NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

OBJECT RECOGNITION IN SUPPORT OF SOF OPERATIONS

by

William L. Clark

June 2021

Thesis Advisor:	Alex Bordetsky
Co-Advisor:	Steven J. Mullins
Second Reader:	Douglas A. Borer

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE	*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2021	**3. REPORT TYPE AND DATES COVERED** Master’s thesis

4. TITLE AND SUBTITLE OBJECT RECOGNITION IN SUPPORT OF SOF OPERATIONS	**5. FUNDING NUMBERS**
6. AUTHOR(S) William L. Clark	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000	**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) SOCOM S&T, MacDill Air Force Base, Florida 33621; Office of Naval Research, Arlington, VA 22217	**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.	**12b. DISTRIBUTION CODE** A

13. ABSTRACT (maximum 200 words)

Current and future operational environments will increasingly require Special Operation Forces (SOF) to be more self-sufficient while operating in contested and politically sensitive regions where situational awareness can be degraded. This project continues Semi-Autonomous Threat Learning Alert System (SATLAS) efforts to integrate artificial intelligence-enabled small unmanned aerial systems into SOF teams to increase situational awareness and survivability. Specifically, we focus on directing prototype development and evaluating the ability of object recognition software to detect and categorize trained entities including weapons, personnel, and vehicles. Collaborating with commercial industries, we conduct simulation and field experiments to measure the ability of the Surveillance, Persistent Observation and Targeting Recognition (SPOTR) object recognition software to meet the technical requirements of the SATLAS project and operational requirements of SOF teams. We evaluate SPOTR based on accuracy, number of entities detected, and range of detection and recommend methods to improve its performance and meet our determined operational requirements. We advance the SATLAS project and set conditions for subsequent student teams to continue these efforts.

14. SUBJECT TERMS SATLAS, machine learning, artificial intelligence, drone	**15. NUMBER OF PAGES** 101
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	UU

NSN 7540-01-280-5500 Standard Form 298 (Rev. 2-89) Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

OBJECT RECOGNITION IN SUPPORT OF SOF OPERATIONS

William L. Clark
Major, United States Army
BA, North Carolina State University, 2009

Submitted in partial fulfillment of the
requirements for the degree of

**MASTER OF SCIENCE IN DEFENSE ANALYSIS
(ASTRONAUTICS)**

from the

**NAVAL POSTGRADUATE SCHOOL
June 2021**

Approved by: Alex Bordetsky
Advisor

Steven J. Mullins
Co-Advisor

Douglas A. Borer
Second Reader

Douglas A. Borer
Chair, Department of Defense Analysis

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Current and future operational environments will increasingly require Special Operation Forces (SOF) to be more self-sufficient while operating in contested and politically sensitive regions where situational awareness can be degraded. This project continues Semi-Autonomous Threat Learning Alert System (SATLAS) efforts to integrate artificial intelligence-enabled small unmanned aerial systems into SOF teams to increase situational awareness and survivability. Specifically, we focus on directing prototype development and evaluating the ability of object recognition software to detect and categorize trained entities including weapons, personnel, and vehicles. Collaborating with commercial industries, we conduct simulation and field experiments to measure the ability of the Surveillance, Persistent Observation and Targeting Recognition (SPOTR) object recognition software to meet the technical requirements of the SATLAS project and operational requirements of SOF teams. We evaluate SPOTR based on accuracy, number of entities detected, and range of detection and recommend methods to improve its performance and meet our determined operational requirements. We advance the SATLAS project and set conditions for subsequent student teams to continue these efforts.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

AGL	Above Ground Level
AI	Artificial Intelligence
ARSOF	Army Special Operations Forces
ATAK	Android Tactical Assault Kit
COTS	Commercial Off the Shelf
DARPA	Defense Advanced Research Projects Agency
DCLT	Detect, Categorize, Localize, and Track
DL	Deep Learning
DNN	Deep Neural Network
DOD	Department of Defense
FMV	Full Motion Video
HEO	Hyper Enabled Operator
H-GCS	Handheld Ground Control Station
ISIS	Islamic State of Iraq and Syria
ISR	Intelligence, Surveillance, Reconnaissance
MDO	Multi-Domain Operations
ML	Machine Learning
PC	Probability of Categorization
PD	Probability of Detection
PSI	Parameter Space Investigation
RPUAS	Rucksack Portable Unmanned Aerial System
SATLAS	Semi-Autonomous Threat Learning Alert System
SFOD-A	Special Forces Operational Detachment - Alpha
SOCOM	Special Operations Command
SOF	Special Operations Forces
SPOTR	Surveillance, Persistent Observation and Targeting Recognition
SUAS	Small Unmanned Aerial System
UAS	Unmanned Aerial System
VTOL	Vertical Take Off and Landing
YOLO	You Only Look Once

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

This project would not have been possible without the support of several individuals and organizations. First and foremost, I would like to thank my wife, Laura, for supporting me through this journey. You are the bedrock of our family and your flexibility and sacrifice are the only reasons this project was completed.

I would also like to thank Steve Mullins. Your advice, perspective, and willingness to sacrifice so many of your red pens are truly responsible for this project. Thank you for giving shape to this project. You are an incredible asset to the SATLAS team and the Naval Postgraduate School.

Additionally, I would like to thank Joe Kehoe and Ha Duong at AeroVironment for providing the technical knowledge and experiment support necessary to move this project in the right direction. Without your help, I might still be looking for experiment ideas.

Finally, I would like to thank the Office of Naval Research and the U.S. Special Operation's Command Science and Technology Directorate for funding the research and prototype development in our quest to try and save lives.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

The demand for U.S. forces, specifically Special Operations Forces (SOF), to address asymmetric and non-state actors is expected to increase over the next decades. Multidomain operations (MDO) anticipates that SOF will be required are required to conduct expeditionary operations independent of responsive support systems and in increasingly complex environments where adversaries and terrain can threaten operational security and restrict freedom of maneuver. The 2018 *National Defense Strategy* (NDS) and the Army's MDO doctrine project that SOF will need to operate in contested environments where all types of support will be more difficult compared to Iraq and Afghanistan,[1] where the U.S. has enjoyed air supremacy, mobility, and technically superior equipment. Army Special Operations doctrine clearly defines the necessity of "operations requiring unique modes of employment, tactical techniques, equipment, and training often conducted in hostile, denied, or politically sensitive environments…and characterized by a high degree of risk."[2]

In such conditions, SOF teams will need to be more self-sufficient in all aspects of their extended operations from intelligence to fires to logistics. The strategic environment highlighted in the *National Defense Strategy* warns that it may be infeasible to operate large unmanned aerial systems (UAS) to support small units in the presence of peer or near peer adversaries in politically sensitive or deep areas.[3] In that context, large, expensive, and theater-level UAS cannot be relied upon to provide dedicated intelligence, surveillance, and reconnaissance (ISR), which will be essential when operating in these

[1] Department of the Army, *The U.S. Army in Multi-Domain Operations 2028,* U.S. Army TRADOC Pamphlet 525-3-1 (Washington, DC: Department of the Army, 2018), 17–21, December 6, 2018, https://api.army.mil/e2/c/downloads/2021/02/26/b45372c1/20181206-tp525-3-1-the-us-army-in-mdo-2028-final.pdf

[2] Department of the Army, *Army Special Operations,* ADP 3–05 (Washington, DC: Department of the Army, 2019), ix, https://armypubs.army.mil/epubs/DR_pubs/DR_a/pdf/web/ARN18909_ADP%203-05%20C1%20FINAL%20WEB(2).pdf

[3] Department of Defense, *Summary of the 2018 National Defense Strategy of the United States of America.* (Washington, DC: Department of Defense 2018. National Defense Strategy), 2–3, Department of Defense, 2018. https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf

environments. There is an urgent need for real-time, organic ISR capability at the team level.[4]

Nonetheless, many SOF units continue to be equipped with small unmanned aerial systems (sUAS) such as the RQ-11 Raven, that possess outdated technological capabilities, require a dedicated operator and a large, fixed position from which to operate. Those systems are already inadequate, but no replacement sUAS with state-of-the-art capabilities has been fielded. We propose that advanced sUAS platforms could enhance situational awareness and survivability at the small unit level. Fortunately, military and corporate research continue to make significant advances in developing drone technology and machine learning (ML) algorithms for the Department of Defense (DOD) that promise to simultaneously enhance the autonomy of small units and reduce the operator's cognitive load.

The warfighter gap that we address in this research is the lack of a sUAS platform that meets the emerging needs at the SOF team level. Many current efforts to enhance UAS with Artificial Intelligence (AI)/ML are focusing on large theater level assets, such as the MQ-9 Reaper.[5] We opine that this prioritization, intentional or not, incurs unnecessary risk by not also adequately focusing on adopting some of the abundant, affordable, and man-packable commercial-off-the-shelf (COTS) sUAS to meet small unit requirements.

A. BACKGROUND

The first phase of the Semi-Autonomous Threat Learning Alert System (SATLAS) project was undertaken by Midgett et. al. from June 2019 to December 2020 and explored this concept through quantitative analysis and by identifying key capability requirements.[6]

[4] Department of the Army, *Army Futures Command Concept for Special Operations 2028,* Army Futures Command Pamphlet 71-20-4: (Washington, DC: Department of the Army, 2020), 18, https://api.army.mil/e2/c/downloads/2021/01/05/bdd61c44/20200918-afc-pam-71-20-4-afc-concept-for-special-operations-2028-final.pdf

[5] David Hambling, "U.S. To Equip MQ-9 Reaper Drones with Artificial Intelligence," *Forbes*, (December 2020), https://www.forbes.com/sites/davidhambling/2020/12/11/new-project-will-give-us-mq-9-reaper-drones-artificial-intelligence/?sh=37723b3c7a8e.

[6] Midgett et. al., "Semiautonomous Threat Learning Alert System," (master's thesis, Naval Postgraduate School, 2020), 46–52, December 2020.

They acquired sponsor funding from the Office of Naval Research and SOCOM's Science and Technology department and conducted market research to ascertain the current state of development of COTS and government sUAS and relevant software. Then they identified four priority capability requirements critical to improving situational awareness and decision-making for SOF teams: 1) a versatile, man-packable sUAS platform whose manufacturer would permit the team to integrate selected software functions, 2) AI-enabled object detection capable of on-the-edge processing, 3) autonomous tasking capability, and 4) a common ground control user interface.

They observed that the Defense Innovation Unit's "Blue sUAS" program is working to introduce a man-packable sUAS platform program of record; however, fielding had yet to begin.[7] The Defense Advanced Research Projects Agency has tested the use of object detection software in sUAS and in robots at the squad level and found that it can significantly increase survivability and decision-making, but no efforts towards implementation have been made.[8] Autonomous tasking remains in its early stages of development within the commercial sector[9]; finally, the Android Tactical Assault Kit (ATAK) has proven capable of meeting the user interface demands of SOF.[10] The objective of the ongoing SATLAS research project is to integrate these four pillars into a versatile ISR prototype for SOF teams.

The purpose of this thesis is to continue a longitudinal research project that seeks to design and develop a prototype deep learning (DL) enabled, semi-autonomous sUAS, and to evaluate its feasibility to support small units. It continues advancement towards a proof of concept that integrates hardware and software solutions to enhance SOF team's

---

[7] Department of Defense*, Defense Innovation Unit Announces sUAS Product Availability to Provide Secure, Capable Small Unmanned Aerial Systems for Critical Uses Across the Government*, Washington, DC: Department of Defense, 2020. https://www.defense.gov/Newsroom/Releases/Release/Article/2318799/defense-innovation-unit-announces-suas-product-availability-to-provide-secure-c/

[8] Defense Advanced Research Projects Agency,*With Squad X, Dismounted Units Partner with AI to Dominate Battlespace*, (Washington, DC: Defense Advanced Research Projects Agency, 2019), https://www.darpa.mil/news-events/2019-07-12

[9] Department of Defense, *Eyes of the Army: U.S. Army Roadmap for UAS 2010–2035*, (Washington, DC: Department of Defense, 2010). 7–9, https://fas.org/irp/program/collect/uas-army.pdf

[10] George Seffers, "Army Tactical Assault Kit Always Adapting for New Era," SIGNAL Magazine, October 28, 2020, https://www.afcea.org/content/army-tactical-assault-kit-always-adapting-new-era.

combat effectiveness and survivability by providing an organic ISR platform to support situational awareness.

The specific focus of this thesis is on the second pillar; the addition of an object recognition capability into a sUAS. Our principal task is to manage its integration onto a surrogate platform, and eventually onto the prototype. Working collaboratively with a private firm and DOD affiliates, we direct and monitor the integration and conduct interim performance evaluations of AeroVironment's Surveillance, Persistent Observation and Targeting Recognition (SPOTR) suite in a virtual environment and on a Nibbler sUAS. The Nibbler serves as the surrogate experimental sUAS based on its platform characteristics, its availability, its compatibility with ATAK, and the willingness of the manufacturer to allow us to access its open architecture.

B. RESEARCH OBJECTIVE

In this study we incorporate emerging object recognition software on to a sUAS to enhance the situational awareness of SOF teams. We direct the integration process and conduct simulation experiments to measure the performance of the SPOTR software to detect and categorize entities.

The primary research question is:

Can the integration of SPOTR object recognition into a sUAS achieve the technical parameters to meet SOF team operational requirements?

We evaluate the technical performance of the SPOTR software on a Nibbler sUAS for its potential operational impact. By managing this integration process, we advance the project toward a prototype that meets our specific performance requirements. We also establish a baseline for future student researchers to continue to develop this and the remaining pillars of the prototype with the integration of autonomous features and a common user interface.

C. SCOPE AND LIMITATIONS

We focus on Pillar 2, integrating object recognition capability into the sUAS with emphasis on its technical performance. Future work will address the subsequent pillars of autonomous capabilities and a common user interface. This research remains unclassified to aid following SATLAS project teams.

Software development cycles, the COVID 19 pandemic, and sUAS platform acquisition time constraints limit our ability to conduct tests in realistic operational scenarios. While the technical and operational requirements will vary across the various operating environments in which SOF may operation, we test only in the context of a temperate zone. In our field testing, we use the Nibbler. Our findings may not be applicable to other sUAS platforms and will need to be evaluated separately, depending upon which platform is selected as the prototype. We conduct our simulation experiments in collaboration with the commercial software developer.

D. ORGANIZATION OF THESIS

Chapter II provides a review of prior research, emerging operational doctrine and threat assessments, and the current state of sUAS and object recognition technology. Chapter III explains our research method, beginning with the background for the test criteria. It then describes our research design, tests, and results. Chapter IV presents the analysis of the results. Chapter V summarizes the significant findings and proposes the way ahead.

THIS PAGE INTENTIONALLY LEFT BLANK

II. LITERATURE REVIEW

This chapter is divided into four sections. The first section sets the tactical context with a factual vignette that illustrates the urgent need for a dedicated, highly capable sUAS for SOF teams. The second section explores what the United States has defined as the future operational environment and the evolving demand for technological solutions to better enable SOF teams. The third section examines military applications for sUAS and AI and the final section reviews prior work on this topic.

A. VIGNETTE – THE WARFIGHTER ISR SHORTFALL

In 2018, Special Forces Operational Detachment Alpha (SFOD-A) 0221 was deployed to a remote region in the Middle East and tasked to secure an area approximately 50 square kilometers that was controlled by Islamic State of Iraq and Syria (ISIS) militants. It was given continuous access to theater-level ISR and supporting fires. As the SFOD-A advanced deep into enemy territory over the course of 5 weeks, the distance between their forward outpost and the forward line of troops grew. To maintain offensive pressure, they employed multiple 3-man observation and surveillance teams to detect enemy positions and call for indirect and aerial delivered munitions on targets. These teams were generally more than 10 kilometers from any friendly position.

One day, during early morning hours, the theater-level ISR asset that had been providing surveillance for one of the observation teams was re-tasked to support a different SFOD-A. Without warning, that team was left without any external overhead surveillance support. The SFOD-A's only organic sUAS, a RQ-11 Raven, was inoperable due to broken parts and was also too large to be man-packable. As the 3-man surveillance team increased its security posture to compensate for the loss of the ISR asset, a squad of ISIS militants quietly approached them. The militants were not detected until they were within 75 meters of the team, well within firing distance. The ISIS squad immediately engaged the 3-man team with automatic weapons and hand grenades. As the surveillance team fought back, the SFOD-A launched a quick reaction force to assist, and the combined elements eventually eliminated the militants. During the firefight one soldier on the surveillance team was killed in action and the others were wounded, virtually eliminating the SFOD-A's combat effectiveness.

