor Roosevelt:

> You gain strength, courage, and confidence by every experience in which you really stop to look fear in the face. You must do the thing you think you cannot do.

Turning to the Navy, consider the meaning of naval expression "aye, aye". Its use conveys the self-confidence of the sailor as a warfighter. It acknowledges (1) that an order has been received, (2) that the order is understood, and (3) that it will be carried out immediately. This is different from answering with a "yes" or "yes, sir" because these could mean only agreement without any expectation, intention, or commitment to act.

Aye, aye goes beyond positively acknowledging receipt of the order. There is no "I didn't get the e-mail" excuse permitted. A response is required. Next, to be understood, the order's explicit tasks and the implicit assumptions entrained in the order must be compatible with the person's (or machine's) capabilities. A sailor would not be expected to understand an order to bring a fully rigged warship about, without adequate training, and a Roomba vacuum cleaner would not understand orders to vacuum the ceiling.

Finally, the response of "aye, aye" carries the understanding that the order is to be carried out immediately. To respond this way, the task must be within the person's/machine's capabilities and the environment must be within operating conditions and a plan developed and considered do-able, because the acknowledgment takes ownership of the responsibility to execute the order. Autonomous machines need to be able to make this commitment and monitor its execution. Requesting clarification of an order not meeting these conditions is permissible, even expected. Autonomous machines operating within a military environment would be expected to behave the same way as a new recruit.

Upon completion of boot camp, new members of the military have a new identity that forms the basis of their self-confidence, as demonstrated by this quote from the Navy Training Command's website (NTC, 2013):

> You have completed five-and-a-half weeks of intense training. You have scored well on three academic tests. You have passed personnel inspections in which your appearance, knowledge, and military bearing have been tested. Your compartment and gear have been judged as "ship-shape" during your bunk and locker inspections. You have passed the final Physical Fitness Assessment. You have demonstrated your ability to work as a team at firefighting and basic seamanship. Now at the end of the sixth week of training you will put your skills to the test. Completion of Battle Stations will earn you the respect of your peers. You will become a Sailor in the United States Navy. Your family will see you march at graduation. Your life is about to change, if you can just pass this final hurdle.

Within the context of supervisory control, i.e. an individual supervising multiple UAVs, the operator and autonomous vehicles have their own hierarchical system or team. The operator receives orders from his command and then tasks the UAVs to accomplish the necessary mission. Self-confidence within this hierarchical system exists on multiple levels, with the vehicles effectively acting as the operator's subordinates. The machine needs to behave as a team member in a military environment. The confidence an operator has that he can accomplish the mission and the confidence a vehicle has that it can accomplish its specific task can vary due to differing experiences, assumptions and access to information.

Example: What does "self-confidence" means to SCOUT controlled UAVs.

Within NRL's SCOUT environment an operator is given responsibility of three UAVs, equipped with forward looking infrared (FLIR) sensors, and directed to conduct an intelligence, surveillance, and reconnaissance (ISR) mission. Specifically, the operator is provided a list of time sensitive targets to find, which have varying priority levels, geographic locations, and search area sizes. When the operator creates a route plan and assigns a target to one of his/her 3 UAVs, that UAV implicitly acknowledges the order by reporting how much of the target search area it can cover prior to that target's deadline.

The UAV calculates this coverage area as a percentage using cruise speed, distance to target, and the range of its sensor (e.g. 90% of search area completion). This percentage is analogous to the machine's confidence that it can perform the specified task, i.e. search for and locate the target within the search area before that target's deadline. This confidence, however, must be interpreted by a human operator only within the context of what that machine's software was programmed to consider (again: UAV speed, distance to target and sensor range) and not assume any additional considerations (e.g. weather, airspace, sensor health).

Representative of the real world, the SCOUT environment has restricted and controlled airspace that the operator must request access to use (and can be either approved or denied). As a military team member would, the UAVs alert the operator if its path intersects with controlled airspace, however it assumes that its superior, i.e. the human operator, still wants to fly that path regardless of whether authorization has been granted. The autonomous machine's displayed confidence remains dependent upon the path that the operator input. In this sense, the autonomous machine has been programmed to assume the human team member has superior rank and knowledge of the world, and therefore defers the higher-level decision to the human operator.

