



DEVCOM DAC-TR-2023-078
November 2023

Report on Human Factors Issues Likely to Affect Air-Launched Effects/Autonomous Loitering Effects (ALE)

by Jamison Hicks, Michael Lewis, and Katia Sycara

DISCLAIMER

The findings in this report are not to be construed as an official Department of the Army position unless so specified by other official documentation.

WARNING

Information and data contained in this document are based on the input available at the time of preparation.

TRADE NAMES

The use of trade names in this report does not constitute an official endorsement or approval of the use of such commercial hardware or software. The report may not be cited for purposes of advertisement.



DEVCOM DAC-TR-2023-078
November 2023

Report on Human Factors Issues Likely to Affect Air-Launched Effects/Autonomous Loitering Effects (ALE)

by Jamison Hicks
DEVCOM Analysis Center

Michael Lewis
University of Pittsburgh

Katia Sycara
Carnegie Mellon University

REPORT DOCUMENTATION PAGE

1. REPORT DATE		2. REPORT TYPE		3. DATES COVERED	
November 2023		Technical Report		START DATE 2/15/2022	END DATE 9/15/2023
4. TITLE AND SUBTITLE Report on Human Factors Issues Likely to Affect Air-Launched Effects/Autonomous Loitering Effects (ALE)					
5a. CONTRACT NUMBER		5b. GRANT NUMBER		5c. PROGRAM ELEMENT NUMBER	
5d. PROJECT NUMBER		5e. TASK NUMBER		5f. WORK UNIT NUMBER	
6. AUTHOR(S) Jamison Hicks, Michael Lewis, and Katia Sycara					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Jamison Hicks DEVCOM Analysis Center BLDG 4506 Joker St. Fort Novosel, AL 36362				8. PERFORMING ORGANIZATION REPORT NUMBER DEVCOM DAC-TR-2023-078	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report reviews human factors research on the supervision of multiple unmanned vehicles (UVs) as it affects human integration with Air-Launched Effects (ALE). Research finds that an operator cannot manually control more than 2–4 UVs even if assigning waypoints and controlling them sequentially. If UVs are allowed to coordinate autonomously and we let the operator control them as a group, this problem goes away but returns if the operator must monitor for targets. If automated target recognition (ATR) is introduced and the operator needs to approve or veto targets, then larger numbers of UVs can be supervised. However, now that the operator is no longer monitoring, the broader context is lost and the operator is unlikely to maintain situation awareness (SA). Displaying marked targets and UVs on a map can support SA; however, if the operator must also make direct use of multiple sensor feeds such as video, the changes in perspective among UVs are very difficult to reconcile. If views from multiple UVs are fused in a synthetic display eliminating the problem of multiple perspectives, resolution must be drastically reduced to fit on the screen. Now the operator must continuously zoom in and out to view the battlespace with sufficient resolution, which can also interfere with SA. Choices among these alternatives depend on the capabilities the Army desires. If “flooding the zone” is a desired strategy, then pairing a map display with ATR and operators who respond to the alarms independently could support it. If intelligence, surveillance, and reconnaissance is the objective, then pairing a fused display and map might work best. The purpose of this review is to point out these trade-offs so they can be made in an informed way.					
15. SUBJECT TERMS Army aviation, unmanned vehicle, Air-Launched Effects, unmanned control, automated target recognition					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED	UU		41
19a. NAME OF RESPONSIBLE PERSON Jamison Hicks				19b. PHONE NUMBER (Include area code) (334) 255-2206	

STANDARD FORM 298 (REV. 5/2020)

Prescribed by ANSI Std. Z39.18

(U) Table of Contents

List of Figures	iv
List of Tables	v
Executive Summary	vi
1. INTRODUCTION	1
1.1 Air-Launched Effects.....	1
1.2 Characterization of ALE	5
2. HUMAN FACTORS REVIEW OF MULTI-UV CONTROL	6
2.1 HRI Complexity in Multi-UV Control	6
2.2 Control of Quasi-independent UVs $O(n)$	7
2.3 $O(n)$ with Human Allocation of Attention	8
2.4 $O(n)$ with Automated Direction of Attention	10
2.5 $O(1)$ Automating Coordination Among UVs	12
2.5.1 Independent Path Planners.....	12
2.5.2 Cooperative Planners.....	13
2.5.3 Plan Libraries	13
2.5.4 Swarms	14
2.6 Multi-UV Display	15
2.6.1 Synthetic Video: Picture-in-Picture and Map Marking	18
2.6.2 Multi-UAV Fusion	19
3. CONCLUSIONS AND RECOMMENDATIONS.....	22
4. REFERENCES	24
List of Acronyms	31
Distribution List	33

List of Figures

Figure 1.	Raytheon's Coyote Block 3 (left) and Area-1's Altius 700 (right) are exemplars of the ALE to be developed	1
Figure 2.	The hypothesized relationship between task types and command complexity.....	7
Figure 3.	Examples of multi-UAV control using Falcon view. O(1) designating search area, O(n) setting waypoints, O(>n) timed rendezvous requiring multi-UV synchronization.	7
Figure 4.	Neglect tolerance cycle from Crandall et al. (2005)	8
Figure 5.	Canonical control interface with multiple video streams (left), in-focus stream (top right), minimap (bottom right).....	16
Figure 6.	Five common perspectives from Wickens and Hollands (1999). (©2000 from <i>Engineering Psychology and Human Performance</i> by Wickens, Christopher D., and Hollands, Justin G. Reproduced by permission of Taylor and Francis Group, LLC, a division of Informa plc.).....	17
Figure 7.	Synthetic video overlay with PiP (center, including landing gear) and labeled landmarks (right) from Calhoun et al. (2005).....	18
Figure 8.	Frame from video of Sarnoff/SRI Terrasight system circa 2012	20
Figure 9.	(Left) Cameras move across height map leaving textured areas on; (right) aerial imagery from the Gascola test site textured on a height map of the area.....	21

List of Tables

Table 1. W56KGU-20-R-0001 specifications for ALE 3

Executive Summary

This report reviews human factors research on the supervision of multiple unmanned vehicles (UVs) as it affects human integration with Air-Launched Effects (ALE). Research finds that an operator cannot manually control more than 2–4 UVs even if assigning waypoints and controlling them sequentially. If UVs are allowed to coordinate autonomously and we let the operator control them as a group, this problem goes away but returns if the operator must monitor for targets. If automated target recognition (ATR) is introduced and the operator needs to approve or veto targets, then larger numbers of UVs can be supervised. However, now that the operator is no longer monitoring, the broader context is lost and the operator is unlikely to maintain situation awareness (SA). Displaying marked targets and UVs on a map can support SA; however, if the operator must also make direct use of multiple sensor feeds such as video, the changes in perspective among UVs are very difficult to reconcile. If views from multiple UVs are fused in a synthetic display eliminating the problem of multiple perspectives, resolution must be drastically reduced to fit on the screen. Now the operator must continuously zoom in and out to view the battlespace with sufficient resolution, which can also interfere with SA. Choices among these alternatives depend on the capabilities the Army desires. If “flooding the zone*” is a desired strategy, then pairing a map display with ATR and operators who respond to the alarms independently could support it. If intelligence, surveillance, and reconnaissance is the objective, then pairing a fused display and map might work best. The purpose of this review is to point out these trade-offs so they can be made in an informed way.

* “Flooding the zone” refers to a strategy of filling and dominating a region of the field of play in a sport such as football.

1. INTRODUCTION

This report reviews human factors research on the supervision of multiple unmanned vehicles (UVs) as it affects human integration with Air-Launched Effects (ALE). Publicly available information on the ALE program is reviewed herein and used to identify human factors and related research relevant to human supervision of multiple UVs.

1.1 Air-Launched Effects

ALEs (Figure 1) are a concept for operating small, inexpensive, attritable, and highly autonomous unmanned aerial systems that can be tube launched from aircraft. Launch from ground vehicles is planned as well, although Ground-Launched Effects are not yet a requirement (Pomerleau, 2022). ALEs are envisioned to provide “reconnaissance, surveillance, target acquisition (RSTA), and lethality with an advanced team of manned and unmanned aircraft as part of an ecosystem including Future Attack and Reconnaissance Aircraft (FARA) and ALE” (W911W6-21-R-0014, 2020). A primary purpose of ALEs is to extend “tactical and operational reach and lethality of manned assets, allowing them to remain outside of the range of enemy sensors and weapon systems while delivering kinetic and non-kinetic, lethal and non-lethal mission effects against multiple threats, as well as, providing battle damage assessment data” (PEO Aviation, 2020).