After consolidating and reorganizing, the team resumed combat operations several days later. To adjust for the increasing distance between the outpost position and the forward line of troops, the SFOD-A emplaced check points along the route between the two positions. Host Nation forces were unavailable or unwilling to operate the check points, so local militia forces volunteered their services. After some vetting, four militia check points were positioned along the route to provide security. The militia forces came armed with their own weapons and wore modest uniforms, and required a significant amount of support from the SFOD-A. This support required the U.S. forces to interact with militia members on a regular basis through a series of key leader engagements. The frequent leader engagements were assessed to be at low risk despite taking place in one of the most violent areas within the region. Therefore, the SFOD-A was assigned a low priority for support from the limited theater-level ISR assets.

It was not long before the SFOD-A again suffered the consequences of that prioritization. Upon departing a check point after one leader engagement, a 5-man team from the SFOD-A was ambushed by ISIS militants with a long burst of automatic machine gun fire from 50 meters away. Immediately, over half of the team was wounded or killed. They returned fire, began rudimentary triage, and engaged in a firefight as they desperately withdrew to a secure position. The team eventually received close air support and destroyed the enemy position. This author, the Detachment Commander of the SFOD-A, strongly contends that had the team been equipped with an organic sUAS, they would have detected the assaulting squad and prevented or effectively responded to the ambushes.

In an earlier example, an SFOD-A on a mission in Niger in 2017 was not supported by any ISR assets and in that incident four service members were killed in action.[11] The SFOD-A was engaged in a remote region in Africa with nominal external ISR support and was unaware that a large group of ISIS militants were quickly approaching their position.

[11] Alice Friend, *DOD's Report on the Investigation into the 2017 Ambush in Niger*, (Washington, DC: Center for Strategic and International Studies, 2018), 2–6, https://www.csis.org/analysis/dods-report-investigation-2017-ambush-niger.

The absence of surveillance support significantly limited their situational awareness and the SFOD-A was rapidly overwhelmed, became separated in the confusion, and suffered the loss of four Soldiers' lives.

It is the assessment of the SATLAS project team that as SOF are increasingly tasked to combat asymmetric threats in remote and contested regions where they cannot depend upon external support, they will require a reliable, versatile organic ISR asset. We further propose that this warfighter shortfall is urgent and brings unacceptable risk to SOF teams, yet a solution should be rapidly and inexpensively achievable.

B. THE FUTURE STRATEGIC ENVIRONMENT

In 2014, Kremlin forces wearing green uniforms absent any rank or insignia encircled government buildings, secured key sites, and successfully annexed Crimea from Ukraine.[12] Months of Russian-backed protests, propaganda campaigns, and economic pressure had set conditions that allowed Russia to seize the region with limited response from Ukraine or the West.[13] Russia's acts of irregular or hybrid warfare and China's whole-of-government approach to spread influence and expand its footprint to undercut the United States,[14] caused the United States to reevaluate its capability to deter or respond to near-peer adversaries.

The 2018 *National Defense Strategy* shifted to prioritize Great Power Competition while remaining committed to the defeat of violent extremist organizations.[15] The evolving doctrine of Multi-Domain Operations (MDO) is the U.S. Army's response to this re-prioritization as it attempts to grasp a future operational environment that includes the use

[12] Michael Kofman et al., *Lessons from Russia's Operations in Crimea and Eastern Ukraine*, RR 1498 (Santa Monica, CA: RAND, 2017), 5–16, https://www.rand.org/pubs/research_reports/RR1498.html.

[13] Heather A Conley et al., *The Kremlin Playbook: Understanding Russian Influence in Central and Eastern Europe*, (Washington, DC: Center for Strategic and International Studies, 2016), 1–5, https://www.csis.org/analysis/kremlin-playbook.

[14] Office of the Director of National Intelligence, *Annual Threat Assessment of the U.S. Intelligence Community*, (Washington, DC: Office of the Director of National Intelligence, 2021), 4–11, https://www.dni.gov/files/ODNI/documents/assessments/ATA-2021-Unclassified-Report.pdf

[15]Department of Defense, National Defense Strategy, 1–3.

of military and non-military means by near-peer and asymmetric adversaries to degrade U.S. influence and deny access.[16] MDO highlights four key trends to expect in the future operating environment; 1) adversaries will contest U.S. presence in all domains, 2) smaller armies will fight on larger and more lethal battlefields, 3) nation-states will be less able to impose their will, and 4) near-peer competition will take place below the level of armed conflict. The MDO also introduces three tenets to adapt the American way of war: a calibrated force posture, multi-domain operations, and convergence.[17] Guided by these tenets, ARSOF leadership has assessed the unique capability that its SFOD-A teams will need to provide within this operational concept. They must be regionally aligned, able to employ clandestine infiltration techniques, and operate by, with, and through partner or surrogate forces to reduce large U.S. footprints in sensitive areas.[18]

In addition to the changing threats, the Chief of Staff of the U.S. Army acknowledged the revolutionary impact that emerging technologies, such as artificial intelligence, nanotechnology, machine learning, and robotics, will have on the fundamental nature of war.[19] To prepare SOF teams for the future MDO environment, U.S. Special Operations Command (SOCOM) recently prioritized technological solutions to increase situational awareness and improve decision-making for small teams. The FY 2021 U.S. SOCOM Acquisition, Technology, and Logistics Directorate of Science and Technology announcement emphasizes efforts to hyper-enable SOF teams, recognizing that the future operating environment will require SOF to operate in "satellite denied/disrupted environments, under threat of targeting by high-end military capabilities…..and where increased scrutiny is routine."[20]

---

[16]Department of the Army, *The U.S. Army in Multi-Domain Operations 2028,* 15–20.

[17] Department of the Army, *The U.S. Army in Multi-Domain Operations 2028*, 15–20.

[18] Department of the Army, *Army Futures Command Concept for Special Operations 2028,* 9–11.

[19] Jim Garamone, "Milley Makes Case for U.S. Military Keeping Up With Global, Technology Changes," U.S. Department of Defense, (December 2, 2020): https://www.defense.gov/Explore/News/Article/Article/2432855/milley-makes-case-for-us-military-keeping-up-with-global-technology-changes/.

[20] Department of Defense, *Broad Agency Announcement For Technology Development and Advanced Technology Development* (Washington, DC: Department of Defense, 2020).

U.S. SOCOM defines the Hyper Enabled Operator (HEO) as a Special Operations professional equipped with technology that enables more timely and accurate decision-making while increasing situational awareness and minimizing cognitive overload.[21] To achieve an HEO, U.S. SOCOM has designated several focus areas for future research. Among these are autonomy-enabled ISR/battlefield situational awareness; group 1 sUAS payloads; edge computing to support localized SOF teams; and sensor algorithms to locate, classify, characterize, and identify items of interest.[22] These particular focus areas greatly informed the original conceptualization of the SATLAS project.

While still in the early stages of development, there is significant potential within these areas to create HEOs capable of defeating asymmetric and near-peer adversaries. Unfortunately, the Army appears to remain locked in a culture that prioritizes large, costly, and expensive platforms despite acknowledging a future operational environment in which smaller, less expensive, unmanned UAS will be most beneficial.[23] Additionally, the Army's byzantine acquisition system is overly complex and unable to field technologies in any reasonable time.[24] These disconnects and the proliferation of COTS technologies capable of meeting the demands of SOF teams have motivated bottom-up approaches from the field to solve this critical shortfall.

C. EMERGING TECHNOLOGY

While no single technology will help to create HEOs, there are technologies whose integration could hyper enable operators through the real-time collection of data, distilling the data down into mission relevant information, disseminating the data to personnel able

[21] Department of Defense, *Broad Agency Announcement For Technology Development and Advanced Technology Development.*

[22] Department of Defense, *Broad Agency Announcement For Technology Development and Advanced Technology Development.*

[23] Liam Collins and Harrison Morgan, "Affordable, Abundant, and Autonomous: The Future of Ground Warfare," *War on the Rocks* (April 21, 2020): https://warontherocks.com/2020/04/affordable-abundant-and-autonomous-the-future-of-ground-warfare/.

[24] Jennifer McArdle, "Simulating War: Three Enduring Lessons from the Louisiana Maneuvers," *War on the Rocks* (March 2021), https://warontherocks.com/2021/03/simulating-war-three-enduring-lessons-from-the-louisiana-maneuvers/.

to use it, and, ultimately, strengthening our capability to compete in a MDO environment.[25] While the vision of the SATLAS project team is to integrate four of these technologies, including a sUAS platform, AI-enabled object recognition, autonomy, and a common user interface, this thesis focuses on developing Pillar 2.

1. sUAS

Militaries have experimented with UAS since the 19th century to increase situational awareness and reduce what Carl Von Clausewitz described as “friction”.[26] While the explosion in the use of drones may appear to be a recent phenomenon, the first use military use of a drone actually occurred in 1849, when Austria used a hot air balloon to bomb Venice during the First Italian War of Independence.[27] Though the earlier applications of UAS platforms were nearly indistinguishable from missiles, the experimental use of drones continued through nearly all major conflicts.[28] Significant technological advancements in the late 20th and early 21st centuries have led to the use of drones by the United States in both lethal and nonlethal applications to address the asymmetric threats that have risen in the post-9/11 era.[29] The capability of UAS platforms to provide near real time intelligence to military commanders and civilian senior leaders has led these systems to become associated with the American style of war.[30]

Many U.S. SOF teams are still equipped with the RA-11B Raven as their organic UAS; however, advances in the civilian market have enabled research into smaller, more capable systems. The Raven has been fielded in its current form since 2006 and has been

[25] Department of Defense, *Broad Agency Announcement For Technology Development and Advanced Technology Development*

[26] Carl Von Clausewitz, *On War*, Translated by Michael Howard and Peter Paret, (Princeton University Press, 1976), Book 1, Chapter 7.

[27] Higinio Gonzales-Jorge et. al., *Unmanned Aerial Systems for Civil Applications: A Review,* (MDPI: July 2017), 1–2, https://doi.org/10.3390/drones1010002

[28] Jack Miller, “Strategic Significance of Drone Operations for Warfare,” *E-International Relations*, (August 2013), https://www.e-ir.info/2013/08/19/strategic-significance-of-drone-operations-for-warfare/.

[29] Milena Sterio, “The United States’ Use of Drones in the War on Terror: The Legality of Targeted Killings under International Law,” *Case Western Reserve Journal of International Law* 45 (2012), 198–200, https://scholarlycommons.law.case.edu/cgi/viewcontent.cgi?article=1072&context=jil.

[30] Miller, “Strategic Significance of Drone Operations for Warfare,” 2013.

used extensively by U.S. SOF teams in nearly every region.[31] The emergence of advanced commercial sUAS, such as the Chinese-made DJI, and their potential uses on the battlefield have led to significant efforts by the DOD to field similarly capable systems. Notably, UAS platforms have been substantially developed to include Vertical Take-Off and Landing (VTOL), complex camera systems, and smaller and more agile platforms. As a result, the DOD launched the Rucksack Portable UAS (RPUAS) program of record[32] while the Defense Innovation Unit (DIU) initiated the Blue UAS program, both designed to introduce inexpensive, rucksack-portable, VTOL-capable sUAS systems to the battlefield.[33]

In 2013, the U.S. Army published capability requirements for the next generation of sUAS designed to enhance the situational awareness of small unit commanders.[34] The document described production threshold requirements, the minimum capabilities necessary for a sUAS to be considered for the program of record, and production objective requirements, the desired capabilities for a sUAS to move forward towards product development.[35] Key production objectives and production threshold requirements include; a handheld system capable of being launched without dedicated devices and a compact system capable of being transported in a rucksack. Additional requirements can be seen in Table 1.

---

[31] Department of Defense, *RQ-11B Raven Small Unmanned Aircraft System (SUAS)* (Washington, DC: Department of Defense, 2016), https://asc.army.mil/web/portfolio-item/aviation_raven-suas/

[32] Department of Defense, *Capability Production Document For Rucksack Portable Unmanned Aircraft System (RPUAS) Increment* 2, (Washington, DC; Department of Defense, 2013).

[33] Department of Defense*, Defense Innovation Unit Announces sUAS Product Availability to Provide Secure, Capable Small Unmanned Aerial Systems for Critical Uses Across the Government* (Washington, DC: Department of Defense, 2020) https://www.defense.gov/Newsroom/Releases/Release/Article/2318799/defense-innovation-unit-announces-suas-product-availability-to-provide-secure-c/#:~:text=The%20Defense%20Innovation%20Unit%20(DIU,options%20to%20the%20U.S.%20Government.

[34] Department of Defense, *Capability Production Document For Rucksack Portable Unmanned Aircraft System (RPUAS) Increment 2.*

[35] Department of Defense, *Capability Production Document For Rucksack Portable Unmanned Aircraft System (RPUAS) Increment 2* (Washington, DC: Department of Defense, 2013).

Table 1. RPUAS Capability Requirements[36]

Joint Capability Area	Production Threshold	Production Objective
2. Battlespace Awareness 2.1 Planning & Direction 5. Command & Control 5.1 Organize 5.2 Understand 5.5 Direct	30-minute time of flight and a flight range of 3km	45-minute time of flight and a flight range of 5km
2. Battlespace Awareness 2.1 Planning & Direction 3. Force Application 3.1 Maneuver	- Payload with sufficient resolution for an operator to have a 90% Probability of Detection (PD) of a man-sized target during day at 300m and 200m at night. - 90% PD of a vehicle during day at 400m and 300m at night - 90% PD of a vehicle or person during day at 250m and 150m at night	- Payload with sufficient resolution for an operator to have a 95% PD of a man-sized target during day at 600m and 450m at night - 95% PD of a vehicle during day at 800m and 600m at night -95% PD of a vehicle or person during day at 600m and 450m at night
3. Force Application 3.1 Maneuver	- Contain one handheld ground control station (H-GCS) - Weigh no more than 19lbs	- Contain one handheld H-GCS - Weigh no more than 8lbs
2. Battlespace Awareness 3. Force Application 3.1 Maneuver 7. Protection 7.2 Mitigate	- Inaudible at 200 feet AGL with a background noise of 65 dBA	- Inaudible at 100 feet AGL with a background noise of 35 dBA
2. Battlespace Awareness 2.1 Planning & Direction 2.2 Collection 3. Force Application 5. Command & Control 6. Net-Centric	- The system must have a modular FMV payload containing day, night (passive infrared), and laser illuminator - The system must have the ability to observe a stationary object/location while sending and receiving C2 data to/from the H-GCS - The H-GCS display must have the detail/quality of a Soldier with 20/20 vision observing the object.	- The system must have a modular FMV payload containing day, night (passive infrared), laser illuminator, and Laser Target Marker). - Both day and night cameras must have optical zoom capability

36 Department of Defense, *Capability Production Document For Rucksack Portable Unmanned Aircraft System (RPUAS) Increment 2*.

Additionally, a key component of both the RPUAS and Blue UAS programs is ensuring the security and safety of any potential sUAS systems introduced into the DOD. To secure the integrity of UAS platforms and mitigate the risks of foreign intrusions into the supply chain, SRR and Blue UAS were both designed to comply with Section 848 of the *National Defense Authorization Act* (NDAA) for Fiscal Year 2020.[37] This section prohibits the use or procurement of foreign-made UAS platforms and attempts to prevent foreign adversaries, including China, from compromising sensitive areas of DOD research.[38] As a result of 18-months of testing and supply chain inspection, DIU and the DOD announced the availability of five UAS options able to be purchased by DOD organizations in September 2020 including: Altavian, Parrot, Skydio, Teal, and Vantage Robotics.[39] The SATLAS project team was able to acquire three Altavian M440s for testing and prototype development and plans to eventually integrate the Skydio X2D.

