JANUARY 2017

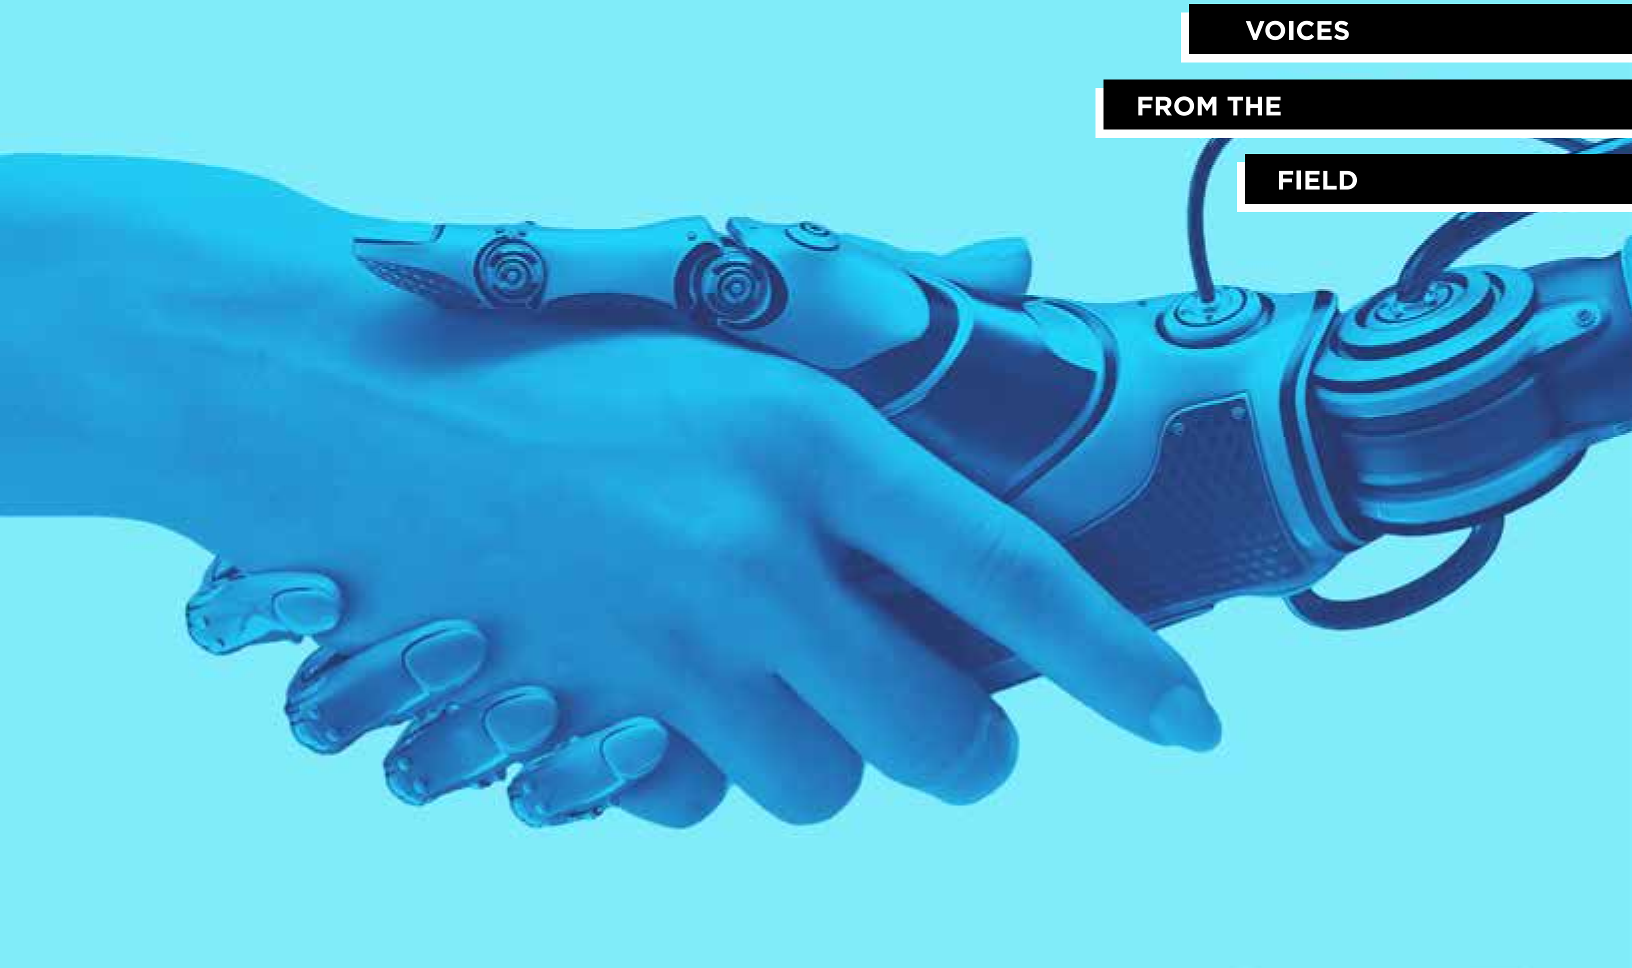

PATRIOT WARS

Automation and the Patriot Air and Missile Defense System

Dr. John K. Hawley

Center for a
New American
Security

About the Author

DR. JOHN K. HAWLEY is an engineering psychologist with the U.S. Army Research Laboratory's Human Research and Engineering Directorate. He has more than 35 years of experience with Patriot and other Army air and missile defense systems.

Acknowledgements

I would like to thank Paul Scharre, Shawn Brimley, Maura McCarthy, and Melody Cook for their assistance in the development of this report. I am solely responsible for the contents of the article, and any errors of fact or omissions are mine. The views expressed in this article are mine and do not necessarily reflect those of the U.S. Army Research Laboratory, the U.S. Army, or the U.S. Department of Defense.

Ethical Autonomy Series

"An Introduction to Autonomy in Weapon Systems," by Paul Scharre and Michael C. Horowitz (February 2015)

"Meaningful Human Control in Weapon Systems: A Primer," by Michael C. Horowitz and Paul Scharre (March 2015)

"Autonomous Weapons at the UN: A Primer for Delegates," by Paul Scharre, Michael C. Horowitz, and Kelley Sayler (April 2015)

"Autonomous Weapons and Operational Risk," by Paul Scharre (February 2016)

About the Voices from the Field Series

The Voices from the Field series is designed to feature timely analysis gleaned from current or former U.S. government practitioners who have recently served either overseas in strategic areas of the world or who have participated in critical policy deliberations. This series also highlights observations and current assessments based on field research conducted across the globe by CNAS experts.