Key features of the ALE are adherence to the Modular Open Systems Architecture so that components rely on nonproprietary interfaces and the Scalable Control Interface, an application programming interface intended to allow control through any command device adhering to the protocol, including both air and ground forces.




Figure 1. Raytheon’s Coyote Block 3 (left) and Area-1’s Altius 700 (right) are exemplars of the ALE to be developed

Over the past 3 years, in a series of experiments and demonstrations known as Experimentation Demonstration Gateway Event (EDGE) and Project Convergence (Convergence), ALE prototypes have demonstrated increasing sophistication and capability. At Convergence-20, six ALEs covered a range of 62 km relaying RSTA information back to manned assets at standoff distances. Four were air launched while two demonstrated a new ground launch capability (Judson, 2020). In EDGE-21, Altius 600 ALE prototypes were launched as a small swarm and controlled from a Blackhawk helicopter. By EDGE-22, four waves of seven ALEs each (28 total) were ground launched, demonstrating the “flood the zone*” capability envisioned by Brigadier General Rugen, head of the modernization team for Future Vertical Lift (FVL) (Parsons, 2022). “Most importantly, the ALE performing these behaviors were controlled by a single operator,” the Aviation and Missile Center said in a pair of tweets (Parsons, 2022). “12 months ago these same operations would have required 3-4 people, which is a testament to the software’s ability to greatly reduce the operator workload.” It is presumed that this refers to a seven-ALE cohort, which is a notable achievement.

In interviews, General Rugen has outlined the vision for the ALE program, including keeping manned aircraft at a standoff distance safe from attack, fully autonomous flight, including sense and avoid and path planning, and ALE as a service accessible through Tactical Assault Kits. “The aircraft knows how to fly,” Rugen said. “We don’t need to fly the aircraft [by remote control]; that’s what we can automate” (Freedberg, 2020). He would also like them to be numerous enough and expendable enough to flood the zone (Freedberg, 2020). Dustin Engelhardt, the requirements lead for ALE in the Army’s Aviation Platforms Requirements Determination Directorate, presents a complementary vision: “What we don’t want is what we typically see with UAS and manned systems right now: where it’s a steady stream of high-definition video, it’s very data-heavy, network-heavy and requires a lot of input from the operator,” Engelhardt says. Instead, the ALE will “execute a series of preprogrammed and dynamic instructions and report back as needed to the FARA” (Trimble, 2020).

Public reports of the expected capabilities of ALEs vary widely, so we will rely on the ranges laid out in W56KGU-20-R-0001 (2020), the Air Launch Effects Request for Information (RFI). Table 1 and subsequent bulleted lists are taken directly from the document dropping items not relevant to human-ALE interaction.

* “Flooding the zone” refers to a strategy of filling and dominating a region of the field of play in a sport such as football.

Table 1. W56KGU-20-R-0001 specifications for ALE

Air frame	ALE small	ALE small	ALE large	ALE large
	Objective	Threshold	Objective	Threshold
Metric				
Max total ALE Platform Weight (lb)	50	100	175	225
Combat Range (km)	150	100	650	350
Min. Operational Airspeed (KTAS)	0	30	0	70
Min. Dash Airspeed (KTAS)	205	120	600	350
Mission payload endurance (h)	1.00	0.50	1	0.5
Total flight time (h)	1.50	0.70	2	1.3

Envisioned operations especially relevant to human interaction from W56KGU-20-R-0001 (2020, pp. 9–10) and re-labeled for future reference are as follows:

- OP-1: Self-organizes through Artificial Intelligence (AI) enabled teaming to maximize coverage of assigned mission area and tasks.
- OP-2: Provisioned with the ability to transfer control from the launching platform to another supported element, via permissions-based scalable control interface, ensuring ALE capabilities are maximized across all domains.
- OP-3 DILR [Detection, Identification, Location, and Reporting]: Passively collects and shares information on threat disposition to provide situational awareness and enable development of Common Operating Picture (COP).
- OP-4 DILR: Passive and active sense & geolocate to provide targetable locations to ALE swarm & ownship.
- OP-5 DILR: Conduct battle damage assessment (BDA) and report to ownship.

Sensor capabilities reported in W56KGU-20-R-0001 especially relevant to human interaction include:

2 Passive EO/IR DILR

2.1 Payloads should consist of an EO/IR [electro-optical/infrared] sensor unit, Inertial Measurement Unit (IMU), and Aided Target Recognition (AiTR) processor.

2.2.2 Aided Target Cueing/Detection (AiTD) for automatic cue to operator indicating presence of potential targets of a selected target size and/or profile within a scene. Post processes near real time. Allows operator to adjust false alarm rate by selecting target detection probability.

2.2.3 Aided Target Recognition of Detected Targets (AiTR) for automatic aid to operator to support recognition of target-like objects detected. Performing on scanned and/or point-and-stare imagery generating potential target locations at increased recognition resolution. Allows operator to adjust false alarm rate by selecting target detection probability.

2.3.1 Sensor payloads to function in day, night and Degraded Visual Environments (DVE).

2.3.3 Output of payloads capable of selecting either Full Motion Video (FMV) or single image target pictures to include target coordinates.

5 Decision-Aiding/Making Software/Algorithms

5.2 Decision aiding are decisions that may be sent upstream to aid in execution of mission level tasks or support situation awareness for the common operating picture (i.e., maneuvering for optimized effect). Decision making are decisions made local to the system to change its own behavior or execute a payload function (i.e., perform disrupt techniques X then Y).

5.3 Decisions will be sensor and effector focused and lead to behaviors that drive execution of an effect or execution of an action needed to create an effect, including:

- Algorithms that autonomously characterizes known/unknown threat emitters
- Algorithms that autonomously determine the best technique or combination of techniques to decoy/disrupt the threat, including # of agents, location of agents and time on target
- Algorithms that monitor countermeasure effectiveness to further optimize response

5.4 Algorithms should enable employment with a single agent or across a team of agents in a distributed environment (i.e., Multi-platform sensor data correlation to make more informed decisions, or coordination across multiple platforms for cooperative engagement)

5.5 Decisions may be based on fixed rules for pre-optimized behaviors, or may employ adaptive reasoning per changing environment/states.

5.6 State changes include but not limited to the following and can occur within the Blue Force or Red Force:

Change in mission/function (New target resulting in change from Decoy to Disrupt)

Change in techniques (from Masking to Attack)

Change in position (to improve technique effectiveness or compensate for degradation, etc)

Change in threat mode and/or threat position

Change in # of ALEs (attrition or acquisition of agents)

5.7 Decisions should take into account platform geometries and operational constraints, for example roll rate, turn radius, pitch angle, etc.

6 Command, Control, Communications (C3)

6.3.5 Radio System Data Links- Provide link capability for ALEs to operate as a single asset or as part of a team of ALEs or a Swarm to execute the same mission or different missions. Include type of communications system(s) used (including LOS, BLOS, point-point, etc.) and protocol(s).

1.2 Characterization of ALE

From these descriptions, ALE has at least two distinct mission profiles and probably many others. One is as an asset controlled by and acting on behalf of a manned platform. This requires relatively few (two to five) ALEs that can be tasked individually or as a team. There is leeway in how they may be controlled and how their data may be consumed, but as General Rugen indicated, they are not to be flown in any literal sense beyond providing waypoints or heading to an autopilot (Parsons, 2022). The second profile involves swarms of ALEs (seven to...?). These will need to be controlled as a group (as argued in the next section) but must also incorporate decision making to perform local actions, such as joint attacks or battle damage assessment by diverting small numbers of ALEs from the ongoing group activity. Human factors topics in control of multiple UVs involving the use of glyphs (multidimensional icon displays) to convey system status or Gantt chart displays for scheduling have been excluded because it is presumed that scheduling will be handled by automation, and the short flight duration and likelihood of attrition make health monitoring less important. In fact, Fabijanowicz (2020) finds in a Joint Dynamic Allocation of Fires and Sensors study that only an average of 10% of ALEs can be expected to survive.

In the Human–Robot Interaction (HRI) review that follows, a range of issues likely to affect human–machine integration in deployment of ALE is examined. As the Edge and Convergence experiments and demonstrations show, some of these capabilities are already at Technical Readiness Level 7. These primarily relate to control and tasking of ALEs and relay of targeting information from individual platforms, which our analyses suggest should be readily achievable. Other capabilities, such as maintaining a common operating picture and situation awareness (SA), may be more difficult.

To help tie this review back to ALE, capabilities and characteristics from W56KGU-20-R-0001 (2020) Air Launch Effects RFI will be referenced with square brackets, for example, [2.2.2 Target Cueing].