2. Object Recognition

As sUAS platforms now play a significant role in increasing the situational awareness of ground forces, the employment of object recognition algorithms is even more feasible. Object recognition is a longstanding problem within computer vision but is an area in which deep learning techniques have enabled recent break throughs.[40] The goal of object recognition is to determine whether specific objects from given categories (e.g., humans, vehicles, or weapons) are present in an image. Deep learning allows for the use of multilayered algorithms to enable software to learn the representation of images through multiple levels of abstraction.[41] This type of learning allows algorithms to be trained using

[37] The Department of Defense*, Defense Innovation Unit Announces sUAS Product Availability to Provide Secure, Capable Small Unmanned Aerial Systems for Critical Uses Across the Government.*

[38] National Defense Authorization Act for Fiscal Year 2020, S.1790, 116th Congress (2019-2020) https://www.congress.gov/bill/116th-congress/senate-bill/1790/text.

[39] The Department of Defense*, Defense Innovation Unit Announces sUAS Product Availability to Provide Secure, Capable Small Unmanned Aerial Systems for Critical Uses Across the Government*.

[40] Li Liu et al., "Deep Learning for Generic Object Detection: A Survey," *International Journal of Computer Vision* 128, no. 2 (February 2020): 261–318, https://doi.org/10.1007/s11263-019-01247-4.

[41] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep Learning," *Nature,* no. 7553 (May 2015): 436–44, https://doi.org/10.1038/nature14539.

known data sets on which objects to recognize and can be tailored to best meet the needs of the operator.

The use of object recognition in military applications has been considered in several projects and studies, including the highly publicized Project Maven[42]; however, this capability has yet to be implemented nor have the requirements for DOD use been clearly defined.[43] ISR systems capable of accurately recognizing potential threats on the battlefield have the potential to significantly improve the survivability of SOF teams.

One common critique of using object recognition to assist military operations is the question of accuracy. Object recognition performance is often evaluated using three criteria: frames per second, precision, and recall. Precision is a measure of how frequently an algorithm correctly recognizes an object.[44] Recall measures whether the algorithm recognized an object every time it should.[45] As an algorithm's recall rate increases, its precision often decreases creating the precision-recall curve (Figure 1). The area underneath the precision-recall curve is referred to as the average precision and indicates correct guesses on object recognition based on the correct number of times the algorithm should have made a guess. Thus, an algorithm can be set on a sensitivity continuum to either make only correct guesses while missing opportunities in which it should have guessed (false negative) or to make guesses at every opportunity which will result in inaccurate guesses (false positives) or somewhere in the middle. This has significant implications for use by the DOD, especially with human-in-the-loop systems. An operator

---

[42] Lucy Suchman, "Algorithmic Warfare and the Reinvention of Accuracy," *Critical Studies on Security* 8, no. 2 (May 3, 2020): 175–87, https://doi.org/10.1080/21624887.2020.1760587.

[43] Doaa Mohey El-Din, Aboul Ella Hassanein, and Ehab E. Hassanien, "An Automatic Detection of Military Objects and Terrorism Classification System Based on Deep Transfer Learning," in *Proceedings of the International Conference on Artificial Intelligence and Computer Vision (AICV2020)*, ed. Aboul-Ella Hassanien et al., Advances in Intelligent Systems and Computing (Cham: Springer International Publishing, 2020), 594–603, https://doi.org/10.1007/978-3-030-44289-7_56; Zhi Yang et al., "Deep Transfer Learning for Military Object Recognition under Small Training Set Condition," *Neural Computing and Applications*, no. 10 (October 1, 2019): 6469–78, https://doi.org/10.1007/s00521-018-3468-3.

[44] Shivy Yohanandan, "MAP (Mean Average Precision) Might Confuse You!," Medium, June 9, 2020, https://towardsdatascience.com/map-mean-average-precision-might-confuse-you-5956f1bfa9e2.

[45] Yohanandan, "MAP (Mean Average Precision) Might Confuse You!".

can opt to set the level of sensitivity to mitigate risk during high consequence scenarios[46] in order to sacrifice some precision to increase recall abilities.

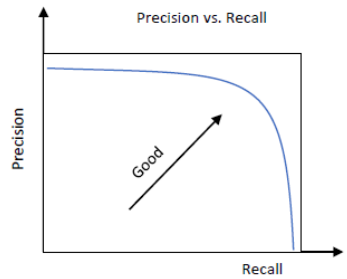

Figure 1. Precision-Recall Curve47

D. PRIOR WORK

This section provides an overview of research that informed this thesis. These research efforts are addressed individually in the following sections.

1. Deep Learning and sUAS Object Recognition

Recent advancements may allow deep learning (DL) algorithms to enable UAS platforms to detect threats and autonomously avoid obstacles. The use of DL techniques to recognize patterns from raw data captured by onboard cameras could increase the autonomous functionality of UAS platforms.[48] Differing from classic machine learning, DL does not require the use of descriptor labels to categorize data. Instead, DL techniques can both identify and categorize data simultaneously. In essence, DL algorithms process raw video from camera systems or sensors onboard UAS and determine if the data represents a threat. Fraga-Lamas et. al. proposed a cloud-based communication system that

[46] DataLabeler L, "Human-in-the-Loop Machine Learning Approach," Medium, March 24, 2020, https://datalabeler.medium.com/human-in-the-loop-machine-learning-approach-b130102b94e5.

[47] Joe Kehoe, email message to author, May 7, 2021.

[48] Paula Fraga-Lamas et. al., "A Review on IoT Deep Learning UAV Systems for Autonomous Obstacle Detection and Collision Avoidance," *Remote Sensing* 11, no. 18 (January 2019): 2144, https://doi.org/10.3390/rs11182144.

can contain the required deep learning algorithms for object recognition. The UAS could then automatically reposition itself to avoid objects, maintain its stealth, or alert the operator. Deep learning-enabled object recognition and cloud-based communication techniques are major considerations in this thesis' proposed solution to increase situational awareness for SOF teams.

2. SQUAD X Experiments

Between 2018 and 2020, the DOD's Defense Advanced Research Projects Agency (DARPA) conducted a series of experiments, labeled SQUAD X, to test the impact of enhancing U.S. Marine Corps squads with DL-enabled robots.[49] DARPA also employed a cloud-based communication architecture connected to ground stations that controlled aerial and land-based robots. Both robots employed a version of AeroVironment's (formerly Progeny Solutions) Surveillance, Persistent Observation and Targeting Recognition (SPOTR) suite. While focusing on developing modular and open architecture-based software, AeroVironment integrated computer vision analytics, namely object recognition, and machine learning through deep neural networks to develop their SPOTR technology.[50] SPOTR has been tailored for use in unmanned applications through its embedded and edge-processing configurations and, once integrated into a UAS' computer vision, enables the system to recognize and detect threats. SPOTR employs complex algorithms and training sets to detect, recognize, categorize, and track potential threats. Detection implies that SPOTR discovers than an object is present through processing camera footage. The object is then recognized as an object that SPOTR has been trained to identify. It is then categorized as the specific trained object (a weapon, person, or vehicle in our case). SPOTR then follows the object and tracks its movement. As seen in Figure 2, potential threats are identified to the operator by outlining the object with a green box in the user interface. A single platform can track multiple threats.

---

[49] Defense Advanced Research Projects Agency, "With Squad X, Dismounted Units Partner with AI to Dominate Battlespace," (July 2019): https://www.darpa.mil/news-events/2019-07-12

[50] Daniel Midgett, email message to author, September 4, 2020.

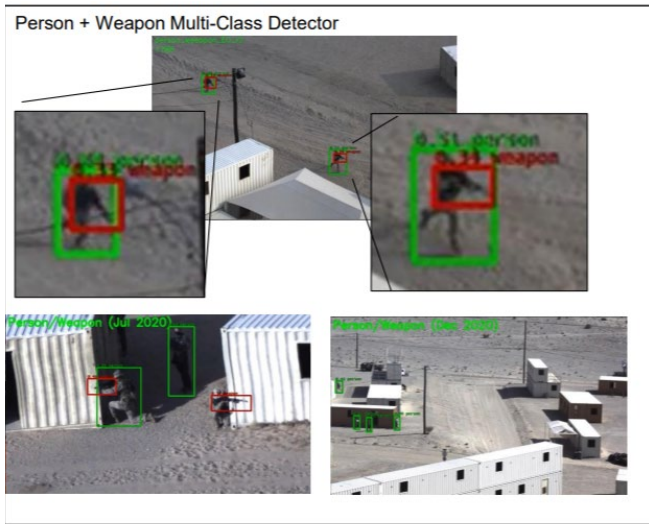

SPOTR Object Recognition Model[51]

During a series of four experiments, AeroVironment's SPOTR software was integrated into UASs to enable U.S. Marines tasked to execute a series of operations in urban areas.[52] DARPA discovered that not only could the object detection algorithms significantly increase the situational awareness of ground forces, it also provided the forces enough digestible information to adjust their ground scheme of maneuver without cognitive overload when the data was shared among the team through the Android Tactical Assault Kit (ATAK).[53] Additionally, feedback from the Marine squad members demonstrated the user-friendly nature of the system. The users also gained trust in the reliability of the system after a single iteration.[54] SQUAD X demonstrated that it is feasible to employ object recognition capability in a human-in-the loop system to increase situational awareness for ground forces. The tool must not cognitively overload the operator, and trust in the tool can be built through reliable detection of enemy threats. SPOTR is the software that we evaluate for its object recognition capabilities.

---

51 AeroVironment, "SATLAS SPOTR Metrics".

52 Defense Advanced Research Projects Agency, "With Squad X, Dismounted Units Partner with AI to Dominate Battlespace," (July 2019): https://www.darpa.mil/news-events/2019-07-12.

53 Keenan Kline, email message to author, November 12, 2020.

54 Keenan Kline, email message to author, November 23, 2020.

3. Object Recognition and Super-Resolution

Developing trust in object recognition systems requires accurate and reliable detections across all conditions. Traditional object recognition is usually performed when images that the algorithm is trained to detect occupy a large portion of the image frame. This technique is largely dependent on the quality of the training data, or group of images used as samples to train deep neural networks on what to detect, and the quality of the image observed. All object recognition models require a certain number of pixels to recognize an item. As the quality of observed images decreases, the risk of mis-categorization or completely missed items increases. Nearly all object recognition training data introduce some level of image degradation to allow for accurate detection when items are small or are partially obscured. This is especially important when considering sUAS which are equipped with smaller, less capable cameras to minimize weight and remain man packable. Additionally, SOF operations often occur in environments with limited visibility, and required standoff that stress traditional camera systems, incurring risk to the team.

In 2019, Christoph Borel-Donohue and Susan Young proposed super-resolution as a technique to increase accurate object recognition rates on degraded images.[55] They used low-resolution images of parking lots and fed them into the You Only Look Once (YOLO) object recognition network to test the number of cars and trucks detected. After the low-resolution frames were scanned by YOLO, they processed the images through a super-resolution algorithm to produce full size images then rescanned them. Their findings show that the number of detections improved two-fold after the images were super-resolved and matched the detection results of an image four times its size. Each object recognition model requires different image sizes and qualities for accurate detection; however, introducing a super-resolution algorithm can help increase detection accuracy while avoiding increased hardware or payload requirements. This suggests that super-resolution algorithms may provide a method to increase object recognition capabilities while avoiding increased sUAS payload weights. This technique informed our analysis of the relationship between

[55] Christoph Borel and S. Young, "Image Quality and Super Resolution Effects on Object Recognition Using Deep Neural Networks," in 2019 SPIE *Conference for Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications.* (SPIE, 2019), 7–8, https://doi.org/10.1117/12.2518524.

object recognition algorithms and onboard camera systems and presents opportunities to improve the range of SPOTR while minimizing weight.

4. AI Robots and Increased Combat Effectiveness

There are several important implications for the use of DL-enabled sUAS organic to a SOF team. In 2020, Midgett et. al. analyzed simulation data from the Maneuver Battle Lab at Fort Benning, GA to estimate the additional combat strength gained by augmenting U.S. Mechanized Infantry platoons with AI-enabled robots and precision strike capabilities.[56] The lab provided data from four controlled simulations in which the simulated engagements did not include AI-enabled robots, and from nine experiment simulations that included the robots. The robots employed deep neural network (DNN) AI and comprised armed ground robots, armed aerial robots, and company-level precision strike robot systems.

Midgett et. al. analyzed the data from simulated engagements between friendly mechanized infantry platoons and enemy forces that ranged in size from a platoon to a battalion. They found a strong statistical significance ($r = 0.96$) between the use of armed aerial robots and precision strikes and improved platoon-level combat effectiveness. Furthermore, they assessed that the addition of AI robots and precision strikes accounted for 92% of the platoon's increase in combat effectiveness. They were unable to determine with any statistical significance whether the use of ground robots impacted combat effectiveness but found that the aerial robot systems could explain 56% of the observed effects of engagements. As seen in Figure 3, this finding suggests that there may be merit in reexamining the traditional rule of only attacking when one holds a three to one force advantage over an opponent. For the purposes of this thesis, Midgett et. al. demonstrated the significant advantage that AI-enabled sUAS might provide to small teams and strengthened our notion to focus on integrating inexpensive systems for SOF teams.

[56] Midgett et. al., "Semiautonomous Threat Learning Alert System," 46–52.

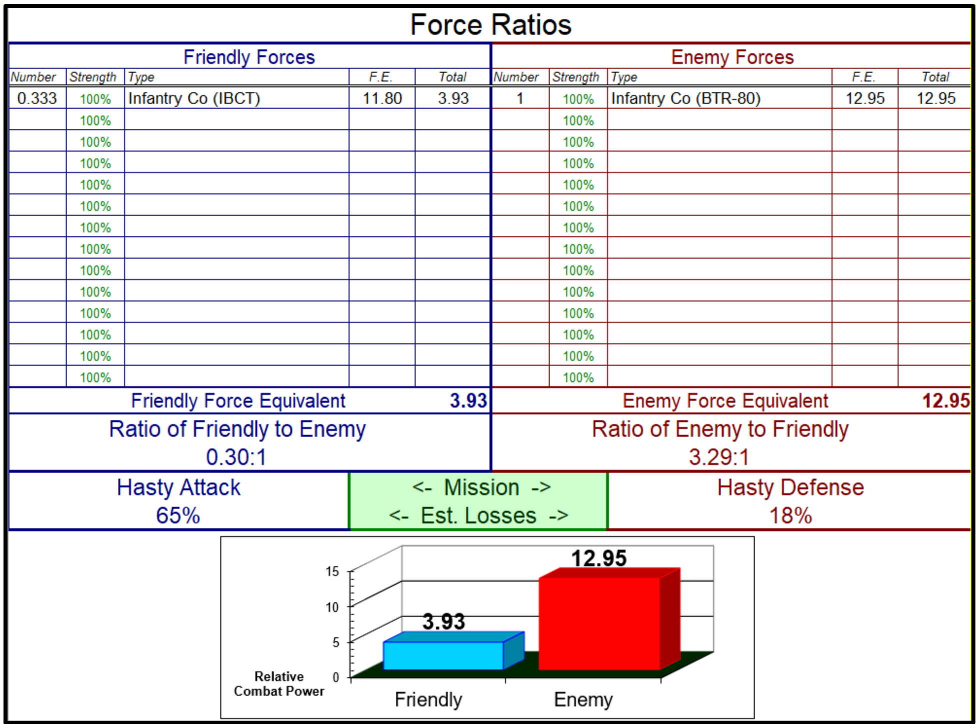

Figure 2. Correlation of Forces Calculator Results of AI-Enabled Platoons57

Since the concept of using an sUAS platform with an object recognition tool to increase SOF situational awareness is only theoretical, the first step in this thesis is to assess whether the integration of object recognition software onto a versatile and robust UAS platform is feasible. This thesis focuses on Pillar 2 while managing the integration into a specific sUAS platform. The next chapter describes the research process, experiment design, execution, and the results.

57 Midgett et. al., "Semiautonomous Threat Learning Alert System," 46–52.

III. EXPERIMENT DESIGN

This chapter describes the series of experiments conducted to evaluate the object recognition integration and our findings. We tested the performance of AeroVironment's SPOTR object recognition algorithm to determine the feasibility of integrating it into a surrogate sUAS platform to enhance situational awareness for SOF units.

We conducted our experiments in four phases. During Phase 0 we participated in simulation software validation with AeroVironment engineers, and to become familiar with their progress in coding SPOTR's ability to detect and categorize individuals and vehicles in a fully simulated environment. In Phase 0 we sought to determine whether the simulated environment that we had selected could adequately and reliably represent the ability of SPOTR to process raw data from a ArduCopter sUAS platform and categorize moving entities as individuals, weapons, or vehicles. We refined and scoped our draft research questions and measures of performance in order to transition to Phase I.