Cover Photo
Shutterstock/Adapted by CNAS

Preface

By Paul Scharre

Automation and autonomy are core components of the Department of Defense's "third offset strategy," designed to reinvigorate American military technological dominance. Effective collaboration between humans and machines is central to harnessing the advantages of automation and autonomy. As Deputy Secretary of Defense Robert Work has explained, it is "human-machine collaboration and combat teaming" that will turn rapid advances in autonomy and artificial intelligence into operational game changers.[1] As the U.S. military begins to grapple with the challenges of incorporating ever more sophisticated autonomous systems into the force, it can draw upon decades of experience with highly automated combat systems, including some that use lethal force. Unfortunately, this experience has not always been positive.

In 2003, during the initial stages of the Iraq invasion, the U.S. Army's Patriot air defense system was involved in two fratricide incidents, shooting down a British Tornado and a Navy F-18 fighter jet. The Patriot is a highly automated system, and the causes of the fratricides were a complex mix of human and machine failures. As automation and autonomy become increasingly incorporated into weapon systems, the lessons learned from the Army's experience with Patriot are vital for understanding the role of humans and automation in lethal systems.

Dr. John Hawley is an engineering psychologist with the U.S. Army Research Laboratory's Human Research and Engineering Directorate and has extensive experience in Patriot operations. He has more than 35 years of experience in human-machine interactions in air and missile defense systems, and led the Army's internal efforts to improve vigilance in Patriot operators following the 2003 fratricides. In this report, Dr. John Hawley shares his perspective reflecting on lessons learned over several decades of experience in human-machine integration in combat systems. These hard-won lessons provide valuable insights into the roles of human and machine intelligence in combat systems and best practices to avoid future accidents.

PAUL SCHARRE is a Senior Fellow and Director of the Future of Warfare Initiative at the Center for a New American Security. From 2008–2013, Mr. Scharre worked in the Office of the Secretary of Defense where he played a leading role in establishing policies on unmanned and autonomous systems and emerging weapons technologies.

Introduction

The use of automation in the modern workplace has had many consequences, both positive and negative, both intended and unintended. Automation in various forms is increasingly being used in a range of weapons systems such as the Army's Patriot air and missile defense system. Moreover, it has become commonplace in aircraft flight control systems, and in prototype self-driving cars that have been traversing streets and highways for several years. Applications of automation in future weapons systems and related uses are expected to proliferate and grow in the years to come. Many observers are calling for a candid discussion of appropriate roles for automation in military systems. This is particularly true now that some of these systems are approaching the threshold for autonomous operations.

To some observers, the use of automation in many of the applications cited above is relatively new. These observers write about such developments as if they are recent, and as if we collectively do not have much experience with automation applied to the development of autonomous or near-autonomous systems. That's not altogether true. Some potential applications of automation technology, like self-driving cars, are relatively new, but other applications, such as near-autonomous air and missile defense systems or extensive flight deck automation in aircraft, have been around for quite some time. Moreover, we have a fair amount of operational experience with existing systems, and that experience has not all been positive. When I read the descriptive literature and claims for some of the newer applications of automation, such as self-driving cars, I find myself wondering whether their proponents either are not aware of our history with these older systems, or tend to view experiences with older systems as not relevant to their "new" and more advanced uses of this technology. Perhaps the idea is, "We're better now, and that old stuff doesn't apply." It is true that automation technology is getting better, but the latter assertion is not necessarily true. There are lessons and pitfalls associated with the use of automation in older systems that apply directly to what can be expected with newer applications. A number of these lessons apply to the humans' residual role in system control, and how difficult that role can be to prepare for and to perform.

What follows is a mostly personal story. I have been in the somewhat unique position of having had a long-term, hands-on association with an early application of automation in weapon system control. The application in question is the Patriot air and missile defense system. The next portion of this paper traces my personal history with Patriot going back more than 35 years. During this time, my views regarding automation and autonomy have evolved considerably, based on extended hands-on experience with that system. I'll state upfront that I'm not as optimistic regarding the safe and effective use of automated and near-autonomous systems as I once was. In this respect, the paper also outlines a number of lessons and cautions derived from my experiences with Patriot. I think these apply to many of the potential applications of automation technology currently being discussed. They go beyond the technology employed and also apply to the personnel and organizations charged with safely and reliably using that technology. In fact, the technology component may be the easiest of all to address. I have observed first hand that human aspects of automation are often the most difficult to resolve.

TERMINOLOGY

Automation: Automation refers to control of a system by mechanical or electronic devices that take the place of human observation, information processing, decision making, or effort.

Automation Technology: Automation technology refers to the technical enablers (e.g., information technology, software, artificial intelligence) underlying automation.

Manual Control: Most or all aspects of task performance are performed by human operators.

In-the-Loop Control: System control exists on a continuum ranging from manual control to full system autonomy. A control mode in which humans retain selected key functions and make all or most decisions is referred to as in-the-loop control. Human operators are an integral part of the system's control loop.

Supervisory Control: The system controls all aspects of operations automatically, but human operators can set goals and intervene as needed. Under a supervisory control regimen, the human operator does not control the system directly. Rather, the operator receives system status information from a machine intermediary (typically a computer). The human operator monitors this control information and intervenes when necessary to keep system performance within desired limits.

On-the-Loop Control: Another term for supervisory control. The operator sets goals, monitors system actions, and intervenes when necessary.

Autonomy: The system's on-board control algorithms provide for full control of all aspects of system operations without human guidance or the ability to intervene.

Near-Autonomous: A term sometimes used to denote a high level of supervisory control. Human operators set goals and monitor system performance loosely, but retain the ability to intervene as judged necessary in critical aspects of system operations, such as overruling a track engagement decision.

Terms are defined as they are used in this paper. These specific definitions may not be identical to definitions used in official Department of Defense documents or other CNAS publications.

A Personal History with the Patriot System and Its Automation

Patriot was one of the first tactical systems in the U.S. Department of Defense's (DoD) inventory to employ what is now termed "lethal autonomy" in combat. Lethal autonomy refers to a system that is capable of applying lethal force with little or minimal direct human oversight. My initial contact with the Patriot system was in the late 1970s. I was fresh out of graduate school with a PhD in psychology but had some experience with predecessor air defense systems, such as Nike Hercules and Hawk, as an air defense officer in the early 1970s. Patriot was a somewhat different experience. The system has two operating modes: semi-automatic and automatic. Patriot in semi-automatic mode is slightly more automated than its immediate predecessor the Hawk system, but still on that I would term the "main line" of evolutionary development for air defense systems of its class. That is, the system provides more computer-based engagement support than its predecessors, but Patriot in semi-automatic mode is still very much an operator-in-the-loop system. Patriot in automatic mode represented a significant jump in capability. In that sense, there was a discontinuity between Patriot in semi-automatic mode and Patriot as it could be used in automatic mode.

Patriot's automatic mode is quite different. So different, in fact, that I once asked one of the prime contractor's systems engineers where they got the engagement-control algorithms used in the system's automatic mode. He replied that they had been adapted from the engagement control logic of the Safeguard system. Safeguard was the first operational U.S. anti-ballistic missile (ABM) system. The system was deployed briefly beginning in the early 1970s and then traded away as part of one of the first treaties limiting U.S. and Soviet ABM systems. Remnants of the old Safeguard system still exist at Ft. Bliss, Texas, and at isolated sites in Montana and North Dakota.