2. HUMAN FACTORS REVIEW OF MULTI-UV CONTROL

2.1 HRI Complexity in Multi-UV Control

In computer science, the notion of computational complexity—the time that must be used to solve a problem as a function of the size of its input—has proved fruitful for weeding out bad algorithms. High-complexity algorithms may work for small problems but fail or grow inefficient for even slightly larger ones. The task of controlling multiple UVs is similar to an algorithm in that the operator must perform sequences of decisions and actions. If the UVs are performing independent activities, the operator can devote the same attention to each in turn, resulting in a complexity of Order n , written $O(n)$, because each of the “ n ” UVs requires the same set of actions, and the total operator effort, therefore, will be proportional to the number of UVs.

Another benefit of independence is that more UVs can be controlled simply by adding more operators because the UVs and their operators will not interfere with one another. A different form of control, such as designating a region to be searched by drawing it on a map, can command an arbitrary number of UVs with a single act. Because the number of actions the operator must take is independent of the number of UVs, control of this sort is $O(1)$ and has a constant effort. Dependent tasks such as box pushing, by contrast, can be arbitrarily difficult with command complexity, $O(>n)$, because dependencies among UVs create cascading demands. When one robot pushes one corner of a box, for example, the operator must control the other robot to push the other corner to straighten its path, after which the first robot needs attention again ad infinitum. Wang and Lewis (2007) demonstrate that operators controlling heterogeneous box-pushing unmanned ground vehicles (UGVs) by waypoints become rapidly saturated with no time left for additional tasks, a result replicated by Kaminka and Elmaliach (2006a; 2006b). In agreement with earlier studies (Parker, 1994), these experiments reconfirm that direct human control over tightly coordinating UVs does not scale beyond two to three and should be avoided for any larger number of UVs.

$O(1)$ tasks require fully autonomous coordination on the part of the UVs but impose only a constant demand on the human operator. In general, $O(1)$ control is appropriate where a large number of UVs must be tightly coordinated with a relatively simple goal, such as formation following or area search or a pre-computed plan to be executed. $O(n)$ tasks, such as approving targets or identifying victims, are UV-centric tasks that can be performed independently by one or more operators and impose a predictable additive demand. $O(>n)$ tasks, by contrast, cannot be specified simply and, depending on the task, could require arbitrarily large control on the part of the operator. Figure 2 illustrates these hypothesized relations. Figure 3 provides examples of control at each level of complexity. Given the complexity of tasks envisioned for future multi-UV teams, it is

unlikely that they can be effectively controlled at any single level of complexity. Instead, we envision a future in which control over UVs will at various times, and sometimes simultaneously, be performed by different people at different levels of complexity. While levels of complexity are not identical to levels of autonomy (LOA), they characterize the degree to which coordination among UVs has been automated and whether or not a human-in-the-loop system can be scaled to a larger number of UVs.




Figure 2. The hypothesized relationship between task types and command complexity



Figure 3. Examples of multi-UAV control using Falcon view. O(1) designating search area, O(n) setting waypoints, O(>n) timed rendezvous requiring multi-UV synchronization.

2.2 Control of Quasi-independent UVs O(n)

O(n) tasks are plentiful and likely to account for a substantial portion of an operator's necessary interactions. Foraging tasks, such as search and rescue or reconnaissance, where UVs are relatively sparse and thus unlikely to interfere with one another or employ automated path planning, form a broad class of applications that can be treated as O(n). Tasks involving individual UV status, such as monitoring for problems or failures or to make adjustments, are also O(n) because they can be performed without reference to other UVs. These tasks, in which the human controls the allocation of attention to UVs, can be modeled by the Neglect Tolerance model, proposed by Crandall and Goodrich (2001), which describes a round-robin control scheme in which each UV is serviced in turn.

A third class of $O(n)$ tasks that are among the most crucial performed by a UV operator are those where the UV requests human judgment or decisions, such as target identification or weapons release. In this case, human capacity to control UVs is determined by the arrival rate of requests rather than Neglect Time, the time a UV can be left unattended. Taken together, foraging, status, and judgment tasks are likely to significantly contribute to the control of multi-UV teams. It is, however, crucial to point out that treating tasks, whether control or monitoring, as $O(n)$ presumes they are quasi-independent; therefore, any intelligence or actions they provide are strictly local and only loosely correlated (in the head of the operator) with other UVs.

2.3 $O(n)$ with Human Allocation of Attention

According to the Neglect Tolerance model, an operator's interaction with a UV can be described as a sequence of control episodes in which the operator interacts with the UV for a period of time (interaction time [IT]), raising its performance above some upper threshold after which the UV is neglected for a period of time (neglect time [NT]) until its performance deteriorates below a lower threshold when the operator must again interact with it. The duration of the IT/NT intervals depends on the task requirements, which determine appropriate performance thresholds. Lowering performance thresholds shortens IT and lengthens NT, which also increases with automation. This model is very general and can describe a range of human-multi-UV systems where the human must monitor or control. Olsen's fan-out model (Goodrich & Olsen 2003; Olsen & Wood, 2004; Crandall et al., 2005) fixes these parameters to give an upper bound on the number of independent homogeneous UVs that a single human can manage. This upper bound is computed by identifying how long an individual UV can be neglected, and then determining how many other UVs could receive interactions during this neglect interval. Thus, span of control is given by $FO = \frac{NT}{IT} + 1$. This fan-out model of span of control has been validated for teleoperated UGVs (Olsen & Wood, 2004) and waypoint-following unmanned aerial vehicles (UAVs) (Cooper & Goodrich, 2008). The fan-out model is illustrated in Figure 4.




Figure 4. Neglect tolerance cycle from Crandall et al. (2005)

One evident oversimplification of the model is the assumption of fixed IT and NT. Goodrich et al. (2009) demonstrate that for even relatively small variations in IT/NT duration, operators may need to be idle up to 90% of the time if unserved UVs are to be avoided. Another mismatch between the fan-out model and the tasks performed by operators in the supervisory control of multi-UV systems lies in the other competing things the operator may do during what is ostensibly neglect time available for interacting with the next UV. Cummings et al. (2007) propose an augmented model, taking into account wait time (WT),

$FO = \frac{NT}{IT+WT} + 1$, where WT is summed across UVs. They attribute wait time to three causes:

- WTI: operator time spent in planning before taking action,
- WTQ: time spent by UVs in degraded states waiting in a queue for operator attention, and
- WTSA: time spent before operator notices degraded state.

The model was validated by Donmez et al. (2010) by categorizing operating intervals containing user interactions (IT and WTI), intervals during which a UV was in a degraded condition and another being serviced (WTQ), and intervals during which one or more UVs persisted in a degraded state without operator intervention (WTSA).

How many UVs an operator can manage in these ways has been a recurring theme and subject of many studies. Trouvain and Wolf (2002) and Trouvain et al. (2003), for example, set the number at four to five but observed that human intervention led to worse performance than automation alone. Hancock et al. (2007) asked, “Is the UAV control ratio the right question?” and came up with the same one-UV answer. Valero et al. (2010) found a drop-off between two and three, as did Lewis et al. (2010) between four and eight. Other researchers, such as Humphrey et al. (2007), have found a break between six and nine. Where this number may actually fall is not intrinsic and ultimately depends on what aspects of the task have been automated and to what degree; see Lewis (2013) for a detailed discussion of the effects of LOA on the controllable number of UVs.

For the use case of ALEs controlled from a helicopter, it is probably safe to assume that a single unmanned aerial system (UAS) could be manually controlled, two to three ALEs could be controlled using waypoints or adjusting headings (Valero et al., 2010), and four to eight could be monitored if path planning is automated (Lewis et al., 2010). It should be noted, however, that while increasing the number of UVs beyond one led to improvements in measures such as targets found and area searched, there were diminishing returns with targets missed and inefficiencies in trajectories increasing with

the number of UVs. While the RFI did not call for this capability, it would seem unwise to preclude operators from using ALEs in this way, such as manually monitoring video to search for out-of-distribution targets or directing ALEs to new headings. In cases where UASs must be redirected to pop-up targets or as a backup to failed or disrupted automation, manual navigation should be acceptable for one to three ALEs with monitoring for three to six.

2.4 O(n) with Automated Direction of Attention

Automated direction of attention is a mainstay of conventional automation regimes in applications ranging from nuclear power plants to aircraft cockpits to hospital Intensive Care Units. Rather than rely on constant human monitoring of systems that may be too complex for operators to understand in their entirety, alarms provide a cue to direct attention to particular subsystems or processes requiring operator interaction. In the case of ALEs, AiTR can alert the operator to initiate action in response to detected targets [2.2.2 Target Cueing]. Alarms, however, are vulnerable to common causes; e.g., an overloaded annunciator system in a nuclear power plant following an off-normal event or a long queue of detected targets awaiting servicing when ALEs overfly a densely populated position.