An experienced operator might decide to fly a direct path through a restricted area while waiting for approval or denial, but know, that access to that region is typically granted only 75% of the time. While waiting, the operator might be mentally preparing a contingency plan to ensure restricted airspace is not violated in the event of access denial or no response. As humans, we consider previous performance on a task to calibrate our confidence. Current UAVs lack "memory" or a historical dataset to mine and compare success or failures in similar situations. As a result, an experienced operator and current autonomous machine's confidence could vary significantly.

These simple examples demonstrate that the self-confidence of an autonomous machine is directly influenced (and sometimes limited) by its input and experience. An autonomous machine can only consider what it is programmed to consider and nothing more. Therefore, it is critical to ensure that all human operators have adequate training and understanding of the assumptions and input factors considered in automatic machines. The machine's programming was different from the operator's boot camp and the machine's memory of similar previous missions is non-existent.

2. What factors influence self-confidence?

As a reference, our military's initial training establishes a basis for the self-confidence of our new military members. This training provides experiences unique to the military that the graduate will have as a basis for self-confidence. For example, all our military services include a set of general orders. The following are from the Navy's Recruit Training Command (NRTC, 2012):

1. To take charge of this post and all government property in view.
2. To walk my post in a military manner, keeping always on the alert, and observing everything that takes place within sight or hearing.
3. To report all violations of orders I am instructed to enforce.
4. To repeat all calls from posts more distant from the guard house than my own.
5. To quit my post only when properly relieved.
6. To receive, obey, and pass on to the sentry who relieves me all orders from the Commanding Officer, Command Duty Officer, Officer of the Deck, and Officers and Petty Officers of the Watch only.
7. To talk to no one except in the line of duty.
8. To give the alarm in case of fire or disorder.
9. To call the Officer of the Deck in any case not covered by instructions.
10. To salute all officers and colors and standards not cased.
11. To be especially watchful at night and during the time for challenging, to challenge all persons on or near my post, and to allow no one to pass without proper authority.

Trainees are required to memorize these orders and they stand watches early in their military service gaining confidence in their ability to successfully carry out these orders. These general orders provide a list of capabilities that a military member is expected to have (1-4, 8, 10, & 11), some limitations on their behavior (5, 7), an opening for additional orders (6), and what to do for situations not covered by the other orders (9).

A set of general orders along the same lines could serve as the basis of self-confidence for military autonomous systems. Here is a possible set:

1. Each military autonomous machine is supervised by one and only one human operator at all times, but it must recognize orders from higher authorities and handle them the same way a junior military member is instructed to. (The machine must acknowledge the new order but if there is a conflict between the new orders and the previous orders, the machine must communicate that to the source of the new order.)
2. Each machine is expected to immediately acknowledge and perform its orders (in the sense of responding "aye, aye").
3. If unable to perform these general orders or a part of its current specific orders, the machine shall immediately notify its supervisor. It would be appropriate to be prepared to provide the reason why it cannot if asked.
4. The machine must be able to receive changes to current orders or new orders as specified in its current orders (allowing operators to intervene in the execution of pervious orders).
5. The machine may automatically navigate to specified waypoints by the best route (as defined by the orders).
6. The machine will operate officially attached subsystems as specified by its orders.

7. The machine will keep safe itself, its orders, its data, and its officially attached subsystems within the proscribed limits. (For reference, Asimov's Three Laws of Robotics (Asimov, 2004) would be beneath these like self-preservation is to humans.)

These general orders provide the basis for the factors affecting self-confidence. The autonomous machines must be able to evaluate its capabilities with respect to these general orders and its current orders continuously and, in accordance with #3, report its status if unable to perform the orders.