Safeguard was a near-autonomous system. Get a green light to initiate the missile engagement process, and the system mostly took over from there. The computer fought the air battle. That was a reasonable choice, given Safeguard's mission and operational context: Fight the first salvo of the Battle of Armageddon at the edge of space. However, that level of automation was not an appropriate operating mode for Patriot's mission and operating environment. Patriot operates in the more cluttered and ambiguous lower-tier region of the air defense operational environment. The potential for track classification and identification mistakes is considerably greater for Patriot than it was for Safeguard. The Army did not fully grasp the impact of these differences, and to some extent still does not. The major problem with Patriot is that the system's automatic feature is mostly an all-or-none operating mode. In automatic mode, there are few "decision leverage points" that allow the operators to influence the system's engagement logic and exercise real-time supervisory control over a mostly automated engagement process.

Beginning in the late 1970s and continuing through Patriot's initial fielding in January 1984, I was involved in a series of system development studies for Patriot. During that time, there was a school of thought in Army circles that using Patriot in automatic mode would be a preferred operating concept. Our early work lent support to the argument that automatic was not a suitable operating mode for Patriot against conventional air threats. Patriot's engagement algorithms were too "brittle" for the system's engagement context. Used in this context, "brittle" refers to the machine's inability to handle unusual or ambiguous tactical situations reliably. The term is now commonly used to describe automation limitations. The basic issue with brittleness is that computer-based algorithms operate in a black-and-white world; they have a little capacity to handle gray or ambiguous situations. That task falls to human operators, if they have the time and expertise to do so. When Patriot was initially fielded, tactical usage guidance directed that the system not be employed in automatic mode. The automatic mode was included with Patriot because it was available from Safeguard, and there were potential Cold War-related situations in which a mostly automated air defense system might prove useful. Safeguard was intended to be used in a nuclear war context in which all bets are off, so to speak, and risk tolerance is very high. That was not the case for Patriot. The initial version of the system was upgraded several times, beginning in 1988, to provide a limited ability to engage short-range tactical ballistic missiles. These upgraded versions were referred to as Patriot Advanced Capabilities 1 and 2 (PAC-1 and PAC-2). The current version is denoted as PAC-3.

The major problem with Patriot is that the system's automatic feature is mostly an all-or-none operating mode.

Patriot operations

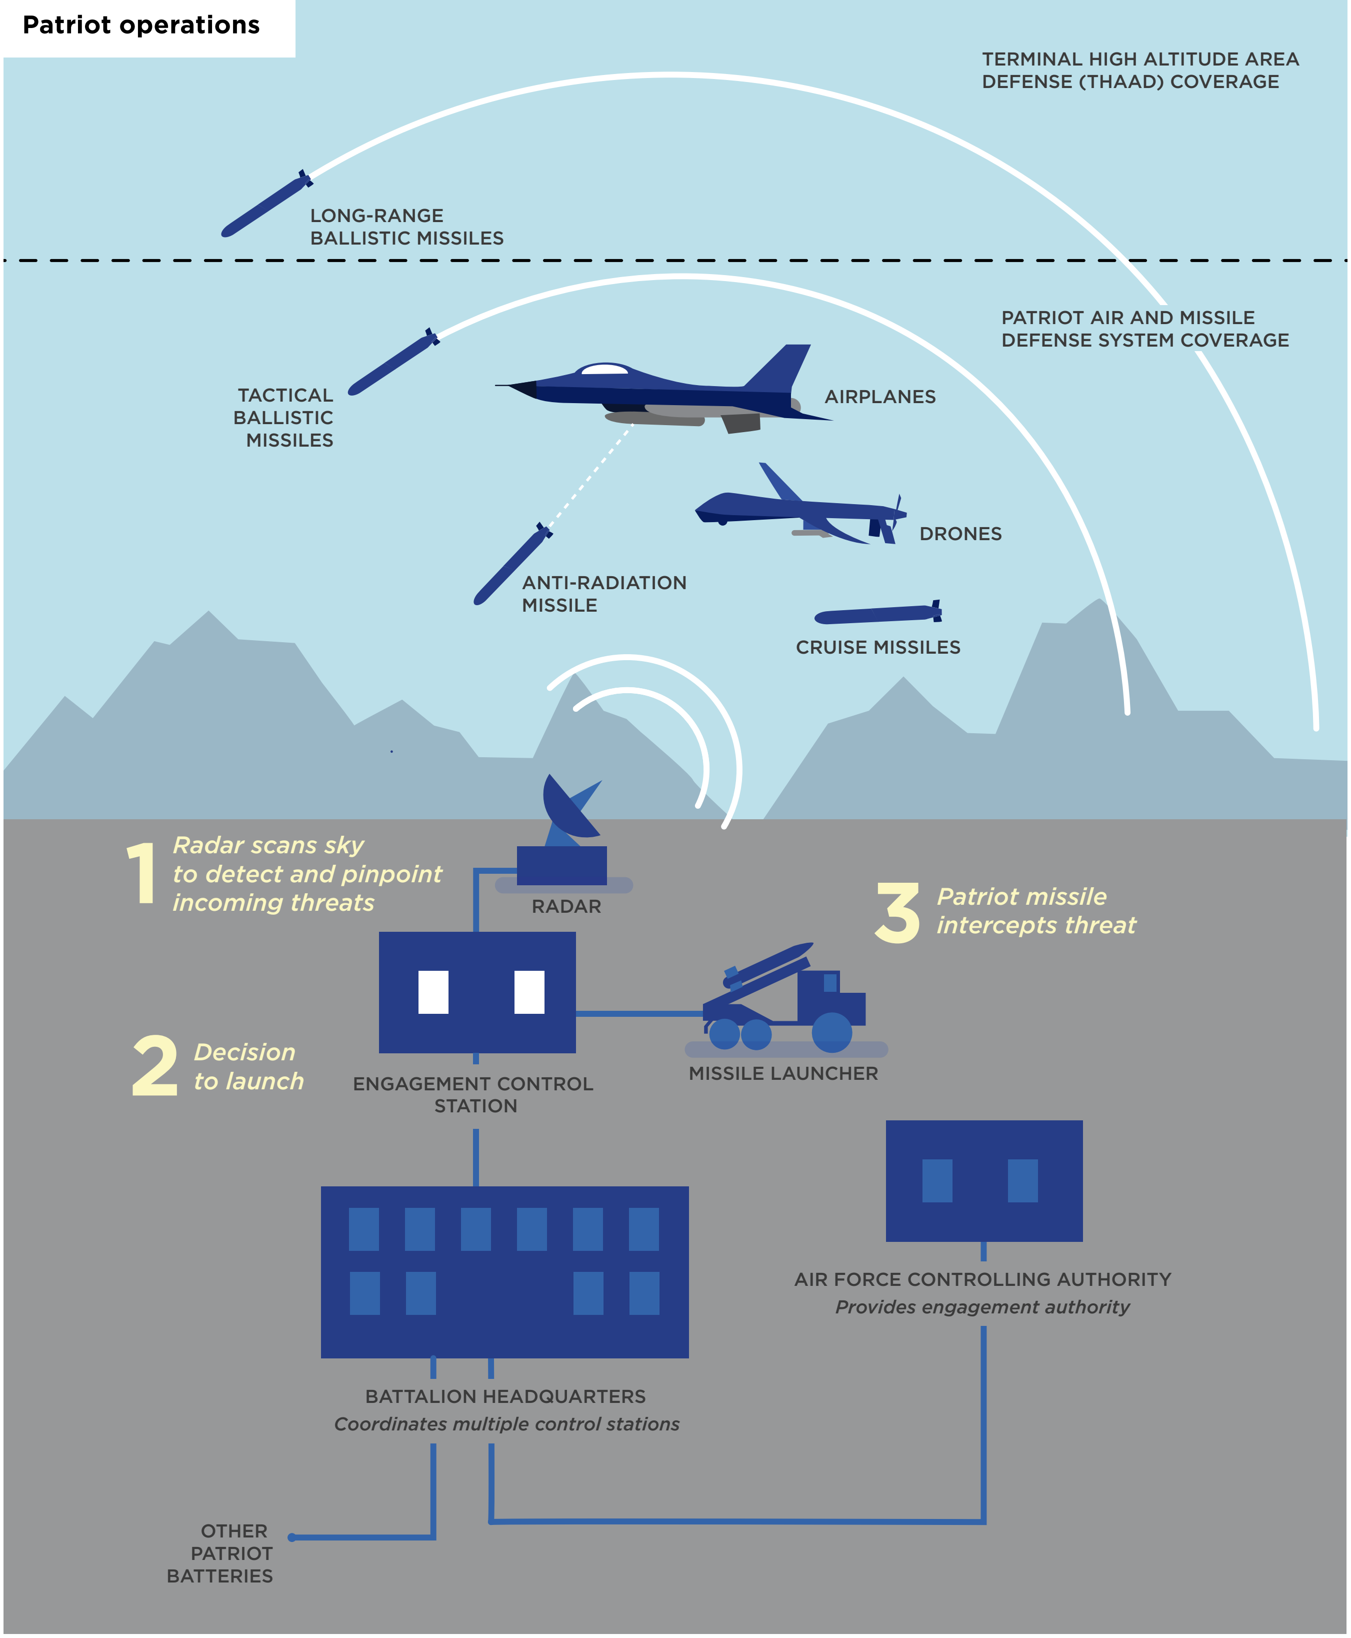
TERMINAL HIGH ALTITUDE AREA DEFENSE (THAAD) COVERAGE
LONG-RANGE BALLISTIC MISSILES
PATRIOT AIR AND MISSILE DEFENSE SYSTEM COVERAGE
TACTICAL BALLISTIC MISSILES
AIRPLANES
DRONES
ANTI-RADIATION MISSILE
CRUISE MISSILES
1 Radar scans sky to detect and pinpoint incoming threats
RADAR
3 Patriot missile intercepts threat
2 Decision to launch
ENGAGEMENT CONTROL STATION
MISSILE LAUNCHER
AIR FORCE CONTROLLING AUTHORITY
Provides engagement authority
BATTALION HEADQUARTERS
Coordinates multiple control stations
OTHER PATRIOT BATTERIES