The effects of automation on task performance depend jointly on what aspects of a task are automated and how they are automated. These distinctions have commonly been organized as levels, stages, or degrees of automation. The oldest of these, Tom Sheridan's (Sheridan & Verplank, 1978) LOAs, focuses on locus of control (human 1–5, shared 6, or automation 7–10) and proceeds from aspects of the task involving information to those dealing with actions. A revision of this taxonomy, Stages of Automation (Parasuraman et al., 2000), emphasizes the progression from information to action by postulating four stages, any of which can be automated:

- Information acquisition: example-providing band-pass filter to minimize clutter and enhance targets on a display
- Information analysis (example): detecting and drawing boxes around potential targets
- Decision selection (example): recommending attack on target matching known signature
- Action implementation (example) execute decision to attack or clear

Guidance based on this model (Lee & See, 2004; Onnasch et al., 2014; Parasuraman et al., 2000), later termed degree of automation, favors higher degrees of automation at early stages and lower degrees of automation at later ones, largely in the interest of keeping the human engaged and aware of developing situations. As the number of UVs

increases, this may no longer be practical, although, as the model suggests, it is quite feasible; for example, automating navigation across the four stages while limiting the automation of monitoring to information analysis, such as drawing boxes around potential targets [2.2.2 Target Cueing]. For ALEs, automation could occur at any of these stages, but the most likely stages are information analysis, decision selection, and action implementation. If only information acquisition/analysis is automated, the human may be required to monitor, making it a human-driven task.

Most research in LOA (Sheridan & Verplank, 1978) for multi-UV control (Goodrich et al., 2001), however, has focused on the two intermediate levels: management by consent (MBC), in which the UV proposes an action for human approval, and management by exception (MBE), in which the human must intervene to veto a machine-chosen action. In the canonical case, a UV plans its trajectory [5.5 Adaptive Reasoning], detects a potential target [2.2.2 Target Cueing], recognizes the signature [2.2.3 Aided Recognition], seeks human approval at this point or selects an action [5.5 Fixed Rule], then seeks human approval. While support for this sequence is not explicit in the RFI, the support in [2.2.3 Aided Recognition] for human inspection of higher-resolution imagery for potentially recognized targets implies such an intermediary role. In MBE, by contrast, the human must stop a programmed execution.

Under these regimes, the human only interacts with the ALE when interaction is needed, so a round-robin sequence of IT and NT periods is no longer necessary. The interaction is instead limited by the arrival rate of targets—although, as noted earlier, service rates could be extended by adding additional operators.

In a study comparing control of one, two, and four UAVs in a suppression of enemy air defenses (SEAD) task, Ruff et al. (2002) found MBC led to better performance in destroying targets, fewer hits, and higher SA than MBE. Their experiment used a relatively novel form of aiding in which assistance was provided for 10 distinct situations the UAVs might encounter. Cummings and Mitchell (2005, 2006) studied control of one, two, four, and six UAVs again at a SEAD task. In this study, the initial path planning was fully automated with operators intervening during execution to replan aspects of the mission in response to unexpected events. In the two highest LOAs tested in this experiment, MBE produced the best performance and highest SA while MBC was the poorest. Calhoun et al. (2009a) compared three LOAs, the highest of which corresponded to MBE, at a three-UAV rerouting task finding; in this case, MBE led to longer rerouting times than MBC. Another line of research has investigated the joint effects of LOA and automation reliability on multi-UV performance. Ruff et al. (2002) varied the fidelity of automation decisions and found that MBC led to better rejection of bad advice than MBE. Adjustable autonomy (Crandall & Goodrich, 2001), where the operator chooses the level at which to interact with the system, has frequently improved

effectiveness typically by allowing the operator to interleave manual commands with control at higher LOAs (Squire et al., 2006; Valero-Gomez et al., 2011). Studies by Calhoun et al. (2009b) and Calhoun et al. (2011) compared sequential tasks of the same or varying LOAs in control of three UAVs finding that matching LOAs led to better performance than mixed LOAs. As this brief review suggests, MBC is generally preferred, although there may be no best LOA for all multi-UV tasks.

The scaling characteristics of $O(n)$ interactions are favorable to deploying large ALE teams, providing tasks can be broken down in quasi-independent ways into local engagements. So, under Gen. Rugen's flood the zone strategy, if the object is to detect and shoot targets with human oversight to veto civilian casualties, this becomes a scheduling problem of finding the operators needed to meet the demand. This can be extended in a limited way to local multi-ALE actions [5.4 Cooperative Engagement] such as engaging additional ALEs for BDA [OP-5].

2.5 O(1) Automating Coordination Among UVs

One consequence of the conjecture that coordination of UVs is generally too complex for a human operator to handle is that operators will need to manage the coordination of UVs some other way. Using planners, programs that generate plans to satisfy user specified criteria, is a widely adopted solution. We distinguish between Independent path planners, which provide simple standalone path planning, Plan libraries that generate complex, multiphase plans from stored templates and Cooperative planners that construct plans to satisfy an objective function under operator guidance.

Independent planners are most useful when UVs are relatively independent, performing tasks such as foraging based on simple criteria such as moving to the frontier while avoiding other UVs. Plan libraries are useful where complex, variable, interdependent behavior is needed such as missions in which multiple UVs assume different roles at different times such as search, attack, and BDA. Cooperative planners are useful where intent can be expressed relatively succinctly in an objective function and a near-optimal solution is desired.

2.5.1 Independent Path Planners

Independent path planners are usually designed so that the human operator cannot directly alter constraint or objective function parameters, so the planner must be run as a black box. Because the operator lacks the ability to influence the plan that is produced, such systems are usually fully automated or practice some form of MBC or MBE. A common strategy for interacting with independent planners is to allow the operator to adapt the generated plan (Calhoun et al., 2017; Chen & Barnes, 2012;

Velagapudi et al., 2009) often by moving planner-generated waypoints (see Section 2.5.2, Cooperative Planners).

For foraging (search) tasks, which involve both directing UV movement and monitoring for targets, automating path planning has been shown to produce greater improvements in performance than aiding or automating the monitoring task (Chien et al., 2011; Goodrich et al., 2007; Humphrey et al., 2007; Lewis et al., 2010). In a review of multi-UAV studies, Cummings et al. (2007) conclude that automated path planning appears necessary for applications in which operators supervise more than four UAVs [5.7 Navigation].

2.5.2 Cooperative Planners

Operators and planners each have limited visibility. The planner has access to large amounts of sensor and other data inaccessible to the human, while the human has knowledge of factors unrepresented within the planner. Fully automated path planning finds near-optimal solutions but does so largely without user influence and therefore is unable to benefit from an operator's greater visibility or more precise knowledge of intent. In cooperative planning, the operator can influence the contents, not merely the choice of plan, by steering the planning process through altering either the constraints or the objective function. After a plan is produced, the operator may be allowed to alter some aspect of the plan, such as moving a waypoint (Chen et al., 2010; Chien et al., 2010), or the alteration can be passed back to the planner for incorporation in a new plan (Cummings et al., 2012). Constraints may be expressed to a planner either through constraining the planning space or delimiting an admissible space within which to plan. In Lewis et al. (2006), for example, wide-area search munitions were directed to search some areas, avoid others, and jettison into a third by drawing associated regions onto a map presented through FalconView (Figure 3). Cooperative control for ALE in the RFI appears limited to adjusting target detection [2.2.2 Target Cueing] and recognition [2.2.3 Target Recognition]. Under some conditions, allowing human input to influence trajectories has been found particularly useful. The "noodle" (Calhoun et al., 2017), in which the operator manually traces a short path for the planner to follow through regions with hazards the planner has not been programmed to recognize, adds needed flexibility.

2.5.3 Plan Libraries

Planning multi-UV paths to meet geometric objectives, such as coverage or rendezvous, is a complex problem but simple to define mathematically. Other tasks that involve phased actions, such as ingress, search, and egress or triggered patterns with role-based coordination, such as target detection, role recruitment, and BDA, are less

easy to define in terms of an objective function to be minimized. Of all the possible actions that might be taken, however, such “plans” can be specified through a small number of coordinated actions in a stereotyped sequence. If the appropriate coordination plan can be instantiated and executed, then coordination can be automated. The success of this approach relies on the match between the situations for which plans have been prepared and those which are encountered. For complex and frequently used coordinated actions, such as overwatch, attack and BDA, etc., the ability to execute stereotyped coordinated behavior is essential [5.4 Multiple Platforms, 5.5 Fixed Rules].