In addition to general orders, the Navy has an instruction applicable to all members that establishes the standard organization and regulations of the U.S. Navy (Navy 2012). In this instruction duties, responsibilities, and authorities are spelled out for Navy units. Applicable here is the discussion on accountability, paragraph 1.3:

a. Each Sailor, regardless of rank or position, is fully accountable for their actions, or failure to act when required.
b. Leaders and supervisors have a duty to assign clear lines of authority and responsibility, reaching to the deck-plate level, for all activity within their organization.
c. Leaders and supervisors have a duty to provide their subordinates the resources and supervision necessary to enable them to meet their prescribed responsibilities.
d. Leaders and supervisors have a duty to hold their subordinates accountable, and to initiate appropriate corrective, administrative, disciplinary, or judicial action when sailors fail to meet their responsibilities.

While these paragraphs can be the basis of guidance for the supervision of military autonomous machines, they also describe what the machine should be able to expect or confidently request from its supervisor.

Example: Factors influencing self-confidence in current and future UAVs.

An important factor influencing an individual's self-confidence is his or her metacognitive regulation, i.e., assessment of his or her knowledge, ability and understanding of the task relevant factors relating to a goal (Livingston, 1997). The military's training within these general orders and operational experience influences an individual's assessment of his or her capabilities. UAVs, on the other hand, base their assessment of ability (i.e. self-confidence) to perform specific tasks on predefined algorithms, which do not yet consider historical data of that UAV's performance.

When a machine computes self-confidence on an assigned task, the human operator should be able to investigate the factors that influenced the machine's output. This is not a trivial challenge though, since simply providing a human operator access to the algorithm might be analogous to providing a Swedish operator an explanation in Portuguese. The contents of the algorithm must be provided in a human readable and easily digestible format. Additionally, the machine should be informed of the factors that influenced success (or failure) on previous similar missions. This is a current limitation of UAVs, as they conduct each new mission with no information about the success of a previous mission. Future machines will need to consider previous experience in order to operate with greater self-confidence and independence.

In addition, machines should prompt their human operators to provide clarification or additional input when calibrating their self-confidence in accomplishing a specific task. To illustrate this idea with a simple example, suppose a UAV is directed to collect imagery in a specific geographic region before nightfall and the UAV can reach, cover that entire area, and return to base before sunset. Given information about the distance, speed, and size of the imagery collection area, the UAV would calculate 100% image acquisition/mission success before sunset. However, assume the operator failed to provide the weather forecast for the area and heavy cloud cover adversely impacted the sensor's data quality. The autonomous machine should, in this example, learn from the error and ensure future imagery collection missions prompt the operator for the latest weather forecast, prior to the machine providing an estimate of self-confidence in mission success. Just as people conduct after action reports, future autonomous machines should also review and learn from each mission, and identify factors impacting mission success, i.e., machine self-confidence.

In addition to an individual's or machine's assessment of their capability to perform a task, is their assessment of whether they *should* complete the task. Explicit input regarding the mission goals and priorities is needed in order to determine self-confidence in one course of action versus another. For example, violating airspace restrictions or running out of fuel are certainly situations that a machine should avoid, however certain missions where human lives are at risk, for example, might dictate that these factors are no longer relevant. In a situation where troops are in need of medical supplies, a vehicle's ability to safely return to base may no longer be a priority, i.e. human life would override considerations of vehicle "life". How these mission goals/priorities are provided to the system and how the machine is told to weigh certain rules becomes critical in these kinds of situations. Furthermore, providing easily understandable reasoning and communications is necessary to ensure that a human does not override a decision due to a failure to understand the intent.

3. How can/should self-confidence be computed and communicated?

The computation of self-confidence is, of course, dependent on the functions the system is expected to be capable of performing and the system's ability to detect possible external obstacles to its successful performance of its orders. So, a system must be able to detect and recognize foreseeable hazards. Not all obstacles may be foreseeable and therefore the system also needs to be able to evaluate its performance independent of externalities. This concept could bring up the question on consciousness.