Elements of the Patriot air and missile defense system, along with higher headquarters, arrayed against various threats. The Patriot covers aircraft, cruise missiles, and tactical ballistic missiles, while the Terminal High Altitude Area Defense (THAAD) system covers longer-range ballistic missiles.

Patriot Decision Cycle

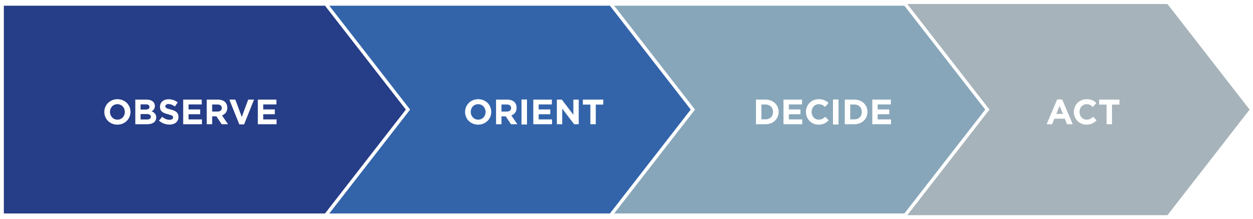

What is it? Whose is it? Radar detects and classifies object Humans apply outside information and context	Is it hostile? Is it a valid target? Establish situational awareness Apply rules of engagement	Engage? Decision whether or not to fire **Semi-automatic mode:** Human operator must authorize engagement or system will not fire **Automatic mode:** System will fire unless human operator halts engagement	System fires and missile maneuvers to target Human operator can choose to abort missile while in flight

Patriot's anti tactical ballistic missile capability was first used operationally against Iraqi Scuds (crude tactical ballistic missiles) during Operation Desert Storm in the early 1990s. Ballistic missile engagements mandate a higher level of automated support than provided by Patriot's semi-automatic mode. The nuts and bolts of the ballistic missile engagement process are too complex and time-limited for direct, in-the-loop human participation. I won't address the issue of Patriot's operational success against Scuds during Desert Storm, but the genie was now out of the bottle.[2] Patriot in automatic mode had demonstrated the potential for a desirable and timely new capability. When the Patriot operators flipped that switch to automatic and engaged the first Scud during Desert Storm, a brave new era of lethal autonomy was initiated. There were a number of anecdotal reports of fratricide "close calls" attributable to track classification and identification problems using the system's automatic mode during Desert Storm, but nothing out of the ordinary actually occurred. The Army left Desert Storm very full of itself regarding Patriot and its capabilities. Self-congratulation led to complacency, which led to unwarranted trust in, and reliance on, the system's automatic operating mode.