The most well-developed of these approaches is Playbook (Miller and Parasuraman 2007), a family of systems that integrate a hierarchical task network planner. Recent experiments using a realistic three-UAV simulation (Fern & Shively, 2009; Shaw et al., 2010) compared three levels of delegation (waypoint, UAV specific delegation, team delegation) and found that the highest level (team delegation) and flexibility (in choice of level) lead to the best performance. Shaw et al. (2010) also examined effects of failures requiring manual control and found that operators responded more rapidly in the “plays” condition, suggesting that the technique did not foster complacency or automation bias. Versions of this system underwent usability testing at the Air Force Research Laboratory (Calhoun et al., 2012; Miller et al., 2012) to improve design and incorporate new features.

2.5.4 Swarms

Swarms (Kolling et al., 2015) are distinct from conventional multi-UV systems in that they use a high level of autonomy, inspired from natural systems or control theory, to control the local interactions of their members, thus removing the need for a human or other centralized actor to coordinate interactions. The advantages of this form of coordination are its robustness to loss of members or other disruptions and its ability to scale to very large numbers of UVs. Swarm behaviors arise from interactions among local control laws and, whether described as physico-memetic or biologically inspired, rely on consensus algorithms that adjust a member’s parameters to match that of its neighbors. This gives rise to behaviors such as flocking (moving as a group), rendezvousing (converging within a region), dispersing (moving away from one another), or formation following. Swarm behavior can be influenced by changing the variables in the control algorithms either as a hybrid system switching between control laws (Kolling et al., 2013) or through influence over selected leaders (Goodrich et al., 2011).

Swarm algorithms are particularly well suited for coordinating large numbers of ALEs. As coordination emerges [OP-1 Self-organizes], ALEs are commanded and coordinated

as a group [5.4 Multiple Platforms], and group behavior is unaffected by attrition or addition of ALEs [5.6 Change in #]. $O(1)$ control as a swarm [5.4 Multiple Platforms] can be paired with $O(n)$ responding to detected targets [2.2.2 Target Cueing, 2.2.3 Target Recognition] with plan library for local multi-ALE response [5.4 Multiple Platforms, 5.5 Fixed Rules]. Chandarana et al. (2021) introduces a model of this process with mission planning for a swarm, in which subswarms break off to prosecute identified targets, after which they reintegrate with the main swarm.

A small number of ALEs controlled from a helicopter could be coordinated either directly $O(n)$ if their tasks are relatively independent or through $O(1)$ automated coordination. Control using a swarm algorithm [OP-1 Self-organizes, 5.4 Multiple Platform, 5.6 Change in #] would probably be most natural as the operator could command them as a group [5.4 Multiple Platform] through control over heading and velocity. Where greater autonomy is desired, as in algorithms seeking to maximize information gain or to favor predetermined areas of interest [5.3 Autonomously Determine], it might be desirable to incorporate plan libraries [5.5 Fixed Rules] for expected tasks of greater complexity than moving between regions or searching independently.

2.6 Multi-UV Display

While automating navigation and directing human attention through alarms, depending on the arrival rate of requests, control of relatively large numbers (10–20?) of UVs could be allowed; this control is of a peculiar sort. The operator could specify $O(1)$ heading, waypoints, etc., for the team and approve or veto $O(n)$ targets as they are discovered, but the system does nothing to support a more coherent understanding of the battlespace. A common strategy to support better SA is to pair a minimap (small map that is not the primary display) showing the locations of the UVs with a second window displaying electro-optical and other data from the UVs. The operator now must cross-index the locations and orientations of UVs with their data streams in order to generate SA. Figure 5 shows a multi-window display of this sort, in which the operator searches to find victims in thumbnails displaying video streams from 24 moving UGVs. After identifying a victim, the operator must mark its location on the minimap. This display is from an experiment comparing different ways of sharing control among operators (Lewis et al., 2011); however, even in the simpler single operator case, performance is poorer (fewer victims marked) with eight UVs than with four (Lewis et al., 2010).




Figure 5. Canonical control interface with multiple video streams (left), in-focus stream (top right), minimap (bottom right)

The Immersed view from a UV (Figure 6) is simply the data stream it returns and provides information available at the location. The Fixed Plan view shows UV location on a map and allows integrated display of multiple UVs but does not convey any additional information. The Rotating Plan view shows UV locations on a map rotated to align with a reference heading as done in some GPS modes. Map views of these sorts provide an excellent and potentially manipulable display for command and control of UV teams but fail to organize collected EO/IR data in a readily intelligible form and are therefore typically paired with displays of sensor feeds. The Tethered views (think Mario Brothers) try to find a compromise by following the actor around, providing contextual information usually by including some part of the actor in the scene to help the viewer infer orientation and motion.




Figure 6. Five common perspectives from Wickens and Hollands (1999). (©2000 from *Engineering Psychology and Human Performance* by Wickens, Christopher D., and Hollands, Justin G. Reproduced by permission of Taylor and Francis Group, LLC, a division of Informaplc.)

A human is frequently included in the loop of a multi-UV system expressly to monitor and interpret electro-optical (such as video) and other data being gathered by UVs. This can be a difficult task with even a single camera. Gugerty and Brooks (2001) investigated the ability of Air Force recruits to make cardinal direction judgments from aerial imagery. Participants were presented with a north-up map with an icon of an aircraft and a line indicating camera heading. The aerial image contained a building with parking lots on four sides, only one of which had cars. The participants' task was to indicate whether this parking lot was on the north, south, east, or west side of the building. At some camera headings, accuracy was well below 50%. With multiple UVs, this problem is multiplied because each view must be reconciled with the multiple frames of reference.

There are three basic problems for exploiting data returned from multiple sensors or UVs: information fusion to integrate data from multiple viewpoints, providing context to maintain SA, and filtering to prevent operator overload. Synchronous displays from multiple UVs can easily overwhelm an operator who must search multiple video streams

for targets even in response to alarms. If targets are plentiful, the operator will likely miss targets that enter and leave unattended views while occupied with targets in attended views. The related fusion and context problems arise because UVs' multiple fields of view are not aligned and may overlap, forcing the operator to reconcile views from different perspectives by mentally "piecing them together." Taken together, the problems of multisensor monitoring and fusion of viewpoints can pose insurmountable difficulties to operators who must develop and maintain SA through remote sensing (Woods et al., 2004).

2.6.1 Synthetic Video: Picture-in-Picture and Map Marking

In picture-in-picture (PiP) displays (a name taken from the TV format), UV video is projected onto a digital terrain map of the same area providing context through visibility of landmarks and look ahead, assisting in things such as setting out-of-view waypoints. Where imagery is left in place, as the UAS's frustum moves on, the image is said to be mosaicked. In a series of studies, Calhoun et al. (2006) and Draper et al. (2006) investigated the use of synthetic video overlays to supply context to UAV video. Both studies found advantages for marking landmarks in overlays, but the initial Calhoun et al. (2006) experiment found no advantage for the PiP technique. In the follow-on studies shown in Figure 7 (Calhoun et al., 2007; Draper et al., 2006), a terrain elevation overlay was substituted, and PiP led to substantial improvements in performance.




Figure 7. Synthetic video overlay with PiP (center, including landing gear) and labeled landmarks (right) from Calhoun et al. (2005)

Other advanced work in synthesizing displays from digital terrain with single UAV video is found in Cooper (2007). Using five distinct perspectives made possible through synthesized video, the display merges the functions of map, video display, and controller; this configuration supports both in-context interpretation of the video feed and precise control of the UAV, including features such as magnified UAV field of view projected to the side of the flight path and control of roll by manipulating wing angle on screen. In a separate experiment, Drury et al. (2006) found differences in target-marking accuracy favoring PiP and strong participant preference for the contextual display. Morse et al. (2007) demonstrates that local mosaicking, texturing several recent filtered frames of video onto a height map, led to a 45% improvement in target detection with a third of those coming from “historical” sightings after the current view had moved on.

2.6.2 Multi-UAV Fusion

Alberts et al. (1999) proposed the concept of Network-Centric Warfare, which argued for shifting from a Platform Centric view, in which information is gathered and presented from a platform’s perspective, to a Network Centric view, in which all available information from whatever source is channeled to a platform’s displays, a sort of everyman’s common operating picture. From the perspective of an operator supervising a team of UAVs, this would appear to be the sort of display needed to make sense of incoming data streams while the conventional command and control map could continue to serve as a display for directing action.