A military autonomous machine needs to be "conscious" enough to be able to develop an informed opinion of its capabilities. This means it needs three things to develop its self confidence: (1) a self-image (internal model), (2) ability to perceive the current and projected environment, and (3) the ability to compare the orders, its self-image, and the environment to develop its self-confidence. Members of the military are trained, tested, and experienced to develop such capabilities. Military autonomous machines must also be so equipped.

Furthermore, there are two aspects to communicating this self-confidence that should be addressed: the content of such a communication and the method of its delivery. First, as providing your birthday is not an appropriate to answer a request for your age, an autonomous system providing its operator every operating parameter does not communicate its own assessment of its self-confidence. The core content is whether the autonomous system has a justified evaluation of its current capabilities with respect to its orders and foreseeable environment. The operator may want to know all the inputs to this evaluation to be able to make the assessment independently. This brings up the topic of trust.

Consider an elevator. It is trusted autonomy. The system may display where the elevator(s) currently is/are, but you have to make a request to get one to pick you up. When it arrives, you request a floor. Those were inputs to the autonomous system. However, the machine may not take you directly to your floor. The elevator may autonomously stop at intermediate floors to pick up other riders. It does so, not to meet your needs, but to fulfill its mission. If it was working for you, it would go directly to your floor and then return to pick up those other riders. But its mission is to also be efficient. Notice that the elevator does not inform you of its plans to stop at intermediate floors. It also does not inform you when it is being held at a floor or when it is out of service. It does not communicate self-confidence. That will not be tolerated behavior for a military autonomous machine. A military autonomous machine will be expected to report when it cannot perform its mission as ordered.

The second aspect of this question to address is the method of delivery of its self-confidence. The orders for the machine need to address the method of communicating, including communicating its self-confidence. Under many circumstances, the operator may expect continuous reporting of the self-confidence of the machine. If there are communication bandwidth concerns, some frequency of reporting may be specified or only provided when requested. If its mission is intended to be covert, the operator may order that the machine not report its self-confidence at all.

Example: How self-confidence is calculated and communicated by SCOUT controlled UAVs.

Most current small tactical UAVs (STUAS) do not have direct access to airspace information and are reliant on an operator to actively push that information to the UAV. Restricted Operating Zones (ROZs) can pop up unexpectedly and coordination within controlled airspace requires communication with the airspace's manager. As such, current STUASs can rarely reliably compute or convey confidence about their position relative to any airspace restrictions. A typical reason for controlled airspace is to eliminate the need for visual separation (detection and avoidance) of aircraft, which current UAVs are not capable of performing. In fact, most current STUASs lack transponders and automated collision avoidance systems, which not only detect collisions but also re-routes each vehicle. This means that safety in environments in which multiple vehicles are operating can only be assured when the operator works with an airspace manager to maintain separation and pushes that information to the vehicle.

This exemplifies the fact that autonomous machines cannot even begin to act truly autonomously, i.e., without the assistance of a human, unless the machine is provided access to all the same information that a human has. Confidence calculations based on a subset of information can also cause large misunderstandings within a machine and human team if assumptions are not clear. As discussed, the human teammate must be provided the ability to access all the information and assumptions which the autonomous machine used to guide its output. Even with all the same information though, algorithms are notoriously brittle (Cummings, 2006) and new situations which have not been explicitly considered or are fully understood can still lead to inaccurate calculations which require a human operator's skeptical interpretation.

As the saying "garbage in, garbage out" exemplifies: a calculation is only as good as what goes into it. In order to compute confidence on anything, access is required to the right information and input. It follows then, that it is necessary to identify what information is needed to successfully complete a specific mission. As noted earlier, this is an area in which the automation on UAVs falls short, compared to the innate ability of humans, since UAVs do not learn from previous data and missions. In a sense, UAVs compute their confidence in every mission as though it were the first mission they have encountered, regardless of whether they have flown that same mission 1 or 1,000 times previously. Future systems will need to rectify this on order to provide more accurate calculations of self-confidence.