Phase I consisted of discovery experiments using an Alienware laptop loaded with the SPOTR software and the simulated environment. The simulated environment was a synthetic representation of the Urban Training Center at the U.S. Marine Corps base in Quantico, VA. This phase allowed us to examine whether the simulation capabilities through multiple iterations of runs and narrow our research questions while collecting SPOTR performance data. The Alienware laptop provided sufficient processing power to support the simulation and SPOTR software without limiting performance. This allowed us to become familiar with SPOTR in a fully unconstrained processing power environment.

Phase II consisted of initial evaluation experiments to limit the processing load available to SPOTR to that available onboard current sUAS to determine the feasibility of integrating the two. We hypothesized that onboard processing power would be sufficient to reliably run the algorithms and wanted to measure the impacts of a constrained environment on the software.

Phase III was designed to test whether the surrogate sUAS could support the object recognition algorithms. This phase involved transitioning from the simulated environment to a field environment to measure the impact on platform and software performance once integrated. The experiments were designed to be basic but provide evidence to further advance the project to determine the feasibility of eventually introducing the technology into the DOD. Figure 4 describes our experiment design plan.

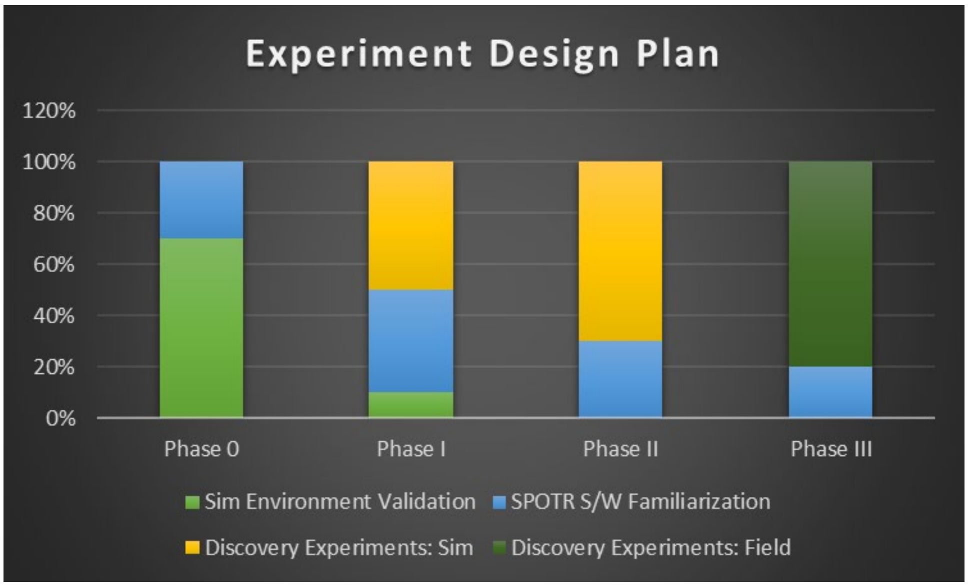

Figure 3. Experiment Design Campaign Plan

Each phase of the experiment campaign was designed with a systems approach in mind and considered the parameter space investigation (PSI) method as well as the DOD's Code of Best Practices for Experimentation. Section A describes the design methodology while Section B details the findings for each phase.

A. EXPERIMENT DESIGN METHODOLOGY

The PSI method, or Multicriteria Analysis, is designed to meet the demands of engineer optimization problems and one that offers important considerations for our project. R. Statnikov et. al. describes this method as a process of correctly identifying a problem and then constructing and analyzing a feasible solution.[58] Important features of this method are

[58] R. Statnikov et. al., "Multicriteria Analysis Tools in Real-Life Problems," *Computers & Mathematics with Applications* 52 (2006) 1–32. https://doi.org/10.1016/j.camwa.2006.08.002

the design variables, functional constraints, and criterion variables. Design variables are often independent variables and the factors thought to be responsible for change in performance. These variables are often limited due to the available resources or technology and often require extensive market research to properly identify. Functional constraints are factors that often fall outside of the intended area of interest yet play an important role in the design and execution of any experiment. These constraints include factors such as the time available, testing sites, and weather, which must also be properly considered. Finally, criterion variables are often those considered to be essential by the designer for the proposed solution to be feasible. These variables may involve a required time of flight, a flight range, or the amount of data able to be processed. Table 2 lists the design variables, functional constraints, and criteria constraints specific to this project. The following section outlines how each impacted this study.

Table 2. Design, Functional, and Criterion Variables[59]

Design Variables	**Functional Constraints**	**Criterion Variables**
sUAS altitude	Software integration timeline	SPOTR accuracy
sUAS Slant Range Distance from entity	Testing Window	Number of entities detected
Number of entities	Test environment	Maximum distance of detection and categorization
Target location		

The following section describes our down select process and how each of these variables and constraints impacted our study and influenced our results.

1. Down Select Considerations

We initially selected specific hardware and software products for our study. In most cases, several options of each product were available to our team but only one could move

[59] R. Statnikov et. al., “Multicriteria Analysis Tools in Real-Life Problems,”1-32.

forward in the experiment phases. In some cases, only one product was available. The following section describes the products that we selected for our experiments.

a. SPOTR Object Recognition Software

As discussed in Chapter II, AeroVironment's SPOTR is the only object recognition software that we considered which is designed to be integrated with sUAS and able to detect potential threats to ground forces. SPOTR is unique is in its ability to detect objects at the edge. The data is streamed directly from the platform to the operator's GCS component and it can be integrated into the onboard computer vision of most sUAS platforms.[60] The SPOTR includes a NVIDIA Jetson TX2 processor, a camera system capable of EO/IR configurations and a Microhard pDDL data radio. Once integrated into a sUAS platform, SPOTR employs custom-tailored DNN and DL algorithms that detect patterns consistent with known targets. As seen in Figure 5, SPOTR is being designed to detect, categorize, localize, and then track (DCLT) potential targets based on identifiers detected during training and validation.

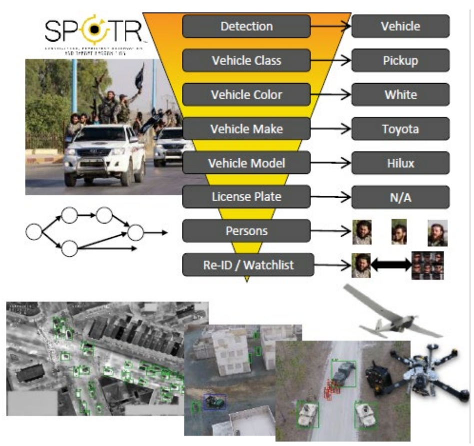

Figure 4. SPOTR Classification Process61

The object recognition algorithms currently used by the SPOTR software support detection in three categories: individuals, weapons, and vehicles. However, the version used

---

[60] Daniel Midgett, email message to author, September 4, 2020.

[61] Daniel Midgett, email message to author, September 4, 2020.

in our laptop-based simulation testing was limited to two categories: individuals and vehicles. This was the only model ready to be downloaded onto a third-party computer.

b. Platform

The SATLAS team originally selected the Altavian M440 as our preferred sUAS platform for prototype development and testing; however, discussions with the manufacturer about integration and the requirement of non-disclosure agreements prevented the Altavian from being available. While AeroVironment (formerly Progeny Solutions) and Altavian worked through the bureaucratic requirements, testing moved forward with the Nibbler sUAS platform.

Although limited on advanced technology, the Nibbler sUAS platform has been successfully fielded with the U.S. Marine Corps. It can be 3-D printed and assembled and repaired at the unit level.[62] Not only does this keep the platform relatively cheap (costing around $2,000 per platform) but it also avoids the extensive DOD acquisition process.[63] The Nibbler is a small, lightweight, VTOL-capable, four rotor quadcopter fitted with a single lens camera capable of visible detection (Figure 6). It has a 20-minute time of flight and was designed to increase situational awareness by conducting reconnaissance in support of small units. It was designed and developed through a collaborative effort between the MITRE Corporation and the U.S. Marines.

---

[62] Megan Eckstein, "Marines' 3D-Printed 'Nibbler' Drone Creating Lessons Learned on Logistics, Counter-UAS," *USNI News*, (September 2017), https://news.usni.org/2017/09/27/marines-3d-printed-nibbler-drone-creating-lessons-learned-logistics-counter-uas.

[63] Bill Eidson, "Nibbler Drone Is an Advanced Manufacturing 'Flagship' for Marines," *MITRE*, (January 2019), https://www.mitre.org/publications/project-stories/nibbler-drone-is-an-advanced-manufacturing-flagship-for-marines.

Figure 5. Nibbler sUAS[64]

Engineers from AeroVironment have worked extensively with the Nibbler and, therefore, integration with this platform was not an issue. The Nibbler is approved for use by the DOD and has been employed almost solely by the U.S. Marines since 2017. However, the Nibbler was not considered by the RPUAS program due to its limited autonomous capability. Nonetheless, it served adequately as the surrogate for our research pertaining to object recognition.

c. Microprocessor

The NVIDIA Jetson TX2 microprocessor is considered the fastest and most power-efficient microchip for advanced AI computing[65] and is the only microprocessor used by AeroVironment. The TX2 has 8 GB of memory and uses 7.5 watts of power to enable on-the-edge AI computing. The TX2 is one of the most common microprocessors employed for AI processing. The TX2 was the only processor used during our experiments. Results on other microprocessors may vary.

d. Onboard Camera

The simulated environment and live experiments employed an onboard camera system equivalent to the GoPro Hero 4 Silver. This is the camera available on the Nibbler and

[64] Bill Eidson, “Nibbler Drone Is an Advanced Manufacturing ‘Flagship’ for Marines,”.

[65] NVIDIA, “Jetson TX2 Module,” NVIDIA Developer, (May 2017), https://developer.nvidia.com/embedded/jetson-tx2.

the specifications of the simulated environment were adjusted to match. The Hero 4 records video in 1080p resolution at 30 frames per second and operates at a super wide, wide or narrow field of view.[66] The GoPro Hero 4 is primarily designed to record objects in close proximity. The lack of zoom capabilities limits its performance as compared to other sUAS onboard cameras. For our research, this restricted the range at which SPOTR was able to detect and categorize entities.

e. Simulation Software

We used AeroVironment's simulation as the environment to evaluate the SPOTR software while integration efforts were ongoing. To examine the capabilities of SPOTR, we loaded the simulated environment on the Alienware laptop. The laptop provided an unconstrained setting for the software to perform its tasks.

A constraint for this study was that the only simulated environment available for experimentation was the one created by AeroVironment. The simulation is designed to replicate the Urban Training Center at the U.S. Marine Corps base in Quantico, VA. It incorporates a small urban area of approximately 150 buildings while the rural areas are limited to open fields. This restricted our tests to take place in a temperate environment only. Additionally, the standard version of the simulation limits the number of detectable objects to 13 individuals, 13 weapons, and 1 vehicle depending on the region scanned. This limited our ability to detect more than 13 objects simultaneously.

2. Design Variables

Design variables are often described as the independent variable.[67] They are the knobs that we turn in order to measure the impacts. We selected the following four design variables specific to this part of the project include: 1) operating altitude of the sUAS, 2) the slant range distance between the sUAS and the entity, 3) the number of entities available for detection and categorization, and 4) target location. During our early test phases, we conducted

[66] GoPro, *GoPro Hero 4 Silver User Manual,* 2014. https://gopro.com/content/dam/help/hero4-silver/manuals/UM_H4Silver_ENG_REVA_WEB.pdf.

[67] R. Statnikov et. al., "Multicriteria Analysis Tools in Real-Life Problems,"1-32.

discovery experiments to determine the appropriate parameters of our independent variables as they had not yet been specified.

a. sUAS Altitude

The altitude of the sUAS was the most easily manipulable design variable and revealed the most about SPOTR's capabilities. In both the simulation and field tests, the altitude of the virtual ArduCopters could be adjusted in five-meter increments using the mission control panel in the application screen. SPOTR is trained to detect entities based on pattern recognition and pixel size. Increasing the sUAS altitude reduced the pixel size of the detection. By adjusting the altitude, we identified optimal performance parameters and limitations of SPOTR to detect and categorize entities based on the available experiment settings and conditions.

b. Slant Range Distance from Entity

Similar to altitude, the distance between the entity and the sUAS was easy to manipulate and impacted the number of pixels available for SPOTR to detect and categorize entities. This distance could be adjusted using the mission control panel to alter the flight path of the drone either manually or by inputting a route for the sUAS to automatically fly. Changes in distance and altitude showed similar impacts on SPOTR's accuracy and were recorded as the slant range distance as discussed in Chapter II. From these, we established initial performance parameters.

c. Number of Entities

An important research question for the SATLAS project involves the maximum number of entities SPOTR is able to detect and categorize simultaneously. To measure this, we manipulated the number of entities available to be detected and categorized. The intent was to identify the limitations of the system and direct the necessary design adjustments to meet our evaluation criteria.

d. Target Location

Target location is important as pixel size and contrast directly impact object recognition capabilities. Within a given test environment, shadows, vegetation, and dark backgrounds negatively impact SPOTR's ability to detect and categorize entities. We measured detection and categorization accuracy in various locations within the available testing environments to evaluate these impacts. Entities within the simulation travel on programmed routes and are not easily adjustable. We were, however, able to select different regions within the simulated test environments that offered varying conditions of shadows and background contrast. Target location proved much easier to manipulate during the field experiment; live targets could travel within any micro-terrain at the test site. These variations enabled us to gather initial data regarding the impacts of vegetation, shadows, and varying backgrounds on SPOTR accuracy. We used those data to guide the software developers.

3. Functional Constraints

Dr. Alberts describes functional constraints as variables or environmental factors accepted by the designer.[68] These constraints impact the relationship between design variables and the criteria constraints and directly influence experiment results. Several functional constraints impacted this research and the results that follow.

a. Software Integration Timeline

This study relied on government funding and commercial industry experts to successfully integrate object recognition software into a sUAS platform. As a result, the experiments were limited to testing the products available within the scope of what could be produced and funded in the given time. This timeline was further constrained by an 18-month academic period. These constraints limited testing window and required the study to employ a simulated environment while observing remotely. This did not allow the team to test the Altavian and Skydio platforms, and required the use of the Nibbler sUAS platform to evaluate the object recognition performance.

[68] David S. Alberts and Richard E. Hayes, *Code of Best Practice for Experimentation*, CCRP Publication Series ([Washington, D.C: DOD Command and Control Research Program, 2002), 68-73, http://dodccrp.org/files/Alberts_Experimentation.pdf

b. Testing Window

As a result of the software integration timeline constraint, the timeline for testing was limited to a period of six weeks. This constraint, and the required coordination with industry partners, meant that the experiments had to remain relatively small and necessitated the use of simulations and subsequent observing field experiments.

c. Test Environment

The testing environment was limited to what was available in the prototype developers' virtual environment simulating a training site at the U.S. Marine Corps base in Quantico, Virginia. The field test environment was an open wooded area in Prince William County, Virginia. We were relegated to temperate zones with limited foliage; follow-on research should be conducted in various environments.

4. Criterion Variables

Criterion variables are defined as the dependent variables of the experiment and are the outputs. These can be thought of as the behaviors the system will perform as a function of technical capabilities and operator settings.[69] Prior to the experimentation phase, we identified several criteria variables that we believed would be of significance in determining the feasibility of employing the object recognition software to increase the situational awareness of ground forces.

This part of the longitudinal study focused on discovery testing and, as such, our criteria variables were defined as: 1) the ability of the object recognition software to accurately detect and categorize entities, 2) the number of entities capable of being detected and categorized simultaneously, and 3) the slant range distance at which targets can be accurately detected and categorized. Our initial research questions addressed the capabilities of the object recognition algorithm.