Around this same time, researchers from Sarnoff Corporation were developing the Aerial Video Surveillance system (Kumar et al., 2001) which, in addition to doing a lot of signal processing from jitter removal to edge enhancement, georeferenced data so that UAV imagery could be integrated back at the base station to synthesize imagery from multiple UAVs. Demonstrations of the Sarnoff technology conducted at Fort Huachuca are described by Hunn (2005, 2006) as fusing video feeds from four UAVs, digital terrain databases, and conventional navigation charts. From Hunn’s description and accompanying screenshots, a 3D view was used to show UAVs and their fields of view projected onto an elevation map (Figure 8). This 3D display appears to have been used as a stand-in for a conventional command-and-control map and used in the same way by operators to interpret selected camera feeds displayed in separate windows.




Figure 8. Frame from video of Sarnoff/SRI Terrasight system circa 2012

Over the next two decades, research using geolocation and image stitching proliferated in the computer vision community creating UAV-generated maps (Li et al., 2014; Ruiz et al., 2018), object tracking (Ramachandran & Sangaiah, 2021), persistent surveillance (Meng et al., 2015), and so on. Rapidly improving processing and sensing hardware such as Go-Pro cameras, GPS, and accelerometers from cell phones have made these capabilities so widely available that a fully autonomous UAS can now be assembled from commercial off-the-shelf (COTS) components with little or no modification (Sun et al., 2016). Because the usefulness of fused displays was so “obvious,” the only thing lacking in the open literature has been human experimentation to examine presentations of these data, how they can be manipulated, and how they can feed back into controlling UAS missions, Abedin et al. (2011) conducted preliminary studies in this area in 2010–11. The SUAVE system created textures from video imagery and pasted them onto a height map of the terrain. As wider areas are covered, the map is continually updated reflecting the areas of the environment that have been viewed.

To gain information from the UV feeds, the user inspects the map. Although the operator’s interaction with this display is asynchronous, the map contains all information that would be currently available from synchronous video, albeit at a lower frame rate. This display is demonstrably $O(1)$ because adding more UAVs without altering the area being viewed will either increase the spatial resolution of the map or increase the temporal resolution of the map due to more frequent updates without affecting the operator’s task. By moving a “window” containing the frusta of multiple UVs over a map, a multi-perspective analog of a tethered display is created allowing the operator to

maintain perspective on both the multiple FOVs and their context. An additional feature of this form of display is that bandwidth can be throttled back to very low rates, such as single frames, without undue disruption as the sliding window can be moved by dead reckoning between updates.

Merging current and historical imagery has both advantages and disadvantages. Providing context and more recent views of a region than available from a terrain database is more “accurate,” but once the UAV moves on, the usefulness of the imagery begins decaying as changes become increasingly likely. Strategies, such as progressively fading imagery as it ages, could convey the growing uncertainty, but at the same time, would deny the user access to the best available information about these regions. Alternatives such as spatio-temporal filtering (Wang, 2013) would allow the user to specify a duration over which imagery would remain at full resolution, search some interval in the past to view what was happening across the FOVs of the team, or sift through video from a region over a longer period (Morse, 2011) to reinterpret past events (Kelley et al., 2016).

Resolution and expanse of imagery is another problem. While conventional video displays provide full resolution for area surveilled along with the ability to zoom in on smaller areas, a multi-UAV display must fit multiple FOVs into the same limited space, necessitating frequent zooming in and out to obtain views with as high a resolution as available from streaming video to fixed windows. As Figures 8 and 9 illustrate, this is a very real problem, and appeals to favorable scaling characteristics cannot obviate the fact that we are trying to display a much larger area on a screen of the same size.




Figure 9. (Left) Cameras move across height map leaving textured areas on; (right) aerial imagery from the Gascola test site textured on a height map of the area

3. CONCLUSIONS AND RECOMMENDATIONS

As the Army moves toward increasingly complex Multi-Domain Operations, the ability to effectively deploy large numbers of unmanned systems becomes crucial. This report analyses human–autonomy interaction through the prism of scalability to better understand how unmanned systems might be deployed and supervised at scales needed to impact modern warfare. Our analysis finds three broad categories of control: autonomous coordination in which UVs are controlled as a team, quasi-independent control in which UVs are controlled largely independently of one another, and coordinated control in which operator(s) must directly coordinate UVs. We conclude that direct coordination is infeasible for realistic applications, leaving automated coordination and quasi-independent control as feasible options.

Quasi-independent control, in which operator(s) direct one or more UVs at largely independent tasks, such as area search with minimally overlapping patterns, offers many advantages in that multiple UVs can be controlled in sequence and scaled beyond a single operator's limits, simply by adding additional operators. Depending on the arrival rates of events requiring operator attention such as targets, UVs might be controlled out of call centers, taking advantage of load balancing to increase the UV-to-operator ratio. This option also has the advantage of relying largely on COTS drone technology with map and first person displays for which there is already a body of research and experience requiring only protection from electronic warfare to be fielded. Applications such as massed attacks to flood the zone are also possible using quasi-independent control. The only proviso is that since sensing and control are UV centered, a general operating picture or situation awareness is difficult for operators to acquire. Though “low tech,” quasi-independent control has many advantages, including robustness to attrition and resilience inherent in possible human takeover under unanticipated circumstances. Many of the planned advanced ALE capabilities including AiTR, cueing and recognition, pre-optimized behavior, and passive sharing of information are fully compatible with quasi-independent control. As the simplest scheme for deploying ALEs, quasi-independent control should be preferred for applications for which it is feasible.

Automated coordination can take varied forms, each of which present different human control issues. Simple forms of coordination, such as consensus algorithms (swarms), are fully scalable, immune to attrition, but limited in behavior. Rule or optimization-based coordination is limited in scale and may be brittle but can produce complex coordinated behaviors. A useful way of thinking about this problem is to use pre-optimized behaviors or “plays” for recurring locally coordinated behaviors, such as Battle Damage Assessment, while coordinating larger groups through scalable methods, such as consensus algorithms. Because the operator is freed from managing UV trajectories,

the primary human factors problem posed by autonomously coordinating UVs is dealing with the data streams they return. Automating trajectories allows refinements (for moderately sized teams 5–20) such as coordinating to maximize coverage, or information gain, or covering locations specified during intelligence preparation of the battlefield. In these cases, it may be very difficult for the operator to piece together multiple views from trajectories they do not control or necessarily understand into a coherent picture. Possible human factors solutions to integrating multiple data feeds are discussed in the report but remain an open problem. As Intelligence, Surveillance, and Reconnaissance is a primary envisioned use for ALE teams, integration of multiple data feeds should be a human factors research priority.

4. REFERENCES

- Abedin, S., Wang, H., Lee, P., Lewis, M., Brooks, N., Owens, S., Scerri, P., & Sycara, K. (2011). SUAVE: Integrating UAV video using a 3D model. *2011 IEEE International Conference on Systems, Man, and Cybernetics*, 2865–2869.
- Alberts, D., Garstka, J., & Stein, F. (1999). Network centric warfare: Developing and leveraging information superiority. Command and Control Research Program.
- Calhoun, G., Draper, M., Abernath, M., Delgadoc, F., & Patzek, M. (2005). *Synthetic vision system for improving unmanned aerial vehicle operator situation awareness* (Report No. AFRL/WS-05-0423). Air Force Research Laboratory.
- Calhoun, G. L., Draper, M. H., Nelson, J. T., & Ruff, H. A. (2006). Advanced display concepts for UAV sensor operations: Landmark cues and picture-in-picture. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(1), 121–125.
- Calhoun, G., Draper, M., & Ruff, H. (2009a). Effect of level of automation on unmanned aerial vehicle routing task. *Proceedings of the Human Factors and Ergonomics Society 53rd Annual Meeting*, 197–201.
- Calhoun, G., Draper, M., Ruff, H., Barry, T., Miller, C., & Hamell, J. (2012). Future unmanned aerial systems control: Feedback on a highly flexible operator-automation delegation interface concept. Infotech@Aerospace 2012, AIAA, June 19–21.
- Calhoun, G., Goodrich, M., Dougherty, J., & Adams, J. (2017). Human-autonomy collaboration and coordination toward multi-RPA missions. In Cooke, Rowe, Bennet & Jorlmon (Eds.), *Remotely piloted aircraft systems: A human systems integration perspective* (pp. 109–136). John Wiley & Sons, Ltd.
- Calhoun, G. L., Ruff, H. A., Draper, M. H., & Wright, E. J. (2011). Automation-level transference effects in simulated multiple unmanned aerial vehicle control. *Journal of Cognitive Engineering and Decision Making*, 5(1), 55–82.
- Calhoun, G., Ruff, H., Draper, M., Wright, N., & Mullins, B. (2009b, April). Levels of automation in multi-UAV control allocation and router tasks. Paper presented at AIAA Infotech@Aerospace, Seattle, WA.
- Calhoun, G., Ruff, H., Lefebvre, A., Draper, M., & Ayala, A. (2007). Picture-in-picture augmentation of UAV workstation video display. *Proceedings of the 51st Annual Meeting of the Human Factors and Ergonomics Society*, 70–74.
- Chandarana, M., Hughes, D., Lewis, M., Sycara, K., & Scherer, S. (2021). Planning and monitoring multi-job type swarm search and service missions. *Journal of Intelligent & Robotic Systems*, 101(3).