At that same time (1992), I was working on an automation applied to command-and-control project in support of the air defense community at Ft. Bliss, which was the home of the Army's Air Defense Artillery Center and School. Over the course of that project, I became familiar with the literature on humans and automation developed up to that point in time. Based on that literature and my previous experience with Patriot and predecessor air defense systems, it became clear to me that the Army was headed for trouble if they were to stay on the course they had chosen after Desert Storm: Employ Patriot in automatic mode. I had a number of conversations with system and training developers on this subject but got nowhere with my argument to be cautious using Patriot in automatic mode. To focus these discussions, I wrote the initial unpublished version of a paper titled "The Human Side of Automation: Lessons for Air Defense Command and Control."[3] That paper summarized the existing literature on humans and automation, and generalized that work to the case of air defense command and control. That project ended in late 1992, and I moved on to other human-systems integration projects. The Army continued to believe and act upon all the automation "myths" described in the contemporary human factors literature.[4]

Fast-forward 11 years to 2003 and Operation Iraqi Freedom (OIF) – the Second Gulf War. Patriot was involved in two fratricide incidents during OIF. (Two out of 11 ballistic missile engagements were fratricides.) The first involved a British Tornado and the second a Navy F-18. I was sitting in my office on Aberdeen Proving Ground, MD, one morning in 2004 when I received an email from our organization's director. His message to me was: "???" He had forwarded an email from our

Soldiers of the 11th Brigade, 43rd Air Defense Artillery fire a Patriot missile as part of Exercise Roving Sands '97 near El Paso, Texas, on April 30, 1997. (Tech. Sgt. James D. Mossman/U.S. Air Force)

on-site representative at Ft. Bliss, who had forwarded an email from the commanding general, Major General Michael Vane, asking a simple question: "How do you establish vigilance at the proper time? 23 hours and 59 minutes of boredom followed by one minute of panic." The director asked me whether I knew what General Vane was talking about. I replied, "Yes, I do. They shot down a couple of aircraft they shouldn't have." After some further discussion, he remarked that I did not seem surprised by those incidents. I replied that I was not surprised. Those outcomes had been in the card deck, so to speak, ever since they first flipped the engagement mode switch to automatic and assumed that was all there was to the conduct of near-autonomous operations. The fratricides were incidents waiting to happen. The director and I later traveled to Ft. Bliss to meet with General Vane. He wanted the Army Research Laboratory (ARL) to conduct a human-factors-oriented assessment of what had happened with Patriot during OIF and what the Army could do to avoid future incidents of that kind.

I won't recount all the observations and recommendations reported out under that project.[5] In brief, however, the Army had committed all the classic "sins" associated with the development and use of automated systems. They had trusted the system in a naïve manner; they had not adequately prepared their operators and crews for proper oversight of automated operations; and they had been unwilling or unable to confront the fact that near-autonomous operations are qualitatively different from old-style manual control (i.e., "on-the-loop" versus "in-the-loop" control). In short, they had failed to adapt to the complex new capability they possessed.

Ineffective Human-Automation Integration in Patriot

Let us now return to the issue of decision leverage points in Patriot's automated control logic. In the human factors literature, this issue falls under the topic of human-automation integration. One of the more interesting aspects of Patriot tactical operations after the first OIF fratricide incident (the British Tornado) was a decision to have fire units drop their launchers to standby mode. That way, the system could remain in automatic engagement mode but not actually engage a track until one or more launchers were returned to ready status. Commanders apparently wanted a "second look" before permitting the system to engage. The second OIF fratricide (the Navy F-18) took place under this modified operating regimen. The system reported a false ballistic missile track later attributable to radar electromagnetic interference. The tactical director at the battalion command and control node gave the order, "Bring your launchers to ready." That directive was tantamount to an order to engage. But that was not what the tactical director intended; he simply wanted to get ready to engage by bringing fire unit launchers to ready status. The subordinate battery fire units were in tactical ballistic missile automatic mode. The tactical director either did not know that, or he did not remember in the heat of impending action that returning launchers to ready status would result in an automatic engagement by the first available launcher. The F-18 was engaged and destroyed.

A later Army board of inquiry recommended that the tactical director be issued a general officer reprimand for not terminating the engagement. In an obvious example of hindsight bias, the board determined that there was sufficient evidence available at the time to have terminated the engagement after missile launch. I thought the reprimand was unwarranted. In both fratricide incidents, the Patriot crews did what they had been trained to do, which was reinforced by the prevailing command climate and widespread, but not generally accurate, beliefs about the system's engagement reliability. In retrospect, I have never believed that the launch crew knew for certain they had engaged the F-18. They shot at what the system initially determined was a tactical ballistic missile. However, that was a false track – there was no ballistic missile. When Patriot's PAC-3 missile approaches its intended target, the missile deploys its own seeker. It is a hit-to-kill weapon. The missile "looked" for the ballistic missile, but there was no ballistic missile. However, it "found" the F-18. The F-18 was simply in the wrong place at the wrong time.

A Patriot missile battery watches a Turkish army base in Gaziantep, Turkey, near the country's southern border in February 2013. U.S. and NATO Patriot missile batteries were deployed to Turkey to assist in defending Turkey in response to the ongoing civil war in Syria. (Glenn Fawcett)

Army "big missile" air defense units such as Patriot function under the operational control of the Air Force. After the second fratricide, the Air Force denied Patriot units any engagement authority, even in self-defense. The Tornado incident was a permissible self-defense engagement against what the system classified as an anti-radiation missile. Under the new rules of engagement, Patriot could engage only when specifically authorized by the Air Force controlling authority. Tactical ballistic missile engagement timelines are often too short for that to be a practical course of action. In essence, that decision took Patriot out of the fight, so to speak. There were no further Patriot launches during OIF, and, luckily, there were no more ballistic missiles to shoot. Similar engagement restrictions on Patriot operations are still in place: the Air Force retains engagement authority for any Patriot shots.

I had a later conversation with the senior officer who had led the Army's inquiry into the OIF fratricides. We got into a discussion of the board's findings and suggested remedial recommendations. He asked what I thought about those findings and recommendations. I replied that the board's conclusions and recommendations were supported by the available human factors research. I also asked him whether he would like to know more about why those incidents should have been expected. I gave him my old working paper on automated command and control –then 12 years old. A few days later, I was having a conversation with a senior warrant officer who had been one of the lead technical specialists on the investigating team. I noticed that he had a copy of my old report on his desk. He looked at me strangely and asked whether I actually had written that report in 1992. I replied that I had. He then stated that the report "predicted everything that happened to Patriot during OIF." I do not think the Army and the Patriot weapons community are alone in this respect. I have seen the same pattern of selective inattention to humans and automation research and experience in the development and use of other Army systems, in reports on mishaps with automated flight control systems, and with recent mishaps and fatalities involving self-driving cars.

Yes, many of the cautions regarding the potential pitfalls of automation and near-autonomous operations were known – even in 1992. But, no, the Army did not act on any of those research results, their own experiences with Patriot during Desert Storm, or subsequent operational tests where some of the same kinds of incidents had occurred. After Desert Storm, the Army proceeded to reduce the experience level of their operating crews; they reduced the amount of training provided to individual operators and crews; and what's worse is that they still have not fully corrected many of these deficiencies.