Detection was defined as a binary yes or no that the software successfully registered that an entity was present. Categorization was defined as accurately classifying the entity as a

[69] Alberts et. al., *Code of Best Practice for Experimentation,* 68–73.

weapon, person, or vehicle. Accuracy was measured as a function of the number of entities (individuals, weapons, and vehicles) correctly detected and categorized versus the number of actual entities present. These measurements were taken at the optimal performance altitude based on the pixel size of the entity and the training data used to recognize the entity. The number of entities capable of being accurately detected and categorized was intended to stress the SPOTR software and determine its limits. Similarly, we manipulated the distance between the drone and the simulated entities to identify the range limits at which the platform could accurately detect and categorize. Simulated environments were used to examine the functionality of this software while acknowledging that real-world performance may vary significantly. Table 3 defines our technical requirements compared to the RPUAS requirements discussed in Chapter II. Our distances are defined as the slant range distance between the sUAS and the entity as we assume SOF teams will avoid operating a sUAS directly overhead to prevent compromising their position. Instead, we assess that a sUAS will operate at a minimum offset distance away from the team.

Table 3. SATLAS and RPUAS Requirements[70]

<table>
<tr><th>Requirement Area</th><th>RPUAS Performance Requirements</th><th>SATLAS Performance Requirements*</th></tr>
<tr><td>Detection Accuracy</td><td>-90% PD:<table><tr><td></td><td>Day</td><td>Night</td></tr><tr><td>Person</td><td>300m</td><td>200m</td></tr><tr><td>Vehicle</td><td>400m</td><td>300m</td></tr></table></td><td>-90% PD:<table><tr><td></td><td>Day</td><td>Night</td></tr><tr><td>Person</td><td>300m</td><td>200m</td></tr><tr><td>Vehicle</td><td>400m</td><td>300m</td></tr><tr><td>Weapon</td><td>200m</td><td>150m</td></tr></table></td></tr>
<tr><td>Categorization Accuracy</td><td>N/A</td><td>-90% PC:<table><tr><td></td><td>Day</td><td>Night</td></tr><tr><td>Person</td><td>100m</td><td>75m</td></tr><tr><td>Vehicle</td><td>200m</td><td>150m</td></tr><tr><td>Weapon</td><td>75m</td><td>50m</td></tr></table></td></tr>
<tr><td>Number of Entities Detected/ Categorized</td><td>N/A</td><td>Squad-sized element (approximately 10–12) detected</td></tr>
<tr><td>Noise</td><td>Inaudible at 200 feet AGL</td><td>Inaudible at 200 feet AGL</td></tr>
<tr><td>Flight Range</td><td>3 km</td><td>3 km</td></tr>
<tr><td>Time of Flight</td><td>30-minutes</td><td>30-minutes</td></tr>
</table>

*This was measured for temperate zones. Additional testing will be required in varying environments.

[70] Department of Defense, *Capability Production Document For Rucksack Portable Unmanned Aircraft System (RPUAS) Increment 2*.

a. Detection

We determined that the RPUAS accuracy requirements for personnel and vehicle detection were sufficient to support SOF operational requirements and, therefore, sufficient for SATLAS experimentation criteria. The detection of personnel at 300 meters places them outside the range of typical small arms fire and allows sufficient reaction time for a team to maneuver the sUAS closer to investigate or take appropriate action. Likewise, detection of vehicles at 400 meters increases standoff distance and allows a SOF team to react accordingly. As discussed in Chapter II, smaller targets create less pixels for the object recognition software to process and are more difficult to detect. While weapon detection at 200 meters places SOF teams within small arms range, it is more realistic given the distance requirements for personnel and vehicle entities. Additionally, the accuracy requirement of a 90% probability of detection (PD) establishes a reliable detection model that will create trust in the system.

b. Categorization

The U.S. Army did not factor in object categorization since it was not a RPUAS technical requirement; however, it is exceptionally important for our project due to its potential to increase situational awareness for SOF teams. We defined the minimum categorization distance for vehicles to 200 meters during day at 150 meters at night with a 90% probability of categorization (PC). Personnel should be categorized at 100 meters during the day and 75 meters at night with a 90% PC and weapons should be categorized at 75 meters during the day and at 50 meters at night with a 90% PC. Categorization requires multiple detections before the SPOTR software is confident enough to label an entity as a vehicle, person, or weapon, thus these distances will be much closer than detection distances. Detection at the specified distances will provide sufficient standoff for SOF teams to maneuver the sUAS to a closer distance to enable successful categorization.

c. Tracks

The U.S. Army also did not define a minimum number of entities that the object recognition software should be able to be detect and track simultaneously. This is also important to SATLAS because we assess that simultaneously detecting 10 entities would provide sufficient situational awareness to deployed SOF teams. A key consideration will be to balance the frame rate for the object recognition software against additional power requirements for the platform.

d. Noise

We determined that the RPUAS noise requirements were sufficient to support SOF operational requirements and sufficient for SATLAS experimentation criteria. Remaining inaudible at 200 feet AGL enables a sUAS to remain outside of the required detection range for both personnel and vehicles. This provides sufficient standoff distance and will mitigate the extent to which a sUAS could compromise a SOF team's location.

e. Flight Range

We determined that the RPUAS flight range requirement was sufficient for SOF operational requirements as well as for SATLAS technical parameters. Flight range remains dependent on a direct line-of-sight between the ground control station and a sUAS. Current state of the art sUAS technology is mostly limited to a 3-kilometer flight range due to the increased signal strength required to operate outside of this range and the complexities involved in operating a sUAS beyond-line-of-sight. Exceeding a 3-kilometer flight range requirement will severely limit the available sUAS.

Additionally, a 3-kilometer range would provide sufficient standoff for SOF teams. At an average walking speed of 3 miles per hour, this range could provide a SOF team nearly 40 minutes to react if the personnel are detected at the maximum flight range of the sUAS. Detecting a vehicle traveling at 20 miles per hour at the maximum flight range will provide 5 minutes to react. Maximum reaction time will often be preferred; however, we assess that these times will provide a sufficient reaction time for SOF teams. This will need to be evaluated in future field experiments.

f. Time of Flight

We also concur with the RPUAS requirement of a 30-minute time of flight. Like flight range, the current state of the art sUAS technology is mostly limited to a 30–35-minute time of flight. Integrating object recognition software and a robust camera into the sUAS platform is expected to increase the power load and decrease flight time. Therefore, increasing this requirement would severely degrade available sUAS and require larger platforms with more battery life which become too big to remain man packable.

There are, however, important implications sUAS flight time will have on SOF operations. Considering the vignette in Chapter II, many SOF operations and ISR requirements exceed 30 minutes. SOF teams will be required to select the most desirable times for organic ISR coverage and carry multiple batteries and charging stations should they use these platforms. This will increase the logistical burden of teams until the technology advances beyond a 30-minute time of flight for VTOL sUAS.

5. Relationship Between Variables

We expected to discover several relationships among the variables. Within the criteria variables, the accuracy, maximum number of entities that can be detected, and maximum range at which accurate detection can occur are all dependent on the onboard camera system. The object recognition algorithms rely on a pixel size ratio to correctly categorize entities as individuals, vehicles, or weapons. And as each sUAS manufacturer uses different cameras, SPOTR could be expected to perform better on certain sUAS platforms than others. Finally, we expected SPOTR to perform better at test sites that had terrain more aligned with the training data used to develop the algorithms. We expect the data from our research to necessitate additional feedback into Pillar 1. The next section describes the phases of our experiments.

B. EXPERIMENT PHASES

The tests in each phase were designed to evaluate SPOTR performance and identify needed adjustments prior to completion of the software integration onto the platform. In Phase 0 we participated in a simulation validation with AeroVironment engineers to determine the limitations of the simulation and better define our research questions and measures of performance. In Phase I we purchased an Alienware laptop that we loaded with a version of SPOTR and a simulated environment to conduct our own experiments in a completely simulated environment. The algorithms were tested using virtual sUAS platforms hovering over the simulated training site and all data processing was done through the laptop. We then moved on to Phase II in which the data processing was done on a Nibbler sUAS platform with an onboard NVIDA Jetson TX2 microprocessor while operating over the same virtual training site. This represented a more constrained processing environment that could better simulate how the software would perform in a real-world scenario. Phase III then transitioned into field testing the Nibbler sUAS over Camp Snyder while detecting role players.

The experiment concluded with Phase III testing and preparation for subsequent Phase IV testing for future student teams. The next sections describe in detail the testing performed for Phases 0 – III as well as the scenarios that were used in Phases II and III.

1. Phase 0 – Simulation Software Validation

a. Setup

Phase 0 began with months of coordination with AeroVironment's engineer team and culminated with a virtual SPOTR familiarization on April 23, 2021. The software verification used a simulated environment resembling the urban training site in Quantico, VA. Two AeroVironment engineers executed the demonstration with the NPS student team and faculty advisors through Zoom.

As seen in Figure 7, the simulation replicates a real-world urban training site to validate the SPOTR software. It primarily used light urban areas and open fields to employ SPOTR. The simulation used two ArduCopter sUAS to scan the designated areas (Figure 8). Additionally, the standard simulation setup had limited the number of entities for SPOTR to recognize to 12 individuals and 1 vehicle. Their location and movements were controlled by the simulation were unable to be manipulated under normal circumstances.

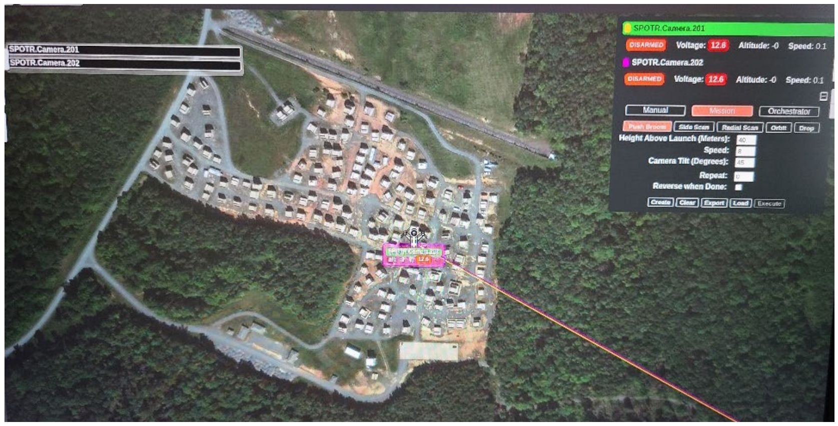

Figure 6. Aerial View of Simulated Environment

Figure 7. ArduCopter sUAS Used in Simulated Environment

The AeroVironment engineers controlled the movement of the ArduCopters and maneuvered them through the simulated environment to demonstrate the capabilities of the simulation and SPOTR. We did not approach Phase 0 with any assumptions or research questions. We had considered measurements that may be beneficial to collect, including accuracy, number of entities detected, distance of detection, and altitude of detection; however, our intent was to determine whether the simulation provided sufficient feedback to effectively evaluate SPOTR during future phases.

b. Execution7

The AeroVironment engineers initially launched the two simulated ArduCopters to demonstrate the control mechanisms, explain the view panels, and display the SPOTR software in action. We were able to observe the overview imagery and the SPOTR-enabled camera imagery (Figure 9). The left screen was used to observe the location of the ArduCopters in relation to the imagery while the right screen displayed the footage

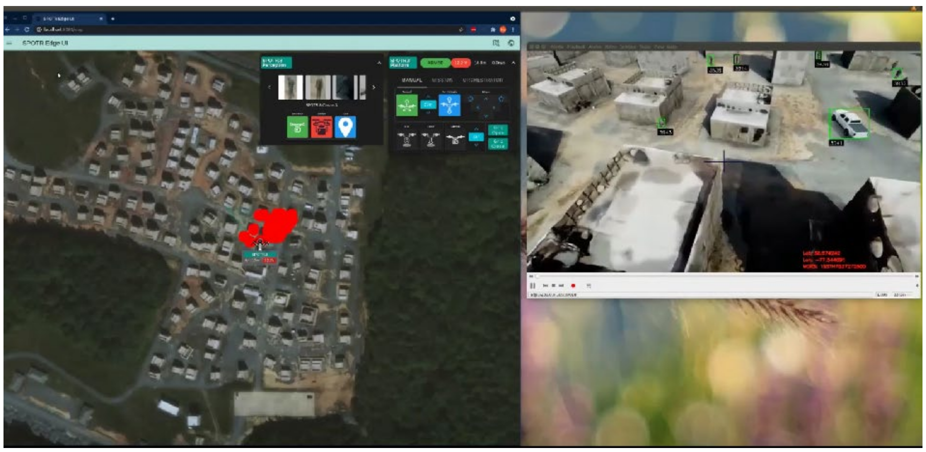

Figure 8. April 23 Demonstration Footage

- through the lens of the sUAS. The ArduCopters were brought to a hover at an altitude of 15 meters AGL to demonstrate the ability of SPOTR to recognize individuals and vehicles in the simulated environment. Highlighted green boxes indicate the detection of an individual or vehicle. The ArduCopters remained over the light urban area while individuals and vehicles moved throughout. We gradually increased the altitude in 5 meter increments up to 40 meters AGL to determine the distance at which SPOTR becomes unreliable.

c. Findings

We determined that the simulation would serve as a sufficient method to evaluate SPOTR during future experiments. It provided an acceptable method to operate SPOTR in a simulated urban environment. SPOTR was able to detect and categorize vehicles and personnel in the simulation despite being developed using live training data. Additionally, the ArduCopters could be maneuvered throughout the simulation to allow SPOTR to scan areas that offered different background contrasts. The simulation does have several limitations, however. The standard configuration limits the number of entities available for detection. Throughout most zones of the simulation, entities were limited to 8–13 personnel and one vehicle. This is a limitation that was unavailable for this thesis. Additionally, the simulation only includes urban terrain and open areas and does not have the ability to generate entities within vegetated regions. Based on the findings in Phase 0, our proposed research questions for Phase I were:

- Can the Nibbler sUAS equipped with integrated SPOTR software accurately recognize and categorize personnel and vehicles as potential threats?
- How many entities can be detected and categorized simultaneously given the current processing power?
- At what slant range can the software reliably recognize entities?

As seen in Figure 10, slant range distance factors into both sUAS altitude and lateral distance from the target and is defined as line-of-sight distance between two points.

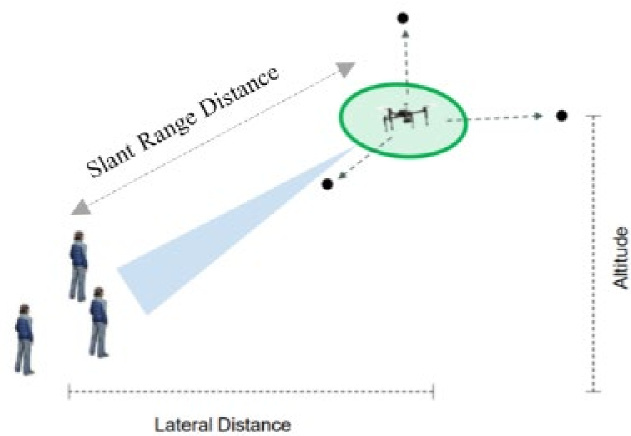

Figure 9. Slant-Range Distance[71]

Phase 0 concluded with revised research questions but confirmed the viability of testing SPOTR functionality in a simulated environment. Given the delays in the software integration, the ability to test SPOTR software in a synthetic environment proved beneficial. This helped us conclude that the simulation software was a viable testing platform to evaluate SPOTR, develop measurable research questions, and identify collectable measures of performance.

2. Phase I – SPOTR Software Familiarization

a. Setup

Phase I began with the purchase of an Alienware Area 51M laptop and the installation of the SPOTR software and virtual environment. We received the laptop in February 2021; however, the findings from Phase 0 enabled us to better tailor our experiments with regards to our initial research questions.

The software installed on the laptop has the same specifications as that observed during the software verification on April 23, 2021. The simulation is loaded with SPOTR's personnel and vehicle detection model and was, therefore, unable to recognize weapons. Having unrestricted access to the software allowed us to conduct multiple experiments to

[71] Joe Kehoe, email message to author, May 7, 2021.

test the capabilities and limitations of the software while avoiding coordination with outside organizations. Our assumptions included:

- Assumption 1: The virtual environment adequately simulates a real-world environment.

- Assumption 2: The virtual entities adequately simulate personnel and vehicles.

Additionally, our initial research questions remained consistent with the findings in Phase 0 and included:

- Can the Nibbler sUAS equipped with integrated SPOTR software accurately detect and categorize personnel and vehicles as potential threats?

- How many entities can be detected and categorized simultaneously in a constrained processing environment?