-
-
- Chen, J. Y., & Barnes, M. J. (2012). Supervisory control of multiple robots effects of imperfect automation and individual differences. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 54(2), 157–174.
- Chen, J., Barnes, M., & Qu, Z. (2010). RoboLeader: An agent for supervisory control of multiple robots. *Proceedings of ACM/IEEE International Conference on Human-Robot Interaction (HRI'10)*, 81–82.
- Chien, S., Mehrotra, S., Brooks, N., Sycara, K., & Lewis, M. (2011) Effects of alarms on control of robot teams. *Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11)*, 55(1), 434–438.
- Chien, S., Wang, H., & Lewis, M. (2010). Human vs. algorithmic path planning for search and rescue by robot teams. *Proceedings of the 54th Annual Meeting of the Human Factors and Ergonomics Society (HFES'10)*, 379–383.
- Cooper, J., & Goodrich, M. (2008). Towards combining UAV and sensor operator roles in UAV-enabled visual search. *Proceedings of ACM/IEEE International Conference on Human-Robot Interaction*, 351–358.
- Crandall, J., & Goodrich, M. (2001). Experiments in adjustable autonomy. *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, 3, 1624–1629.
- Crandall, J., Goodrich, M., Olsen, D., & Nielsen, C. (2005). Validating human-robot interaction schemes in multi-tasking environments. *IEEE Transactions on Systems, Man, and Cybernetics —Part A: Systems and Humans*, 35(4), 438–449.
- Cummings, M., How, J., Whitten, A., & Toupet, O. (2012). The impact of human-automation collaboration in decentralized multiple unmanned vehicle control. *Proceedings of the IEEE*, 100(3), 660–671.
- Cummings, M. L., & Mitchell, P. J. (2005). Managing multiple UAVs through a timeline display. *Proceedings of AIAA Info Tech*.
- Cummings, M. L., & Mitchell, P. (2006). Automated scheduling decision support for supervisory control of multiple UAVs. *Journal of Aerospace Computing, Information, and Communication*, 3(6), 294–308.
- Cummings, M., Nehme, C., Crandall, J., & Mitchell, P. (2007). Predicting operator capacity for supervisory control of multiple UAVs. *Studies in Computational Intelligence (SCI)* 70, 11–37.
- Donmez, B., Nehme, C., & Cummings, M. L. (2010). Modeling workload impact in multiple unmanned vehicle supervisory control. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 40(6), 1180–1190.
-
-

-
-
- Draper, M., Calhoun, G. Nelson, J., Lefebvre, A., & Ruff, H. (2006). Synthetic vision overlay concepts for uninhabited aerial vehicle operations: Evaluation of update rate on four operator tasks. *Paper presented at the NATO-HTM-135 Symposium on Human Factors of Uninhabited Military Vehicles as Force Multiplier.*
- Drury, J. Richer, J., Racklie, N., & Goodrich, M. (2006). Comparing situation awareness for two unmanned aerial vehicle human interface approaches. *Proceedings of the IEEE International Workshop on Safety, Security and Rescue Robotics.*
- Fabijanowicz, J. (2020). *Design of experiments for air launched effects unmanned aerial vehicles* [Master's thesis, Naval Postgraduate School].
- Fern, L., & Shively, R. J. (2009). A comparison of varying levels of automation on the supervisory control of multiple UASs. *Proceedings of AUVSI's Unmanned Systems.*
- Freedberg, Sydney. (2020, September 14). Kill chain in the sky with data army's project convergence. Breaking Defense, <https://breakingdefense.com/2020/09/kill-chain-in-the-sky-with-data-armys-project-convergence/>.
- Goodrich, M. A., McLain, T. W., Anderson, J. D., Sun, J., & Crandall, J. W. (2007). Managing autonomy in robot teams: Observations from four experiments. *Proceedings of the ACM/IEEE international conference on Human-robot interaction, 25–32.*
- Goodrich, M., Morse, B., Engh, Cooper, J., & Adams, J. (2009). Towards using UAVs in wilderness search and rescue: Lessons from field trials, interaction studies. *Journal, special issue on "Robots in the Wild: Exploring Human-Robot Interaction in Naturalistic Environments", 10(3), 453–478.*
- Goodrich, M. A., & Olsen Jr, D. R. (2003, October). Seven principles of efficient human robot interaction. *IEEE International Conference on Systems, Man and Cybernetics, 4, 3942–3948.*
- Goodrich, M. A., Olsen, D. R., Crandall, J. W., & Palmer, T. J. (2001). Experiments in adjustable autonomy. *Proceedings of IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents* (pp. 1624–1629). International Joint Conference on Artificial Intelligence.
- Goodrich, M., Sujit, P., Pendleton, B., & Pinto, J. (2011). Toward human interaction with bio-inspired robot teams. *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics* (pp. 2859–2864).
- Gugerty, L., & Brooks, J. (2001). Seeing where you are heading: Integrating environmental and egocentric reference frames in cardinal direction judgments. *Journal of Experimental Psychology: Applied, 7(3), 251–66.*
-
-

-
-
- Hancock, P. A., Mouloua, M., Gilson, R., Szalma, J., & Oron-Gilad, T. (2007). Provocation: Is the UAV control ratio the right question? *Ergonomics in Design*, 15(1), 7.
- Humphrey, C., Henk, C., Sewell, G., Williams, B., & Adams, J. (2007). Assessing the scalability of a multiple robot interface. *Proceedings of ACM/IEEE 2007 International Conference on Human-Robot Interaction, HRI'07*, 239–246.
- Hunn, B. (2005). The Human challenges of command and control with multiple unmanned aerial vehicles. *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*, 49(1), 20–24.
- Hunn, B. (2006). Video imagery's role in network centric, multiple unmanned aerial vehicle (UAV) operations. In Cooke, Pringle, Pedersen, & Connor (Eds.), *Human factors of remotely operated vehicles* (pp. 179–192). Elsevier.
- Judson, J. (2020, August 24). Army awards air launched effects contracts for future helicopters. *Defense News*, <https://www.defensenews.com/land/2020/08/24/army-awards-air-launched-effects-contracts-for-future-helicopters/>
- Kaminka, G. A., & Elmaliach, Y. (2006a, May). Experiments with an ecological interface for monitoring tightly coordinated robot teams. *Proceedings 2006 IEEE International Conference on Robotics and Automation, ICRA*, 200–205.
- Kaminka, G. A., & Elmaliach, Y. (2006b). Single operator, multiple robots: Call-request handling in tight-coordination tasks. *Distributed Autonomous Robotic Systems 7*, 103–113.
- Kelley, S., Kagel, J., Poyant, J., & Plourde, J. (2016, April 27). Eye in the sky, radio lab. <https://www.wnycstudios.org/podcasts/notetoself/episodes/conspiracy-theorist-radiolab-surveillance>.
- Kolling, A., Sycara, K., Nunnally, S., & Lewis, M. (2013). Human swarm interaction: An experimental study of two types of interaction with foraging swarms. *Journal of Human-Robot Interaction*, 2(2), 103–128.
- Kolling, A., Walker, P., Chakraborty, N., Sycara, K., & Lewis, M. (2015). Human Interaction with robot swarms: A survey. in human-machine systems (HMS). *IEEE Transactions on Human-Machine Systems* (pp. 9–26). IEEE.
- Kumar, R., Sawhney, H., Samarasekera, S. Hsu, S. Tao, H., Guo, Y., Hanna, K., Pope, A., Wildes, R., Hirvonen, D., Hansen, M., & Burt, P. (2001). Aerial video surveillance and exploitation. In *Proceedings of the IEEE*, 89(10), 1518–1539.
- Lee, J. & See, K. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors Journal, Human Factors and Ergonomics Society*, 46(1), 50–80.
-
-