For roughly 20 years, automated air and missile defense systems such as Patriot have operated under a Title 10 mandate similar to what is now in the DoD's policy on the use of automation and autonomous systems.[6] Autonomous and near-autonomous systems must operate under what is termed "positive human control," a requirement that has never been clearly defined. In my experience, that requirement has had little impact on air and missile defense system development or operations. Decision makers, system developers, and users confidently assert that the requirement for positive human control is met even if that means not much more than having a warm body at the system's control station. One of the hard lessons of my 35 years of experience with Patriot is that an automated system in the hands of an inadequately trained crew is a de facto fully automated system. Moreover, I'm no longer sure what the term "adequately trained" means within the context of supervisory control of near-autonomous operations. Humans are very poor at meeting the performance demands imposed by supervisory control. The most problematic of these are the requirement for sustained operator vigilance and developing and maintaining the broad-based situation awareness upon which suitable intervention decisions can be based. In many respects, calling for reliable supervisory control over a complex automated system is an unreasonable performance expectation.

It is only fair to note that the Army made a number of changes in the aftermath of the OIF fratricide incidents. Perhaps the most important change is acceptance of the fact that the system is not always right. Patriot's command and control kill chain was modified to provide additional oversight of engagement decision-making. Training has been modified to include incidents similar to those encountered during OIF. Trainees are encouraged (and instructed on how) to query to the system to confirm or disconfirm its track classification and identification results. However, the length of institutional and collective (unit) training has not been increased substantially. Training times still fall short of what the literature on operator expertise suggests for jobs of Patriot's complexity. For the most part, the training changes that have been made are add-ons or modifications to older training curricula. New approaches to, and objectives for, operator and crew training have been recommended but have not been implemented. Moreover, Patriot operators and crews still do not remain in hands-on air battle management roles long enough to become truly proficient in their jobs. Routine Army personnel practices interfere with the development of essential levels of individual and crew expertise. During operational tests of Patriot software upgrades, incidents of the sort that occurred during OIF still occur. This is particularly true when test events go off-script, and operators are presented with situations they have not previously seen or explicitly trained to address.

Observations, Lessons, and Cautions

There is a tendency among system developers with little background in human performance theory to assume that automation is innately beneficial. For example, one of the purported advantages of self-driving cars is that they might provide considerable benefit in reducing the role humans play in causing car crashes. Research and experience in a number of areas suggests, however, that such expectations might not always prove to be accurate, or might be very long in coming. The paragraphs to follow highlight and discuss problems that frequently occur when automated systems are developed with little regard for the human component.[7] The context of that discussion is air defense command and control, but many of these observations, lessons, and cautions also apply to other areas in which automation might be applied.

Automated Systems Seldom Provide All Anticipated Benefits. Newly automated systems rarely live up to their initial billing. First-time users of automated systems must anticipate a debugging and calibration period during which the system's actual capabilities and limitations are determined. It is often necessary for field users to determine how they should practically employ the system, as opposed to unquestioningly using it the way system developers think it should be used. System developers often fail to anticipate operational problems that an automated system will create. Automation "surprises" should always be expected. Unquestioning acceptance of an automated system opens the door to what has been termed "automation misuse," or unwarranted over-reliance on, and trust in, automation. Automation misuse on the part of Patriot crews was identified as a major contributor to the system's fratricides during OIF.

Increased System Monitoring Load. Automation may change the nature of an operator's job, but it does not always simplify it. Automated systems often are characterized by a proliferation of components brought on by increased system complexity. Under an automation regimen, operators often have less to do moment-to-moment, but as a consequence of an increased number of components, they have more indications of system status to monitor. Vigilance can be a problem. Sustained vigilance involves hard mental work and can be stressful. As Major General Vane stated it in the email that launched ARL's fratricide investigation, "How do you establish vigilance at the proper time? 23 hours and 59 minutes of boredom followed by one minute of panic." It is very difficult for operators to maintain a high level of vigilance over a long period of time during which nothing out of the ordinary is happening. Expecting sustained high levels of vigilance by monitors of automated systems is an unrealistic performance expectation.

False Sense of Security. Belief in the system's infallibility (i.e., it's always right) can lull operators into a false sense of security, with the result that they will not make checks that would otherwise be advisable. Long periods of time during which the system operates successfully have been observed to lead to complacency, and complacency can result in reduced vigilance and a lessening of operational prudence.

Automation Transforms Operators into System Monitors Rather than Active Controllers. Automation does not remove human operators from the system. Rather, it moves human operators from direct, in-the-loop control of system operations to higher-level supervisory control tasks (that is, on-the-loop control). Problems can arise when the automated control system has been developed because it presumably can do the job better than a human operator, but the operator is left in to "monitor" that the automated system is performing correctly and intervene when it is not. Humans are very poor at meeting the monitoring and intervention demands imposed by supervisory control.

Automation does not always lessen operator training requirements.

Out-of-the-Loop Familiarity. When system operator tasks are replaced by automation, the operators' level of interaction or familiarity with the system is reduced. There is considerable evidence that when an abnormal situation does occur, operators will be slower to detect it and will take a longer time to jump back into the control loop and make the appropriate control actions. This problem is sometimes referred to as loss of situation awareness, or SA. Major General Vane's remark about "one minute of panic" was a direct reference to the Patriot crew's mad scramble to "get back into the loop" and reestablish SA, upon which suitable intervention decisions could be made.

There also appear to be longer-term consequences of being removed from direct, in-the-loop control. Operators may lose basic control proficiency as they receive less and less hands-on experience. This situation has been identified as a significant problem for pilots who rely excessively on automated flight control systems. It also will be a problem for "drivers" of future self-driving cars. Any notions that such drivers will be able to rapidly and seamlessly disengage from whatever they are doing and assume control from the vehicle's automation under unusual or ambiguous circumstances are not borne out by past experience with automated systems. Such control transitions will be problematic.

Increased Training Requirements. One of the most common myths about automation is that as a system's automation level increases, less human expertise is required. Contrary to this popular belief, automation does not always lessen operator training requirements. It frequently changes the nature of operator performance demands and increases operator training requirements. Automated systems tend to be more complex than their non- or less-automated predecessors. This increased complexity can make automated system operational skills more difficult to learn and retain. Moreover, operators often must have a deep knowledge of the complex systems under their control to be able to intervene appropriately when necessary. They have to understand how those systems "work," and how the automation's control algorithms dictate system actions.

Frequent simulator sessions or other types of operator in-the-loop training are often posed as means of combating problems associated with skill decay attributable to out-of-the-loop familiarity. There are, however, several inherent problems with the use of simulators to maintain supervisory control proficiency. Perhaps the most serious of these problems is the difficulty of training for extreme situations. These are the situations in which skilled human intervention is necessary. The skills required for performance during extreme situations are not always developed or maintained during routine training or while operating long-term in a supervisory control mode.

Automation Challenges Going Forward

The most problematic aspect of automation and autonomous operations is the human aspect, or human-automation integration. The popular concept of automation is that of a complex of machines performing their intended function with little or no human intervention. That is, a system is controlled either manually or automatically, with nothing in between. Experience has indicated, however, that all-or-none control is the exception rather than the rule. Automated systems that do not leave some residual functions for human operators are rare. It is this "residual functions" problem that leads to the list of operational human-automation integration problems discussed in the previous section. Simply put, human-automation integration in a fast-paced, real-time performance setting such as air defense command and control is a difficult challenge.