- At what slant range distance can the software reliably recognize entities?

b. *Execution*

Phase I experiments were conducted numerous times as we became more familiar with the software; however, our research questions remained the same. To determine how accurately the SPOTR software could detect and categorize entities, the ArduCopters were launched and hovered at the designated optimal slant range distance of 20 meters. We then counted the number of personnel and vehicles that had been observed compared to the number of those detected, and those correctly categorized. This was repeated 5 times to ensure reliability. Each iteration was recorded.

Our second research question was answered by hovering the ArduCopter at the optimal slant range distance for detection (20 meters) and measuring the number of entities detected and categorized simultaneously. Different zones were programmed with varying numbers of entities to detect; therefore, we maneuvered the ArduCopter through each zone to ensure reliability. Each zone was recorded.

Our third research question was answered using slant range to measure the distance between the ArduCopter and each entity being detected to determine the distance where detection accuracy begins to degrade. Slant range accounts for both sUAS altitude and the ground distance from the sUAS to the entity. Scans began at the optimal slant range distance of 20 meters and increased by five meters until reaching the maximum altitude of 60 meters available on the simulation. The ArduCopter remained in position while the number of targets detected was compared to the number of targets observed to determine the accuracy. This was repeated five times and the average detection percentage was recorded. As seen in Figure 11, the SPOTR software indicates a detection by highlighting the entity with a white square while items categorized are highlighted with a green square.

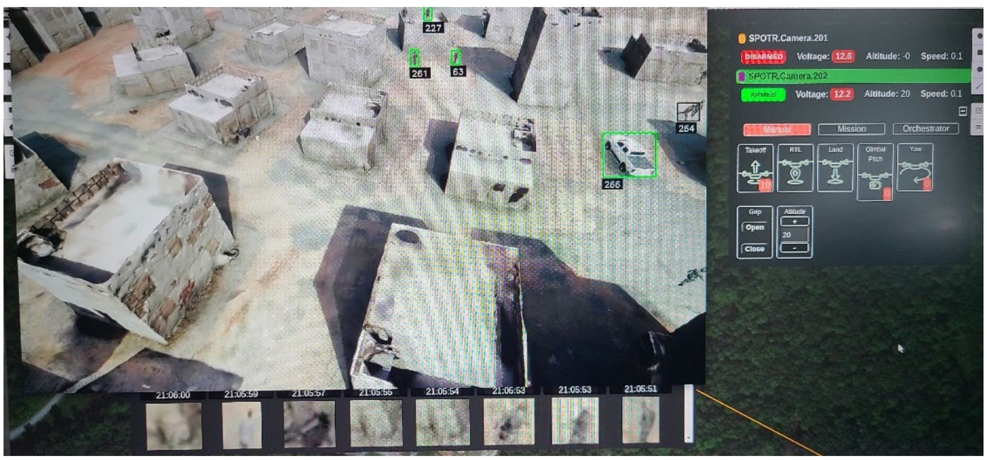

Figure 10. SPOTR Simulation View

c. *Findings*

For each launch to 20 meters AGL, SPOTR was able to detect 100% of the observable entities. This supported first research question that, at the optimal slant range distance, SPOTR is able to accurately recognize and categorize personnel and vehicles in a simulated environment.

To examine our second research question, we maneuvered the ArduCopter between zones to find its limit. Software restrictions did not allow us to fully explore this capability as we were unable to manually increase the number of detectable entities; however, we determined that the software could detect up to 11 entities and categorize 8 simultaneously. SPOTR was not limited in the number of entities it can detect but it is currently limited to 8 simultaneous categorizations. This limit allows the system to maintain a sufficiently high frame rate for additional processing and is a restraint placed on the system by AeroVironment engineers. Future testing is planned by AeroVironment to remove this internal limitation and evaluate the limits of the software, but the current parameters are sufficient for SATLAS.

To measure our third research question, we began at the optimal slant range for detection (10 meters for weapons and 20 meters for personnel and vehicles) and increased the distance by 5-meter increments. The number of entities detected and categorized was compared to the number of entities present and recorded as a percentage. Interviews with AeroVironment engineers indicated that accurate detection would require an approximately 25–35 pixels per meter. This equated to SPOTR being able to detect weapons at approximately 10 meters – 15 meters slant range, personnel could be detected out to 40 meters slant range, and vehicles were detected out to 60 meters.

Figure 12 summarizes the findings for our third research question regarding the maximum slant range at which SPOTR is capable of recognizing personnel and vehicle entities. SPOTR was able to detect 100% of the entities up to a slant range of 40 meters. At 50 meters, the recognition accuracy dropped slightly to 90% while entities at 60 meters were only recognized at 44% accuracy. False positives became evident at 60 meters and buildings began being categorized as vehicles. Entities at 70 meters were recognized with a 33% accuracy which then dropped to 22% at 80 meters. From 80 meters to 100 meters slant range, SPOTR was able to detect entities with a 22% accuracy.

Figure 11. SPOTR Performance (Phase 1)

According to AeroVironment engineers, the simulated environment is primarily intended to test the functionality of the threat detection software. It did, however, enable us to measure initial performance parameters during prototype development despite caveats that real-world performance will likely vary as the training data used to tune the algorithms are better tuned to actual environments. Additionally, Phase I experiments allowed us to measure the performance of the algorithms in an unconstrained processing environment prior to moving to Phase II. Given our findings, we requested for AeroVironment to develop a method to evaluate SPOTR performance while constrained by the processing power of a NVIDIA TX2 microprocessor.

3. Phase II – Discovery Experiments: Simulation

a. Setup

Phase II experiments were conducted on May 7, 2021 and used Zoom to enable screen sharing and communication between AeroVironment engineers in Virginia and NPS students and faculty located on campus at NPS. The experiment was conducted employing the same virtual environment replicating the Urban Training Center in Quantico, VA. The primary difference for Phase II testing included the use of a hardware-in-the-loop set up to constrain the processing power of SPOTR. As seen in Figure 13, a NVIDIA Jetson TX2 microprocessor mounted on a Nibbler was used to conduct the SPOTR processing while the virtual ArduCopter's maneuvered through the simulated environment. Additionally,

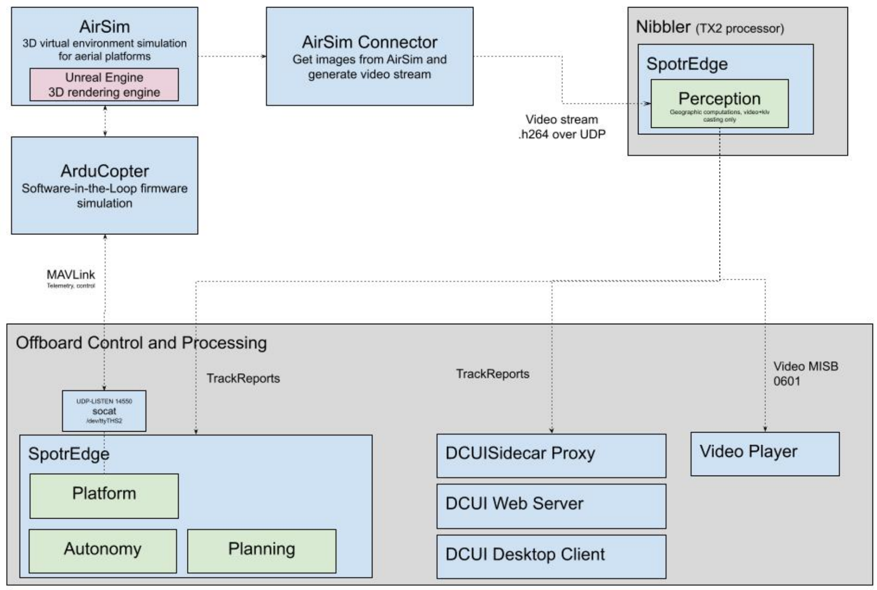

Figure 12. Phase II: Hardware-in-the-Loop Setup[72]

SPOTR's personnel and weapon model was used instead of the personnel and vehicle model employed in Phase I. As seen in Figure 14, the simulated environment was also broken down into 8 zones with varying numbers of targets available in each.

[72]Joe Kehoe, email message to author, May 7, 2021.

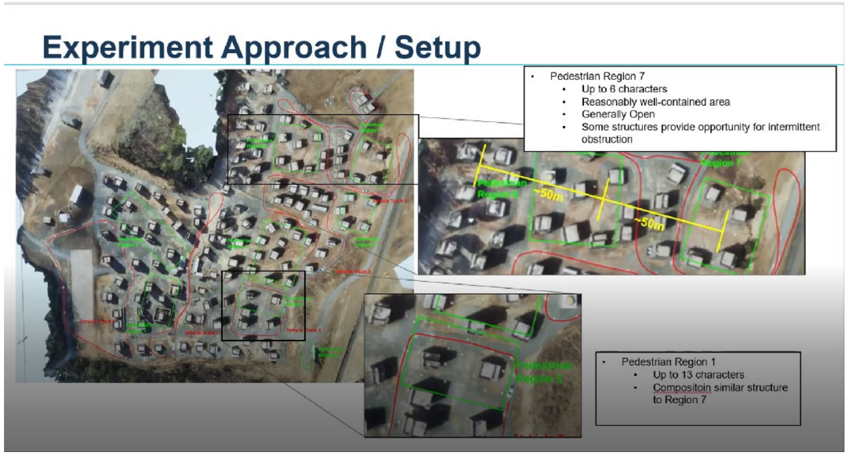

Figure 13. Simulated Environment Zones Used during Phase II Testing73

Incorporating the hardware-in-the-loop design did not alter our assumptions for this phase. We did, however, adjust Assumption 2 to include personnel and weapon entities while removing vehicles due to the change in detection model. The new assumptions were:

- Assumption 1: The virtual environment adequately simulates a real-world environment.
- Assumption 2: The virtual entities adequately simulate personnel and weapons.

Also, because of the varying detection model and the limited processing power, we refined our research questions to included:

- Can the Nibbler sUAS equipped with integrated SPOTR software accurately recognize and categorize personnel and weapons as potential threats?
- How many entities can be detected and categorized simultaneously in a constrained processing environment?

[73] Joe Kehoe, email message to author, May 7, 2021.

- At what slant range distance can the software reliably recognize entities?

b. Execution

Phase II testing was conducted in two separate zones within the simulation. Testing began in Zone 7 which was programmed to include 6 personnel targets all carrying weapons. To answer Research Question 1, the ArduCopters were launched and hovered at the optimal slant range distance of 20 meters. We then counted the number of personnel observed and compared this to the number of those detected as well as those correctly categorized. The ArduCopters were then lowered to an altitude of 10 meters to examine the weapons detection abilities of SPOTR. This was repeated twice for Zone 7.

We continued our experiment in Zone 7 to answer Research Question 2 by adjusting the hovering altitude from 15 meters AGL to 40 meters AGL. The ArduCopter was paused at every 5-meter increase in elevation and the number of personnel detected was compared to the number of personnel targets observed. Figure 15 displays a test iteration of SPOTR in Zone 7 at 20 meters slant range distance. After answering our two research questions in Zone 7, the ArduCopters were moved to Zone 1 to repeat the process.

Figure 14. Phase II Test of Zone 7

Zone 1 was programmed to include 13 personnel targets carrying weapons. To answer Research Question 1, the ArduCopters were brought to a hover at 20 meters altitude to examine the number of personnel targets detected. The ArduCopters were again lowered to 10 meters altitude to measure the number of weapons detected. This was repeated twice for Zone 1.

We addressed Research Question 2 by adjusting the hovering altitude of the ArduCopters between 15 meters AGL and 60 meters AGL. The ArduCopters were halted at every 5-meter increase and the number of personnel targets detected was compared to those observed on the screen.

c. Findings

For each launch to 20 meters AGL in both Zone 7 and Zone 1, SPOTR was able to detect and categorize 100% of the observable personnel. This provided the data for our first research question to assess that, while operating at the optimal slant range distance, SPOTR can accurately recognize and categorize personnel in a simulated environment. However, in Zone 7, SPOTR was unable to detect any weapons despite 6 being present at any given time. Additionally, in Zone 1, SPOTR was only able to detect 2 weapons and correctly categorize none of the weapons while up to 13 were present.

For our second research question, we found that SPOTR was able to simultaneously detect all observable personnel entities while operating at the maximum slant range distance of 20 meters. SPOTR detected all 6 personnel entities in Zone 7 and all 13 personnel entities in Zone 1 when they were within line of sight.

Figures 16 and 17 summarize our findings for our third research question regarding the maximum slant range at which SPOTR can recognize personnel entities. In Zone 7, SPOTR was able to detect 100% of the entities a slant range of 20 meters. At 25 meters, the recognition accuracy dropped to 83%. Recognition accuracy dropped to 67% for entities at 30 meters and 35 meters while recognition accuracy dropped to 0% for entities at 40 meters and beyond. In Zone 1, SPOTR was able to detect 100% of the personnel entities between 20 meters and 30 meters. At 40 and 45 meters, the recognition accuracy degraded to 77%. Recognition accuracy declined to 62% at 50 meters and 15% at 60 meters.

Figure 15. SPOTR Performance for Phase II: Zone 7

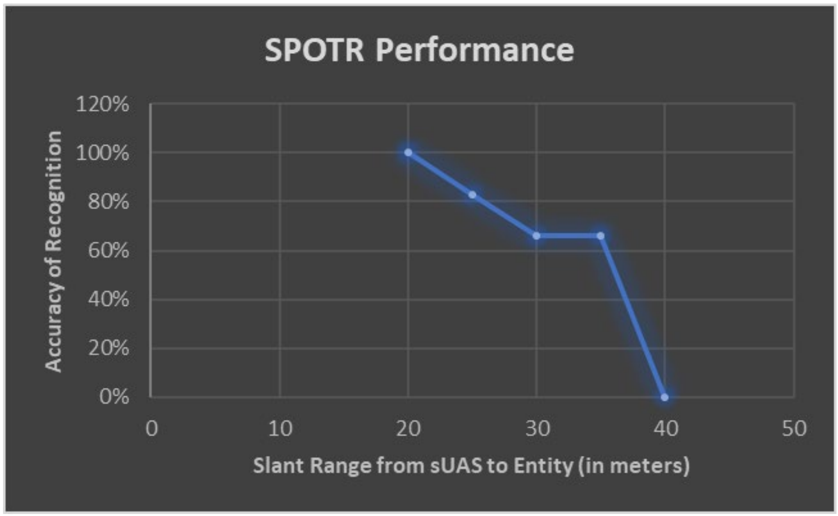

Figure 16. SPOTR Performance for Phase II: Zone 1

Again, the simulated environment is primarily intended to test the functionality of the threat detection software and may not replicate real-world performance. This experiment did, however, provide us another opportunity to measure initial performance capabilities while incorporating a hardware-in-the-loop processing. During the Phase II experiments we measured the performance of the algorithms in a constrained processing environment prior to moving to Phase III.

4. Phase III – Discovery Experiments: Field

a. Setup

Phase III experiments were conducted on May 18, 2021 via Zoom to enable screen sharing and communication between AeroVironment engineers in Virginia and NPS students and faculty. This experiment served as our first field testing of a SPOTR-enabled sUAS. We employed a Nibbler sUAS as the surrogate platform while production development continued on our Altavian M440 prototypes. As seen in Figure 18, the Nibbler contained an internal NVIDIA Jetson TX2 microprocessor; however, the object recognition software was conducted on a separate TX2 microprocessor at the ground control station. Thus, similar to Phase II experiments, Phase III testing constrained the processing power to that available on the TX2 microprocessor.

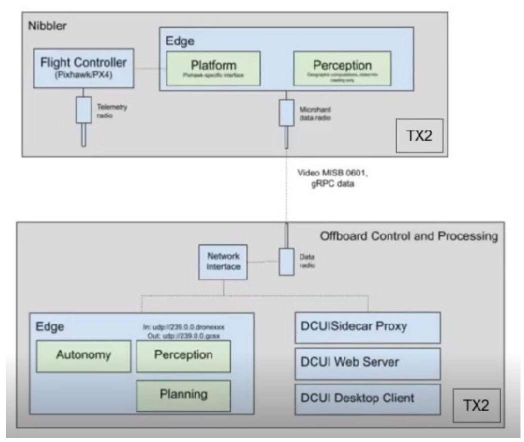

Figure 17. Phase III: Live Testing Setup74

74 Joe Kehoe, email message to author, May 18, 2021.

The experiment was conducted at the William B. Snyder Boy Scout camp in Haymarket, Virginia. As seen in Figure 19, the camp was divided into three zones to evaluate SPOTR's performance in varying environments. Zone 1 included a gravel field that produced significant contrast between the targets and the background. Zone 2 introduced more vegetation that included scrub brush, waist-high brush, and fully grown trees. Zone 3 included a single large building, a grass field, and intermittent trees. For this experiment, AeroVironment was only able to provide 3 individuals to serve as genuine entities. The individuals carried fake weapons that simulated a standard M4 rifle while a single pick-up truck was available for use. These tools allowed us to test both the personnel-weapon detection model and the vehicle-personnel detection model.