-
-
- Lewis, M. (2013). Human interaction with multiple remote robots. In: D. Kaber (Ed), Human Factors Reviews Volume 9, on Human Performance in Teleoperation and Beyond, HFES, 131–174.
- Lewis, M., Polvichai, J., Sycara, K., & Scerri, P. (2006). *Scaling-up human control for large UAV teams*. In N. Cooke (Ed.), The Human Factors of Remotely Piloted Vehicles, Elsevier, pp. 237–250.
- Lewis, M., Wang, H., Chien, S., Velagapudi, P., Scerri, P., & Sycara, K. (2010). Choosing autonomy modes for multirobot search. *Human Factors*, 52(2), 225–233.
- Lewis, M., Wang, H., Chien, S., Ma, Z., Velagapudi, P., Scerri, P., & Sycara, K. (2011). Process and performance in human-robot teams. *Journal of Cognitive Engineering and Decision Making*, 5(2), 186–208.
- Li, J., Yang, T., Yu, J., Lu, Z., Lu, P., Jia, X., & Chen, W. (2014). Fast aerial video stitching. *International Journal of Advanced Robotic Systems*, 11(10).
- Meng, X., Wang, W., & Ben Leong, B. (2015). SkyStitch: A cooperative multi-UAV-based real-time video surveillance system with stitching. In *Proceedings of the 23rd ACM international conference on Multimedia (MM '15)*. Association for Computing Machinery, 261–270.
- Miller, C., Hamell, J., Barry, T., Ruff, H., Draper, M., & Calhoun, G. (2012). Adaptable operator-automation interface for future unmanned aerial systems control: Development of a highly flexible delegation concept demonstration. *Proceedings of AIAA Infotech*, 1507–1527.
- Miller, C., & Parasuraman, R. (2007). Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors*, 49(1), 57–75.
- Morse, B. S. et al., Application and evaluation of spatiotemporal enhancement of live aerial video using temporally local mosaics. 2008 IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, AK, USA, 2008, pp. 1–8
- Olsen, D., & Wood, S. (2004). Fan-out: Measuring human control of multiple robots. *Human Factors in Computing Systems (CHI'04)*, 231–238.
- Onnasch, L., Wickens, C., Li, H., & Manzey, D. (2014). Human performance consequences of stages and levels of automation: An integrated meta-analysis. *Human Factors*, 56(3), 476–488.
- Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 30(3), 286–297.
-
-

-
-
- Parker, L. (1994). Heterogeneous multi-robot cooperation. dissertation, Massachusetts Institute of Technology.
- Parsons, D. (2022, May 16). 28 drone swarm just paved way for simulated Army air assault mission. The Drive, <https://www.thedrive.com/the-war-zone/28-drone-swarm-just-paved-way-for-simulated-army-air-assault-mission>.
- PEO Aviation. (2020, August 21). Air launched effects. https://www.army.mil/article/238407/air_launched_effects_al.
- Pomerleau, M. (2022, September 6). Army wants to develop EW payloads launched from the air and potentially the ground. Fedscoop, <https://fedscoop.com/army-wants-to-develop-ew-payloads-launched-from-the-air-and-potentially-the-ground/>.
- Ramachandran, A., & Sangaiah, A. K. (2021). A review on object detection in unmanned aerial vehicle surveillance. *International Journal of Cognitive Computing in Engineering*, 2, 215–228.
- Ruiz, J., Caballero, F., & Merino, L. (2018). MGRAPH: A multigraph homography method to generate incremental mosaics in real-time from UAV Swarms. *IEEE Robotics and Automation Letters*, 3(4), 2838–2845.
- Ruff, H. A., Narayanan, S., & Draper, M. H. (2002). Human interaction with levels of automation and decision-aid fidelity in the supervisory control of multiple simulated unmanned air vehicles. *Presence: Teleoperators & Virtual Environments*, 11(4), 335–351.
- Shaw, T., Emfield, A., Garcia, A., de Visser, E., Miller, C., Parasuraman, R., & Fern, L. (2010). *Proceedings of the Human Factors and Ergonomics Society 54th Annual Meeting*, 1498–1502.
- Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators*. Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.
- Squire, P., Trafton, G., & Parasuraman, R. (2006). Human control of multiple unmanned vehicles: Effects of interface type on execution and task switching times. *ACM Conference on Human Robot Interaction*, 26–32.
- Sun, J., Li, B., Jiang, Y., & Wen, C-Y. (2016). A camera-based target detection and positioning UAV system for search and rescue (SAR) purposes. *Sensors*, 16(11), 1778.
- Trimble, S. (2020, October 7). Scout US Army plots future air launched effects. <https://aviationweek.com/shows-events/ausa-2020/scout-scout-us-army-plots-future-air-launched-effects>.
-
-

-
-
- Trouvain, B., Schlick, C., & Mevert, M. (2003). Comparison of a map-vs. camera-based user interface in a multi-robot navigation task. *IEEE International Conference on Systems, Man and Cybernetics*, 4, 3224–3231.
- Trouvain, B., & Wolf, H. L. (2002). Evaluation of multi-robot control and monitoring performance. *Proceedings of the 11th IEEE International Workshop on Robot and Human Interactive Communication*, 111–116.
- Valero-Gomez, A., de la Puente, P., & Hernando, M. (2011). Impact of two adjustable-autonomy models on the scalability of single-human/multiple-robot teams for exploration missions. *Human Factors*, 53(6), 703–716.
- Valero, A., Saracini, C., de la Puente, P., Rodriguez-Losada, D., & Matia, F. (2010). Exploratory analysis of operator: Robot ratio in search and rescue missions. *Proceedings of the 2nd Symposium New Frontiers in Human-Robot Interaction*.
- Velagapudi, P., Wang, H., Scerri, P., Lewis, M., & Sycara, K. (2009). Scaling effects for streaming video vs. static panorama in multirobot search. *2009 International Conference on Intelligent Robots and Systems (IROS'09), IEEE*, 5874–5879.
- Wang, H. (2013). Asynchronous visualization of spatiotemporal information for multiple moving targets. Dissertation, University of Pittsburgh.
- Wang, J., & Lewis, M. (2007 March). Human control for cooperating robot teams. *2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 9–16.
- W56KGU-20-R-0001. (2020, December 18). Army Combat Capabilities Development Command (CCDC) command, control, communications, computers, cyber, Intelligence, Surveillance, and Reconnaissance (C5ISR) Center, Air Launch Effects (ALE) Request for Information (RFI).
- W911W6-21-R-0014. (2020, August 13). Army Combat Capabilities Development Command (CCDC) command, control, communications, computers, cyber. Intelligence, Surveillance, and Reconnaissance (C5ISR) Center, Air Launch Effects (ALE) Request for Information (RFI).
- Wickens, C., & Hollands, J. (1999). *Engineering psychology and human performance*. Prentice Hall.
- Woods, D., Tittle, J., Feil, M., & Roesler, A. (2004). Envisioning human–robot coordination in future operations. *IEEE Transactions on Systems, Man, and Cybernetics: Part C*, 34(2), 210–218.

LIST OF ACRONYMS

3D	three dimensional
AI	Artificial Intelligence
AiTD	Aided Target Cueing/Detection
AiTR	Aided Target Recognition
ALE	Air-Launched Effect/Autonomous Loitering Effect
ATR	automated target recognition
BDA	battle damage assessment
COP	Common Operating Picture
DILR	Detection, Identification, Location, and Reporting
DVE	Degraded Visual Environment
EDGE	Experimentation Demonstration Gateway Event
EO/IR	electro-optical/infrared
EW	electronic warfare
FARA	Future Attack Reconnaissance Aircraft
FMV	Full Motion Video
FVL	Future Vertical Lift
GPS	global positioning system
HRI	Human–Robot Interaction
IT	interaction time
KTAS	knots per true air speed
LOA	level of autonomy
MBC	management by consent
MBE	management by exception
NT	neglect time
PiP	picture in picture
RSTA	reconnaissance, surveillance, and target acquisition
RFI	Request for Information
SA	situation awareness
SEAD	Suppression of Enemy Air Defense
UAS	unmanned aerial system
UAV	unmanned aerial vehicle

UGV	unmanned ground vehicle
UV	unmanned vehicle
WT	wait time

DISTRIBUTION LIST

DEVCOM Analysis Center
FCDD-DAH-A/J. Hicks
Daleville, AL

University of Pittsburgh
M. Lewis
4200 Fifth Ave.
Pittsburgh, PA 15260

Carnegie Mellon University
K. Sycara
5000 Forbes Ave.
Pittsburgh, PA 15213

DEVCOM Analysis Center
FCDD-DAD-TP/E. Chatterton
Redstone Arsenal
Huntsville, AL 35898

DEVCOM Army Research Laboratory
FCDD-RLB-CI/Tech Library
2800 Powder Mill Rd.
Adelphi, MD 20783

Defense Technical Information Center
ATTN: DTIC-O
8725 John J. Kingman Rd.
Fort Belvoir, VA 22060-6218