The residual functions issue coupled with the inherent difficulty of integrating humans with automated components has created a situation that has come to be known as the "dangerous middle ground" of automation – somewhere between manual control and full and reliable automation.[8] The current generation of automated systems is not reliably autonomous in the sense that they do not require human intervention at selected critical points in their operation. Contemporary automation technology is not yet that good. The so-called brittleness problem of automata remains an issue. At the same time, it is challenging for humans exercising supervisory control to intervene acceptably when something goes amiss, and they are required to perform some critical function. Vigilance limitations and the out-of-the-loop familiarity problem tend to make adequate intervention problematic. We are thus left on the dangerous middle ground between these two conflicting control dilemmas, and will likely be there for the foreseeable future. Operator "mistakes" like those leading to incidents like the Patriot fratricides during OIF will always be a possibility. Such incidents must be considered "normal accidents" in the sense that Charles Perrow uses that term.[9] The term normal accident indicates that given the system's characteristics, multiple and unexpected interactions leading to failure are inevitable.

The organizations employing autonomous systems also are important with respect to the prudent and effective use of such capabilities. The challenge facing any organization employing autonomous systems is developing the capability for sustained high reliability in a complex and unpredictable operational setting. This is particularly imperative for military organizations employing systems capable of lethal autonomy. Daunting as this challenge might seem, there is a class of organizations that have managed to do just that – maintain high performance in complex and unpredictable environments. Karl Weick and Kathleen Sutcliffe refer to such organizations as "high-reliability organizations."[10] Examples include air traffic control facilities, nuclear submarines, and aircraft carrier deck operations. Unfortunately, Army air and missile defense units employing Patriot did not, and still do not, meet the requirements for inclusion in the list of high-reliability organizations.

Human-automation integration in a fast-paced, real-time performance setting such as air defense command and control is a difficult challenge.

High-reliability organizations are characterized by ways of acting and leadership styles that enable them to manage the unexpected better than most other organizations. High-reliability organizations foster and maintain an attitude of mindfulness or "intelligent wariness." To be mindful is to have an awareness of detail and an enhanced ability to identify and prevent errors that could escalate into an adverse event. Desirable as it might be, acting more like a high-reliability organization is neither simple nor easy. It is difficult for individuals and crews to remain chronically wary about their operations. Moreover, several of the high-reliability organizations noted in the previous paragraph have had to create distinct supporting subcultures that often put them at odds with their parent organizations.

Prudent Use of Automation

The previous sections present a somewhat pessimistic picture of the situation with respect to the safe and effective use of automated weapons systems such as Patriot. That said, I am not opposed to the development and use of air and missile defense systems employing a high level of automation and capable of near-autonomous operations. There are situations in which a high level of automation and near-autonomous operations clearly are required. One such vsituation involves defending against large numbers of incoming ballistic missiles, what analysts refer to as a saturation attack. Human operators performing in-the-loop or too closely on-the-loop in such situations could be overwhelmed and not able to cope effectively with performance demands. Too closely on-the-loop refers to a situation in which operators under-trust the automation and do not permit the system the control latitude the engagement situation demands. This is the flip side of the automation over-trust issue mentioned previously. In a sense, this requirement led to the development of Patriot's automatic mode of operation more than 35 years ago. Recall that Patriot's automatic mode was adapted from the Safeguard system's automatic mode. That mode of operation was entirely appropriate for Safeguard's mission objectives and operating environment. Problems arose when the automatic mode was incorporated into Patriot without a critical consideration of differences between Patriot and Safeguard. That led to imprudent use of Patriot during OIF and contributed to the fratricide incidents.

As implied above, the key to the safe and effective use of highly automated and potentially near autonomous systems such as Patriot is prudent use. So, what does it mean to use a system such as Patriot in a prudent manner? The formal definition of prudent is to act with judiciousness and demonstrate care and thought for the future. This definition is consistent with the mindfulness or intelligent wariness exhibited by high-reliability organizations and with the DoD's requirement for positive human control.

With respect to automated systems such as Patriot, I think there are three fundamental requirements for such systems to be developed properly and employed prudently. First, users must accept the notion that such systems are fallible. Trust in the system's automation must be developed incrementally on the basis of experience, and will always be situation-specific. Crews must learn through experience when the system can be trusted, and when additional scrutiny and system oversight are necessary. The Army clearly violated this requirement with its pre-OIF stance that Patriot crews should trust the system without question.

Second, highly automated systems such as Patriot rely on a high level of user expertise for safe and effective use. Expertise is developed over time using a hands-on instructional regimen that presents trainees with "tough cases" that challenge and expand their skill level and depth of system understanding. Once again, the Army failed to do this prior to OIF. Training was too short given the system's technical and operational complexity. Pre-OIF Patriot training focused too much on getting crews certified to enter the unit's operational crew rotation and too little on corresponding skill development. Training also tended to focus on what the Army's own post-OIF board of inquiry criticized as emphasizing "rote drills" over critical thinking and problem solving. Operator and crew roles were defined and assessed in terms of rote procedural outcomes rather than the mindful exercise of positive control over engagements and lethal assets. Successful execution of rote drills and procedures became the de facto functional measures of training success and readiness to enter the unit's operational crew rotation. With inadequate, rote-drill-oriented training, the operators' performance capabilities are brittle in the same sense that the system's control algorithms are brittle. This situation has been partially corrected, but in my observation, the Army still has a long way to go with respect to training times, methods, and standards along with supporting personnel practices. Avoiding the rote-drills trap remains a challenge for future users of automated systems.

The third requirement for safe and prudent use concerns the way automated systems are developed. In a military setting, the traditional approach to system development can present an obstacle to the deployment of effective automated systems. This obstacle pertains to what might be termed the "irreversible waterfall" from requirements definition through to testing, deployment, and field use. The usual practice in DoD systems acquisition is to define detailed system requirements and specifications up-front and then proceed linearly to system development, developmental testing, operational testing, deployment, and field use. This series of events, once initiated, often becomes the irreversible waterfall mentioned above. Information flows in one direction only, regardless of the downstream consequences for the system, rather than in an iterative fashion where requirements and design solutions can evolve as the technology is developed. As things stand now, most substantive system evaluation is left until formal test events conducted immediately prior to mandated milestone review points and system development is nearly complete. By that time, most degrees of freedom for concept reevaluation or design changes have been lost. The system has, in effect, gone over the developmental waterfall. Program office metrics mostly concerned with schedule and funding dictate this lockstep approach to system acquisition. Rarely is there time or funding to go back and re-conceptualize, redesign, or retest a system. Consequently, development and fielding often go forward with too many loose ends, design rough edges, and unknowns. This problem was and continues to be true for Patriot, and I also observe it in the developmental programs for successor systems. Achieving effective human-automation integration is a tough technical challenge. But an even more daunting challenge to the development and successful use of automation in future military systems will be to modify traditional system development practices.