In addition, the operator should be able to hone the machine's confidence calculations by placing emphasis on different elements of the mission, e.g. provide an estimate of confidence in getting to location X, given the current weather forecast, and ignore need to return to base. In the context of the SCOUT simulation environment, a large portion of tasking involves airspace management and path planning. NRL is currently developing decision support that enables the automation and the human to work as a team and develop various path plans, given different user-input priorities and accepted risk.

For example, if an operator is given 7 targets with different priority levels and deadlines, a conservative plan (only pursue targets with deadlines that enable 100% search of the area) versus risky plan (pursue high value targets, even if only 10% of the search area can be covered before the deadline) would yield entirely different plan recommendations. Depending on the mission and commander's intent, a plan with A) 10 - 30% confidence in finding 3 high value targets, might be considered superior to B) 100% confidence in finding 4 low value targets. This demonstrates the importance of clear communication of output and assumptions, so that the human operator can consider options and courses of action.

Alerting is one mechanism by which automation communicates ability and self-confidence to a user. Alerts implicitly inform the operator that the UAV has encountered a situation it is not programmed to handle, or that it cannot accomplish without violating some pre-determined rule (e.g., a previous order to not violate restricted airspace). Programming simple alerts (e.g., report if engine temperature surpasses X degrees) allows the human operator to attend to other tasking. However, alerts can become problematic if thresholds are too low and cause constant interruptions, or too high and problems are missed (Parasuraman & Riley, 1997).

Default states, such as never violating restricted airspace, could be set in advance of a mission, as a rule, in case the operator is overloaded and cannot immediately attend to an issue. If one of these rules (e.g. do not fly through unauthorized airspace, and instead take the shortest path around the restricted airspace) is performed however, it is critical that the human supervisor is informed and presented with all the factors that the machine considered. This will enable the human operator to override the decision if the mission necessitates making a riskier decision. Failure to properly communicate output of a machine can lead to large errors in human operator interpretation of the situation, similar to the events preceding the Three Mile Island nuclear reactor accident.

Self-Confidence of Autonomous Machines within Teams

Self-confidence of a machine is necessary to be a good teammate (Bolstad & Endsley 1999). Being part of a team includes backing each other up. Self-confidence is involved in monitoring teammates, understanding roles in a mission's accomplishment, and responding appropriately to unexpected direction. As a team member, an autonomous machine should monitor the behavior of the other team members and prompt appropriate action, if directed to do so. Therefore, the machine could appropriately monitor its human operator's behavior to determine when communications are lost or not provided when expected/necessary. The machine part of the autonomous system can have capabilities the human member could rely on, such as reminding the operator when new directions are needed.

An autonomous machine teammate can also perform better if it is included in understanding the shared mission, i.e., the commander's intent (Bolstad & Endsley 1999). It can then have a context in which to evaluate orders and priorities and respond appropriately if an order is beyond its capability. Even in a strict military environment, it is appropriate for a recruit to seek clarification when a commander issues an order with the assumption that it is within the recipient's capability, but the recruit lacks confidence in their ability to complete the order. For example, the recruit receives an order to drive to a specific position that would require crossing terrain outside of what the recruit believes the vehicle can safely traverse. Finally, providing some mission priority would also help an autonomous system and autonomous machine understand whether the mission is more important than the safety of the machine or the machine is more important that a particular mission. In summary, as a member of a team, the autonomous machine needs to be treated as more than just a dumb machine.

Summary Thoughts

Self-confidence is a critical aspect of success for the warfighter and is built up throughout his/her time in the military. The concept of self-confidence within autonomous systems is complex. As the warfighter becomes increasingly dependent upon autonomous machines, acting as their supervisor, the self-confidence of the human operator to perform his/her mission is largely based upon his/her confidence in the machines' ability. In turn each of the machines the human is supervising should develop its own level of confidence in its ability to perform the assigned mission and relay that information back to the operator. This is analogous to the human having to report up their chain of command when they are tasked with a specific order. There are still many research and philosophical questions about how an autonomous machine should develop self-confidence. Does self-confidence require memory so that machine can learn from previous performance and have a more accurate self-assessment? Can the system become confident enough that it, like the warfighter, can question its supervisor?