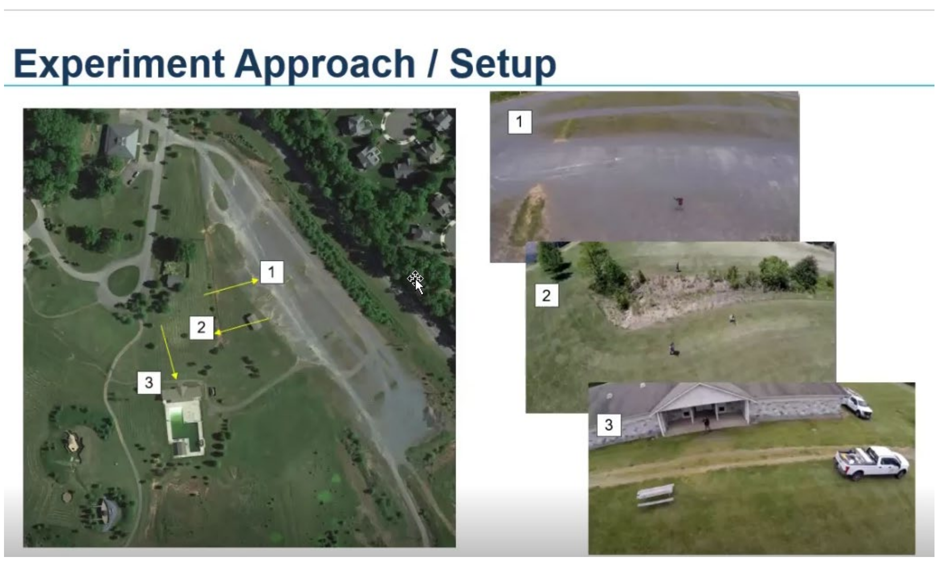

Figure 18. Phase III Testing Zones75

Our assumptions for Phase III were altered since we were able to transition to field testing. Our assumptions for this include:

- Assumption 1: The surrogate platform (Nibbler sUAS) will reasonably replicate the performance capabilities of our Altavian M440 prototypes.

75 Joe Kehoe, email message to author, May 18, 2021.

- Assumption 2: The field environment at Camp Snyder adequately replicates a representative operational environment for SOF. teams.

Our research questions also included:

- Can the Nibbler sUAS equipped with integrated SPOTR software accurately detect and categorize personnel and weapons as potential threats?
- How many entities can be detected and categorized simultaneously in a constrained processing environment?
- At what slant range distance can the software reliably recognize entities?

Due to limitations, we were unable to examine the maximum number of entities that SPOTR is capable of detecting and categorizing during this phase.

b. Execution

Phase III testing was conducted in Zones 1–3 at Camp Snyder. We began in Zone 1 which included a gravel field and 3 personnel targets, 2 of whom were carrying fake weapons. To answer Research Question 1 in Zone 1, the Nibbler was launched and hovered at the optimal slant range distance of 20 meters. We then counted the number of personnel present and compared this to the number of those detected as well as those correctly categorized. The Nibbler was then lowered to an altitude of 10 meters to examine the weapons detection abilities of SPOTR.

To answer Research Question 3 in Zone 1, we incrementally increased the hovering altitude of the Nibbler from 10 meters to 40 meters AGL. The Nibbler was paused at every 5-meter increase in altitude and the number of personnel detected was compared to the number of personnel targets observed. Weapon detection was measured; however, detection was not possible above 15 meters of altitude. After answering our two research questions in Zone 7, the Nibbler and our personnel targets were moved to Zone 2 to repeat the process.

Zone 2 included the same 3 personnel targets and 2 weapon targets but transitioned to a setting that included a grass field, scrub brush, and full-size trees. This setting allowed us to examine the impact of object obscuration on the detection models. Research Question 1 was answered by bringing the Nibbler to a hover at 20 meters altitude to examine the number of personnel targets detected. The Nibbler was then lowered to 10 meters altitude to measure the number of weapons detected.

We then moved to Research Question 3 by adjusting the hovering altitude of the Nibber from 15 meters AGL to 30 meters AGL. The Nibbler was halted every 5-meter increase and the number of personnel targets detected was compared to those observed on the screen. After recording our findings, the Nibbler and our targets were moved to Zone 3 to repeat the process.

Zone 3 included a single-story building, several smaller structures, and sparse trees. The three personnel targets were instructed to move around the structures as we maneuvered the Nibbler, in order to observe the impacts of man-made structures on object detection. Research Question 1 was answered by bringing the Nibbler to a hover at 20 meters altitude to examine the number of personnel targets detected. The Nibbler was then lowered to 10 meters altitude to measure the number of weapons detected.

We then moved to Research Question 3 by adjusting the hovering altitude of the Nibber from 15 meters AGL to 40 meters AGL. The Nibbler was halted at each 5-meter increase and the number of personnel targets detected was compared to those observed on the screen.

c. Findings

In each zone, SPOTR was able to detect and categorize 100% of the observable personnel targets while hovering at the optimal personnel detection altitude of 20 meters. Additionally, SPOTR was able to detect 100% of the weapon targets while operating at the optimal weapon detection altitude of 10 meters. This allowed us to confirm our first research question and determine that SPOTR can accurately detect and categorize personnel and weapons while operating at the optimal slant range distance for detection in a real-world environment.

Figures 20, 21, and 22 summarize our findings for our third research question regarding the maximum slant range that SPOTR can recognize personnel entities. In Zone 1, SPOTR's personnel and weapon algorithm was able to detect 100% of the personnel entities a slant range of up to 30 meters. At 35 meters, the recognition accuracy of personnel dropped to 66% and 0% at 40 meters. In Zone 2, SPOTR was able to detect 100% of the personnel entities at 20 meters. At 25 meters, SPOTR was able to detect and categorize 66% of the personnel targets but dropped to 0% at 30 meters and beyond. In Zone 3, SPOTR detected 100% of the personnel targets up to 25 meters slant range. Detection and categorization dropped to 66% from 30 meters to 35 meters slant range and dropped again to 33% at 40 meters. Phase III results varied from Phase II as we observed minimal detection absent of categorization. Put differently, SPOTR was able to categorize the vast majority of the entities that it detected; however, was unable to detect entities at greater distances as seen in Phase II.

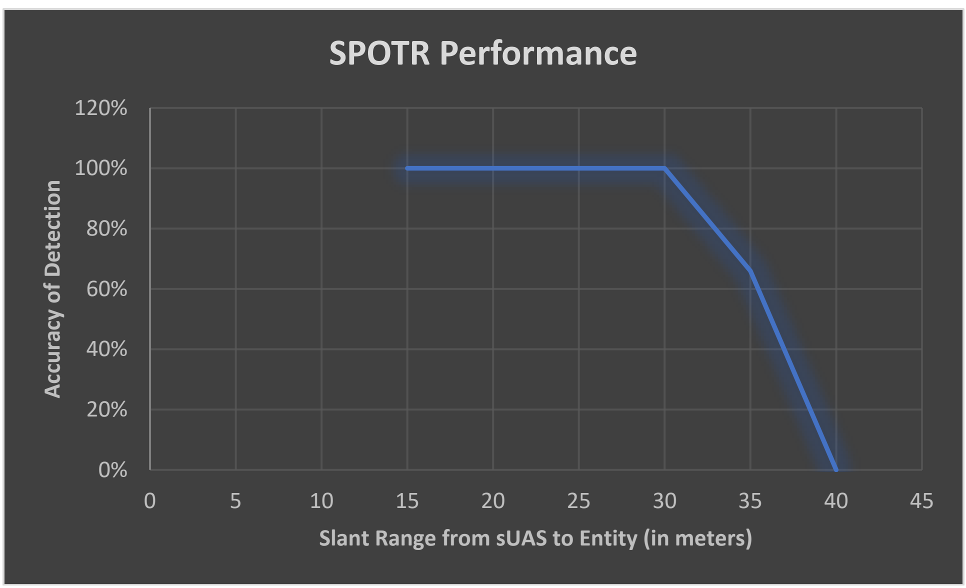

Figure 19. SPOTR Performance for Phase III: Zone 1

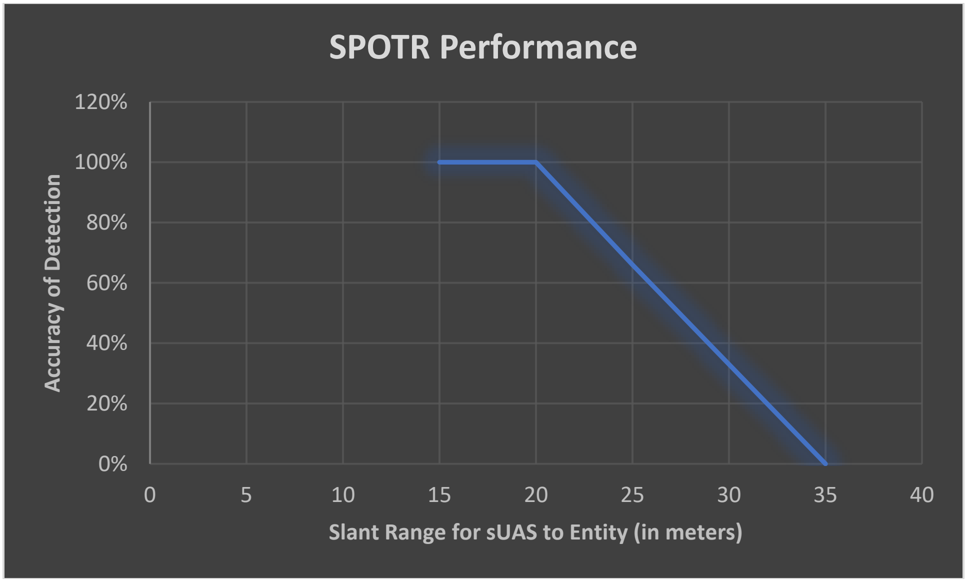

Figure 20. SPOTR Performance for Phase III: Zone 2

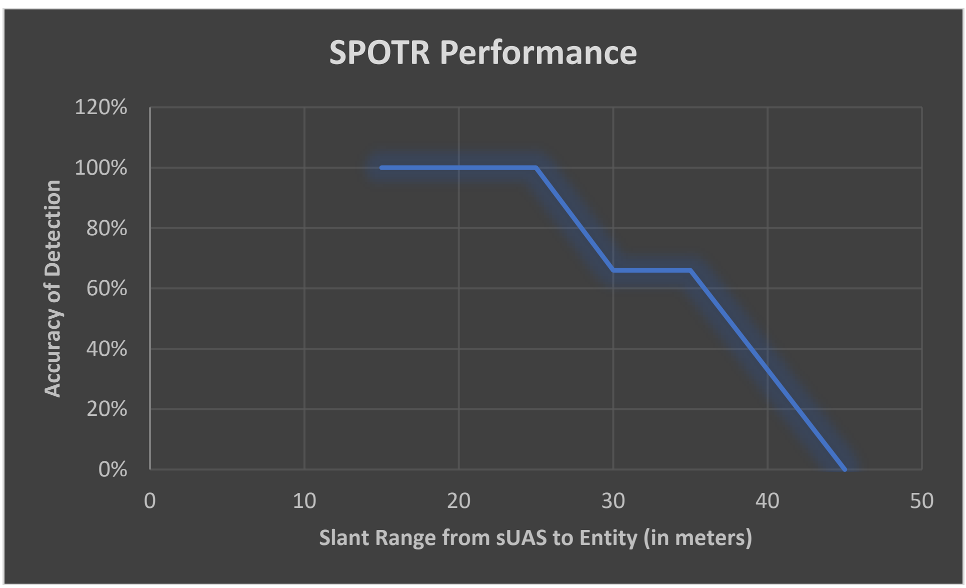

Figure 21. SPOTR Performance for Phase III: Zone 3

Phase III presented performance disparities between zones. Zones 1 and 3 which we presume is based on background contrast. The gravel field in Zone 1 provided significant contrast between the background and the personnel targets and enabled detection out to 40 meters. Similarly, the grass field and buildings in Zone 3 provided contrast that allowed the targets to stand out from the background. Personnel targets could obviously not be detected or categorized once they were obscured by buildings and performance was degraded once the personnel were covered by shadows; however, the background contrast did not appear to degrade performance.

We observed the most significant degradation of detection and categorization accuracy in Zone 2. The Nibbler was forced to drop to an altitude of 10 meters AGL to allow weapons to be accurately detected and the accuracy of personnel detection dropped once personnel targets entered the vegetation even when the Nibbler 20 meters from the target. Shadows also negatively impacted SPOTR's performance in personnel and weapons detection at the 20-meter range. SPOTR could accurately detect and categorize the personnel targets while they were maneuvering in the grass field surrounding the thicker vegetation; however, once they entered the scrub brush or denser vegetation, detection was severely degraded.

The field testing of SPOTR's detection and categorization supported its capability to detect and categorize personnel and weapons while operating at the optimal slant range distance for each object. Also demonstrated was the ability to detect and categorize entities out to a distance of approximately 40 meters when not obscured by vegetation, buildings, or shadows. The variances in performance, especially in Zone 2, reveals the significant impact that object obscuration can have accurate object detection models. Phase III, however, supported SPOTR's ability to process live sUAS camera footage and detect trained entities while operating in a constrained, field environment.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. ANALYSIS AND IMPLICATIONS

This section provides an analysis of the experiment results. We examine the results by research question within each phase of testing and compare the results to the technical and operational requirements for SATLAS discussed in Chapter III. Second, we analyze the implications of these results in the context of the SATLAS project and the ability to employ an object recognition-enabled sUAS within SOF teams.

A. ANALYSIS BY PHASE

1. Phase 0

Phase 0 was designed to align perspectives with regards to the simulation software between the NPS student team and advisors and the AeroVironment engineers. The objective was to assess whether it was feasible to use the simulation to evaluate the performance of the SPOTR software on a surrogate drone. Additionally, Phase 0 provided the SATLAS project team with necessary familiarization with SPOTR. Phase 0 also let us refine measurable research questions and measures of performance in order to manage the software development toward a prototype. We identified data that would be necessary to collect in order to establish performance parameters, including 1) accuracy and recall, 2) quantity of entities detected, 3) distance of detection, and 4) altitude of detection. The following section discusses each of these measurements.

a. Accuracy and Recall

Accuracy was our primary requirement for SPOTR as any object recognition software must be able to detect trained entities accurately and reliably to be of use. Accuracy measurements were not recorded during Phase 0 because this was not part of the Phase 0 test plan; however, SPOTR did successfully detect and categorize individuals and vehicles in the synthetic environment. The algorithm for weapon detection was not used during Phase 0, therefore, we were unable to examine this capability. Despite this shortfall, this phase suggested that the simulation could be used to quantitatively measure SPOTR's accuracy.

b. *Quantity*

We also identified the quantity of entities detected as a possible measurement. This was intended to determine the limitations of the system to better design future tests and assess the feasibility of integrating SPOTR into SATLAS prototypes. We were able to confirm that SPOTR was capable of detecting and categorizing multiple entities simultaneously while operating in the virtual environment.

We did, however, discover that the simulation had internal restrictions regarding the number of entities it could provide. Different zones within the simulation were programmed to provide different target profiles and was limited to 8–10 personnel and one vehicle. This restriction was not easily adjustable and would require significant software engineer support to manipulate. Despite this shortfall, we assessed that the simulation provided a sufficient number of targets to gather initial measurements.

c. *Distance and Altitude*

Other concerns included the distance and altitude at which SPOTR could detect and categorize. Phase 0 showed us that separating lateral distance and altitude to measure the range of SPOTR was not the most accurate method. The simulation provided targets at various lateral distances; therefore, we combined these two research questions to measure the slant range distance between the sUAS and the entity. This allowed for more accurate and reliable measurements in future testing. We identified that SPOTR was able to detect entities at various slant range distances within the simulation. We were also able to manipulate the slant range distance by maneuvering the ArduCopters over the target region which allowed us to measure the slant range distance of detection within the simulation.