Conclusion

Over the past decade, there has been growing interest in the topic of automation and weapon system autonomy. Much of this interest is being driven by developments in computing power, software engineering, artificial intelligence, and similar technical and engineering disciplines. The most recent example of this growing interest is the DoD's "third offset" strategy. Third offset makes extensive reference to human-machine teaming and weapons system autonomy. There also have been many defense-related publications addressing the potential role of weapons system autonomy. Examples of these include the Defense Science Board's 2012 report on the role of autonomy in DoD systems, and *Autonomous Horizons,* published in 2015 by the Air Force Office of the Chief Scientist. In general, these publications are very well done. From a human factors perspective, *Autonomous Horizons* does a particularly good job of laying out the human performance challenges associated with achieving effective human-automation integration. However, I think these and related publications tend to downplay the difficulties associated with meeting those human performance challenges in operational systems. Policy makers could easily be lulled into a false set of expectations regarding the timing, necessary due diligence, and eventual operational reliability of automated and near-autonomous systems. That is a big part of what happened with Patriot. At some level, policy and plans must reflect the limits of technical and operational feasibility. The Patriot case study illustrates how difficult it is to resolve a number of the underlying issues involving human-automation integration and training operators and crews to perform satisfactorily in a supervisory control capacity. It also illustrates the potential negative consequences of doctrine and usage practices being inconsistent with technical realities.

The OIF Patriot fratricides and ARL's deep-dive assessment that followed provided a unique opportunity to examine the performance of a highly automated weapon system in a realistic operational environment. As the previous discussion suggests, the incidents observed during OIF are representative of the kinds of problems that can and will occur with such systems. To a great extent, the OIF fratricide incidents were symptoms of the underlying humans and automation problems discussed throughout this paper. It has often been observed that we in the human factors community know a lot about how a variety of factors (e.g., system design, use of automation, training, crew dynamics) make certain kinds of incidents and erroneous actions predictable. Our ability to predict the timing and number of these incidents and erroneous actions is very weak, but our ability to predict the kinds of errors that will occur is very good. As described previously, such was the case with Patriot. We also are pretty good at telling designers how to avoid those kinds of errors and incidents. Unfortunately, we are better at doing that after prototypes exist than while a system is being developed.

There are few hard and fast rules for achieving effective human-automation integration.

There are few hard and fast rules for achieving effective human-automation integration. Most such rules or design guidance are, in essence, design rules of thumb. Consequently, the degree to which acceptable human-automation integration has been achieved often must be determined empirically on a trial-and-error basis after system prototypes are available. Effective usability work of that kind requires real-time interactions with expert job performers, or as close as we can come to that. It is also true that current DoD system acquisition practices often make it difficult to conduct such usability work as the system is being developed. Developing effective automated systems is far more than simply a technical or engineering challenge. Human factors and organizational considerations such as those discussed herein are vitally important to the safe and effective use of automated and near-autonomous systems.

Endnotes

1. Robert O. Work, Remarks at CNAS National Security Forum, December 14, 2016, http://www.cnas.org/transcripts/work-remarks-national-security-forum.

2. For more on the controversy surrounding the Patriot's performance against Iraqi SCUDs during the Gulf War, see Patrick E. Tyler, "After the War; Did Patriot Missiles Work? Not So Well, Scientists Say," *New York Times,* April 17, 1991, http://www.nytimes.com/1991/04/17/world/after-the-war-did-patriot-missiles-work-not-so-well-scientists-say.html; Theodore A. Postol, "Optical Evidence Indicating Patriot High Miss Rates During the Gulf War," Statement to the House Committee on Government Operations, April 7 1992, https://fas.org/spp/starwars/congress/1992_h/h920407p.htm; Theodore A. Postol, "Postol/Lewis Review of Army's Study on Patriot Effectiveness," September 8, 1992, http://fas.org/spp/starwars/docops/pl920908.htm; Fred Kaplan "Patriot Games," *Slate,* March 24, 2003, http://www.slate.com/articles/news_and_politics/war_stories/2003/03/patriot_games.html.

3. J.K. Hawley, A.L. Mares, & C.A. Giammanco, *The Human Side of Automation: Lessons for Air Defense Command and Control* (ARL TR 3468), (Adelphi, MD: U.S. Army Research Laboratory, 2005).

4. For example, see J.M. Bradshaw, R.R. Hoffman, M. Johnson, and D.D. Woods, "The seven deadly myths of 'autonomous systems,'" *IEEE Intelligent Systems*, May/June 2013.

5. An overview of these results is provided in J.K. Hawley and A.L. Mares, "Human Performance Challenges for the Future Force: Lessons from Patriot after the Second Gulf War," in *Designing Soldier Systems: Current Issues in Human Factors*, ed. P. Savage-Knepshield (Burlington, VT: Ashgate, 2012), 3-34. Related topics are summarized in Hoffman, R.R., Hawley, J.K., Hawley, and J.R. Bradshaw, "Myths of Automation, Part 2: Some Very Human Consequences," IEEE Intelligent Systems, March/April 2014.

6. Section 226 of Pub. L. 100-108 provided, "No agency of the Federal Government may plan for, fund, or otherwise support the development of command and control systems for strategic defense in the boost or post-boost phase against ballistic missile threats that would permit such strategic defenses to initiate the directing of damaging or lethal fire except by affirmative human decision at an appropriate level of authority." For current DoD policy on autonomy in weapons, see DoD Directive 3000.09 (November 21, 2012). *Autonomy in Weapons Systems* (Washington, DC: U.S. Department of Defense).

7. These observations are summarized from material presented in Hawley, Mares, and Giammanco (2005).

8. D.A.Norman,. *The Design of Future Things* (New York: Basic Books, 2007), 113.

9. C. Perrow, *Normal Accidents: Living with High-Risk Technologies* (Princeton, NJ: Princeton University Press, 1999).

10. K.E. Weick and K.M. Sutcliffe, *Managing the Unexpected: Assuring High Performance in an Age of Complexity* (San Francisco, CA: Jossey-Bass, 2001).

About the Center for a New American Security

The mission of the Center for a New American Security (CNAS) is to develop strong, pragmatic and principled national security and defense policies. Building on the expertise and experience of its staff and advisors, CNAS engages policymakers, experts and the public with innovative, fact-based research, ideas and analysis to shape and elevate the national security debate. A key part of our mission is to inform and prepare the national security leaders of today and tomorrow.

CNAS is located in Washington, and was established in February 2007 by co-founders Kurt M. Campbell and Michèle A. Flournoy.

CNAS is a 501(c)3 tax-exempt nonprofit organization. Its research is independent and non-partisan. CNAS does not take institutional positions on policy issues. Accordingly, all views, positions, and conclusions expressed in this publication should be understood to be solely those of the authors.

1152 15th Street, NW Suite 950 Washington, DC 20005
t. **202.457.9400** | f. **202.457.9401** | **info@cnas.org** | **cnas.org**

Center for a
New American
Security

Bold. Innovative. Bipartisan.