NRL is currently exploring how and when to engage the user in situations where the human's decision making is determined to be sub-optimal and automation could potentially work with the user to provide better alternative recommendations. Within SCOUT, the system is continuously evaluating an operator's plan and comparing it to all other possible plans as the environmental context changes. Research questions also abound regarding the communication of alternate options to the human operator. For example, how does the machine develop the confidence to know when interrupting the operator is appropriate? Is the gain from changing plans large enough to warrant interrupting the operator? Has the operator's strategy changed since the mission began and render the alternative suggestion useless?

Furthermore, our fate as limited-capacity processors (Kahneman, 2011) susceptible to error perhaps calls for a restructuring of the traditional unidirectional paradigm of man monitoring machine; and suggests that a bi-directional teammate relationship of monitoring and communication could yield better mission performance. An experienced human team member picks up on cues from their supervisor and for example, knows when an interruption is appropriate. Information about an operator's state, gathered from physiological (e.g., eye gaze) and behavioral (e.g., keystroke) data, may be able to provide autonomous machines with additional information to base their confidence in whether the human-machine team can successfully accomplish their tasking. Mission performance might be augmented if autonomous machines could serve as a trusted team member and confidently provide assistance by monitoring a human's decisions and questioning orders when they contradict the machine's understanding of the mission goals. The future possibilities and their associated research questions are endless.

Acknowledgements

The authors would like to thank the symposium organizers for encouraging thinking on this topic. The effort has furthered work at the Naval Research Laboratory on the interface between human operators and autonomous machines. The Office of Naval Research's 2015 Summer Faculty Research Program (SFRP) sponsored this work, in part. The views and conclusions contained in this document should not be interpreted as necessarily representing official policies, either expressed or implied, of the US Navy.

References

Army (2014) Army Regulation 600-20: Army Command Policy. 6 Nov 2014.

Asimov, I. (2004). *I, robot* (Vol. 1). Spectra.

Bajracharya, M., Maimone, M.W., Helmick, D. (2008) Autonomy for Mars Rovers. Computer 41(12):44-50.

Bolstad, C. A., & Endsley, M. R. (1999, September). Shared mental models and shared displays: An empirical evaluation of team performance. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 43, No. 3):213-217. SAGE Publications.

Cummings, M. L. (2006). Automation and accountability in decision support system interface design. *Journal of Technology Studies* 32(1):23–31.

Defense Science Board (2012) Task Force Report: The Role of Autonomy in DoD Systems. Department of Defense.

Department of Defense. (2009). FY2013–2038 Unmanned Systems Integrated Roadmap. Department of Defense.

Department of Defense. (2012) DOD Directive 3000.09: Autonomy in Weapon Systems. 21 Nov 2012.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Livingston, J.A. (1997). Metacognition: An overview. Available: http://www.gse.buffalo.edu/fas/shwell/cep564/Metacog.htm.

NASA (2013) NASA's Mars Curiosity Debuts Autonomous Navigation. Available: http://www.jpl.nasa.gov/news/news.php?feature=3884.

Navy Recruit Training Command (2012) General Orders. Available: http://www.bootcamp.navy.mil/general_orders.asp.

Navy Recruit Training Command (2013) Recruit to Sailor. Available: http://www.bootcamp.navy.mil/recruit_to_sailor.asp. .

Navy (2012) OPNAV INSTRUCTION 3120.32D: Standard Organization and Regulations of the U.S. Navy. 16 July 2012.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, *52*(3), 381–410.

Sibley, C., Coyne, J., & Morrison, J. (2015) Research Considerations for Managing Future Unmanned Systems. 2015 AAAI Spring Symposium on Foundations of Autonomy and Its (Cyber) Threats: From Individuals to Interdependence. AAAI Press.

TRADOC (2008) Training and Doctrine Command (TRADOC) Pam 600-4: Initial Entry Training Soldier's Handbook. 23 Dec 2008.