Phase 0 enabled us to determine that the simulation provided a sufficient environment to evaluate the SPOTR software. We were able to measure SPOTR's accuracy, the quantity of entities detected, and the slant range distance of detection to support future experiment phases.

2. Phase I

We addressed our three adjusted research questions from Phase 0 in the context of our stated operational requirements. We gathered initial measurements in the synthetic environment and determined the capabilities of the software. Each research question and the analysis and implications are discussed individually in the next section.

a. Research Question I-1

Research Question I-1 was designed to determine the ability of SPOTR to accurately detect and categorize vehicles and personnel. This research question considered only accuracy at the optimal slant range distance of 20 meters. Detection and categorization slant range distance were considered in Research Question 3. Additionally, we were unable to test SPOTR at night; therefore, these requirements are omitted from the below table. Table 4 displays the results of Phase I testing compared to the RPUAS and SATLAS technical requirements.

Table 4. Accuracy Requirements and Accuracy Performance: Phase I*

<table>
<tr><th></th><th>RPUAS Data</th><th>SATLAS Data</th><th>SPOTR Data</th></tr>
<tr><td>Detection Accuracy/Slant Range Distance</td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>300m</td></tr><tr><td>V</td><td>400m</td></tr></table></td>
<td>-90%:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>200m</td></tr><tr><td>P</td><td>300m</td></tr><tr><td>V</td><td>400m</td></tr></table></td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>60m</td></tr><tr><td>V</td><td>75m</td></tr></table></td></tr>
<tr><td>Categorization Accuracy</td>
<td>N/A</td>
<td>-90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>75m</td></tr><tr><td>P</td><td>100m</td></tr><tr><td>V</td><td>200m</td></tr></table></td>
<td>-90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>50m</td></tr><tr><td>V</td><td>75m</td></tr></table></td></tr>
</table>

*” P” indicates personnel. “V” indicates a vehicle. “W” indicates a weapon.

As discussed in Chapter II, accuracy is the measure of when the object recognition algorithm guesses how often it is correct, whereas recall is a measure of the how often the algorithm guesses when it should. We measured these two dependent variables and measured the percentage of entities correctly detected and categorized while the sUAS hovered at the optimal slant range distance. We evaluated the accuracy and recall by recording the percentage of correctly detected and categorized entities compared to the number of entities present, including false positives and false negatives. Our findings indicate that, given an optimal slant range distance and relatively unobscured line of sight to an entity, SPOTR should be able to detect and categorize weapons, personnel, and vehicles above the 90% accuracy requirement.

As seen in Table 4, SPOTR met both the detection and categorization requirements for RPUAS and SATLAS defined in Chapter III as it achieved over a 90% detection and categorization accuracy while operating at the optimal slant range distance in the simulation. Constraints of the simulation prevented the sUAS from operating above 60 meters altitude and thus limited the slant range to approximately 75 meters. While this limited our ability to determine the vehicle detection and categorization range, it did allow us to examine the limits of the personnel detection and categorization algorithm.

Phase I enabled us to conclude that the SPOTR software is accurate enough in the simulation mode to warrant further integration into the SATLAS project.

b. Research Question I-2

The goal of Research Question I-2 was to find the maximum quantity of entities SPOTR is capable of simultaneously detecting and categorizing. Phase I allowed us to collect this measurement while the system operated in an unconstrained processing environment. Table 5 displays the results of Research Question 2 during Phase I.

Table 5. Quantity of Entities Detected: Phase I

	RPUAS Data	SATLAS Data	SPOTR Data
Quantity of Entities Detected	N/A	10	10

The number of entities capable of being detected by SPOTR met the operational requirements of SATLAS. In this phase, the number of entities capable of being detected was limited by the simulation constraints that restricted the number of targets available to 10. This constraint did not allow us to measure the number of entities SPOTR is capable of detecting; however, we were able to observe it detecting all 10 entities at one time. The number of entities capable of being categorized was limited to 8 by a software constraint designed to maintain a higher frame rate for the algorithms. SPOTR was able to reach this maximum limitation of 8 simultaneous categorizations.

Research Question I-2 helped us to conclude that SPOTR should be able to simultaneously detect a sufficient number of entities while operating in an unconstrained and simulated environment.

c. Research Question I-3

Our third research question concerned the maximum slant range distance at which SPOTR could detect and categorize entities. As illustrated in Table 4, SPOTR did not meet the distance requirements for detection for RPUAS or SATLAS and did not meet the distance requirements for categorization for SATLAS. Detection of personnel-sized targets dropped significantly at slant range distances of greater than 60 meters while categorization began to drop when greater than 50 meters. Of note, however, the simulation replicated the wide-angle camera setting of a GoPro Hero 4. Additionally, SPOTR uses real-world training data and simulation performance is not exact to real-world performance.

Two potential explanations for SPOTR's performance are the differences between real-world training data and what the simulation can replicate as well as the wide-angle camera replicated in the simulation. As previously mentioned, the synthetic environment is generally designed to validate the performance of the detection algorithms before moving into live tests. Live training data will perform differently in the virtual setting. Additionally, as mentioned in Chapter II, camera capabilities and pixel size are important factors of object recognition results. The wide-angle cameras replicated did not provide the desired results in the simulated environment. Thus, prototypes will likely need to employ both wide and narrow field of view cameras to reach the technical requirements for both RPUAS and SATLAS. Software improvements, such as super-resolution, may also be required to allow object detection at the required ranges.

3. Phase II

Phase II testing was designed to test our research questions in the same synthetic environment while integrating a hardware-in-the-loop design to constrain processing capabilities to that of onboard future prototypes. We approached this phase with the same three research questions, gather additional measurements in the simulated environment, and determine the feasibility of on-the-edge processing for object recognition. Each research question and the analysis of the findings are discussed individually in the next section.

a. Research Question II-1

Research Question II-1 remained the ability of SPOTR to accurately detect and categorize personnel and weapons. The loaded algorithms for Phase II included personnel and weapon detection but not vehicle detection; thus, vehicle detection was not measured. Additionally, we were unable to test SPOTR at night; therefore, this requirement is omitted from the below table. Table 6 displays the results of Phase II testing compared to the RPUAS and SATLAS technical requirements as well as the results of SPOTR testing in Phase I.

Table 6. Accuracy Requirements and Accuracy Performance: Phase II*

<table>
<tr><th></th><th>RPUAS Data</th><th>SATLAS Data</th><th>SPOTR Data: Phase I</th><th>SPOTR Data: Phase II</th></tr>
<tr><td>Detection Accuracy/Slant Range Distance</td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>300m</td></tr></table></td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>200m</td></tr><tr><td>P</td><td>300m</td></tr><tr><td>V</td><td>400m</td></tr></table></td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>60m</td></tr><tr><td>V</td><td>75m</td></tr></table></td>
<td>-Zone 7: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>0m</td></tr><tr><td>P</td><td>20m</td></tr></table>-Zone 1: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>35m</td></tr></table></td></tr>
<tr><td>Categorization Accuracy</td>
<td>N/A</td>
<td>-90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>75m</td></tr><tr><td>P</td><td>100m</td></tr><tr><td>V</td><td>200m</td></tr></table></td>
<td>-90% PC<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>50m</td></tr><tr><td>V</td><td>75m</td></tr></table></td>
<td>-Zone 7: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>0m</td></tr><tr><td>P</td><td>20m</td></tr></table>-Zone 1: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>35m</td></tr></table></td></tr>
</table>

*”P” indicates personnel. “W” indicates a weapon. “V” indicated a vehicle.

SPOTR met the detection and categorization accuracy requirements for RPUAS and SATLAS as defined in Chapter III and maintained over 90% accuracy when operating at its optimal slant range distance. Interestingly, the constrained processing environment appeared to negatively impact object detection and categorization performance. Of additional interest was the performance disparity between zones during Phase II. SPOTR performance in Zone 1, while degraded from Phase I testing, was significantly better than its performance in Zone 7. Zone 1 contained a much darker background than Zone 7 which allowed the targets to standout. This finding suggests that SPOTR performance improves with increased contrast while operating in a simulated environment. While personnel detection and categorization accuracy met the requirements of RPUAS and SATLAS, Phase II demonstrated that accuracy may be negatively impacted by limited processing power and the lack of contrast between entities and the observed background. This may necessitate considerations for Pillar 1 platform modifications.

b. Research Question II-2

Research Question II-2 sought to find the maximum number of entities SPOTR is capable of simultaneously detecting and categorizing. Software engineers were able to slightly increase the number of targets to 13 in different zones to further stress SPOTR while operating in a constrained processing environment. Table 7 displays the results of Research Question 2 during Phase II as compared to the RPUAS and SATLAS requirements as well as Phase I.

Table 7. Quantity of Entities Detected: Phase II

	RPUAS Data	**SATLAS Data**	**SPOTR Data: Phase I**	**SPOTR Data: Phase II**
Number of Entities Detected	N/A	10	10	Zone 7: 13 Zone 1: 6

The quantity of entities detected by SPOTR met the technical requirements of SATLAS while this phase also demonstrated that detection rate was not negatively impacted by reduced processing power. The number of entities capable of being detected was limited to 13 in Zone 7 and 6 in Zone 1 and SPOTR was able to detect all available targets while operating at an optimal slant range distance of 20 meters. The number of entities capable of being categorized remained limited to 8 by the internal software constraints. SPOTR was able to reach this maximum limitation of 8 simultaneous categorizations while during this phase. From Phase II, we concluded that SPOTR is able to simultaneously detect an adequate number of entities while operating on limited processing power and in simulated environment.

c. Research Question II-3

The third research question regarded the maximum slant range distance at which SPOTR could detect and categorize entities. As seen in Table 6, SPOTR performance did not meet the distance requirements for detection for RPUAS or SATLAS and did not meet the distance requirements for categorization for SATLAS.

Phase II findings demonstrated a reduction in range for SPOTR while performing with reduced processing power. The distances at which SPOTR could detect and categorize entities was cut nearly in half between Phase I and Phase II. Phase II also demonstrated a clear impact on detection and categorization performance in different zones. The only observable difference between the two zones tested was that the backdrop in Zone 7 appeared much darker and had an increased number of shadows from the buildings. Given that the maximum distance for detection and categorization in Zone 1 was nearly double that in Zone 7, this may indicate that the size of an entity is as important as background contrast in a simulated environment.

Phase II suggested that increased processing power will be needed to improve SPOTR's detection range. The SATLAS project may need to seek alternatives to the NVIDIA TX2 that provide increased processing power Additionally, this phase suggested that cameras that are able to emphasize contrast between entities and the background (e.g., EO/IR) may also improve detection range and may be valuable avenues of future research.

4. Phase III

Phase III testing was our first opportunity to evaluate our research questions using a sUAS with integrated SPOTR in a field environment. Similar to Phase II, processing power was limited to that available on a NVIDIA TX2 microprocessor. We intended to evaluate our same three research questions; however, limited role players did not allow us to measure the maximum quantity of entities SPOTR could detect and categorize in a live environment. Phase III used a surrogate sUAS (Nibbler) to gather initial measurements in a field setting to determine the feasibility of on-the-edge processing for object recognition. Each research question and the analysis of our findings are discussed in the next sections.

a. Research Question III-1

Research Question III-1 was the ability of SPOTR to accurately detect and categorize personnel and weapons. Engineers were able to transition between the personnel and weapon detection model and the vehicle and personnel detection model; however, testing primarily focused on the former. Table 8 displays the results of Phase III testing for each zone compared to the RPUAS and SATLAS technical requirements as well as the results of SPOTR testing in Phase I and Phase II.

Table 8. Accuracy Requirements and Accuracy Performance: Phase III*

<table>
<tr><th></th><th>RPUAS Data</th><th>SATLAS Data</th><th>SPOTR Data: Phase I</th><th>SPOTR Data: Phase II</th><th>SPOTR Data: Phase III</th></tr>
<tr>
<td>Detection Accuracy/ Slant Range Distance</td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>300m</td></tr></table></td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>200m</td></tr><tr><td>P</td><td>300m</td></tr><tr><td>V</td><td>200m</td></tr></table></td>
<td>-90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>60m</td></tr><tr><td>V</td><td>75m</td></tr></table></td>
<td>-Zone 7: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>0m</td></tr><tr><td>P</td><td>20m</td></tr></table>-Zone 1: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>35m</td></tr></table></td>
<td>-Zone 1: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>15m</td></tr><tr><td>P</td><td>30m</td></tr></table>-Zone 2: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>25m</td></tr></table>-Zone 3: 90% PD:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>25m</td></tr></table></td>
</tr>
</table>

<table>
<tr><th></th><th>RPUAS Data</th><th>SATLAS Data</th><th>SPOTR Data: Phase I</th><th>SPOTR Data: Phase II</th><th>SPOTR Data: Phase III</th></tr>
<tr><td>Categorization Accuracy</td><td>N/A</td><td>-90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>75m</td></tr><tr><td>P</td><td>100m</td></tr><tr><td>V</td><td>200m</td></tr></table></td><td>-90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>P</td><td>50m</td></tr><tr><td>V</td><td>75m</td></tr></table></td><td>-Zone 7: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>0m</td></tr><tr><td>P</td><td>20m</td></tr></table>-Zone 1: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>35m</td></tr></table></td><td>-Zone 1: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>15m</td></tr><tr><td>P</td><td>30m</td></tr></table>-Zone 2: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>25m</td></tr></table>-Zone 3: 90% PC:<table><tr><td></td><td>Day</td></tr><tr><td>W</td><td>10m</td></tr><tr><td>P</td><td>25m</td></tr></table></td></tr>
</table>

* “P” indicates personnel. “W” indicates a weapon. “V” indicated a vehicle.

Field testing of SPOTR met the detection and categorization accuracy requirements for RPUAS and SATLAS as defined in Chapter III at the optimal slant range distance. Phase III demonstrated less disparity between the distance at which entities are detected versus the distance at which they are categorized. Essentially, once SPOTR was able to detect an object, it could categorize it at the same distance so long as the software was trained for the specific object. This phase also revealed very similar performance outcomes as those observed in the constrained processing environment during Phase II.

This finding reveals the impact of object obscuration on object recognition models while also demonstrating some current limitations of employing these systems in wooded operational environments. While personnel detection and categorization accuracy met the requirements of RPUAS or SATLAS at the optimal slant range distance, Phase III demonstrated the impact of obscuration on object detection accuracy and emphasized the need for improved camera systems to increase detection in varying environments.

b. Research Question III-2

Phase III did not measure the maximum number of entities SPOTR was capable of simultaneously detecting and categorizing as the experiment was limited to the three personnel and two weapons. The testing did reveal that at the optimal slant range distance, SPOTR could simultaneously detect the three personnel and two weapons; however, further testing will need to be conducted to determine the true limits of the system.

c. Research Question III-3

Our third research question considered the maximum slant range distance at which SPOTR could detect and categorize entities. As seen in Table 8, SPOTR did not meet the distance requirements for detection for RPUAS or SATLAS and did not meet the distance requirements for categorization for SATLAS.

The findings from Phase III revealed similar performance metrics to Phase II which suggests that the constrained processing power will reduce the slant range at which entities can be detected and categorized. Live testing also demonstrated SPOTR's ability to categorize at the same distance it could detect entities. In open areas, SPOTR could accurately detect and categorize personnel out to approximately 40 meters. The shadows and vegetation present in Zone 2 cut this range down to approximately 25 meters and demonstrated the severe impacts of object obscuration on object recognition.

Despite the inadequate slant range distance, there are various options to increase the range of detection for SPOTR or other object recognition models. As discussed in Chapter II, pixel size is dependent on the camera system. Adjusting the available payload in favor of dual-camera systems with wide-angle cameras and narrow-angle cameras is one possible solution. As seen in Figure 23, narrow-angle cameras with available zoom options will increase the available pixel size and increase the slant range for detection and categorization. A second alternative is the super resolution algorithm discussed by Borel and Young.[76] Processing camera feed in a super resolution algorithm also increases the available pixel size for object detection but without impacting the onboard hardware. Both techniques offer options to increase pixel size and improve detection range and performance in vegetation.

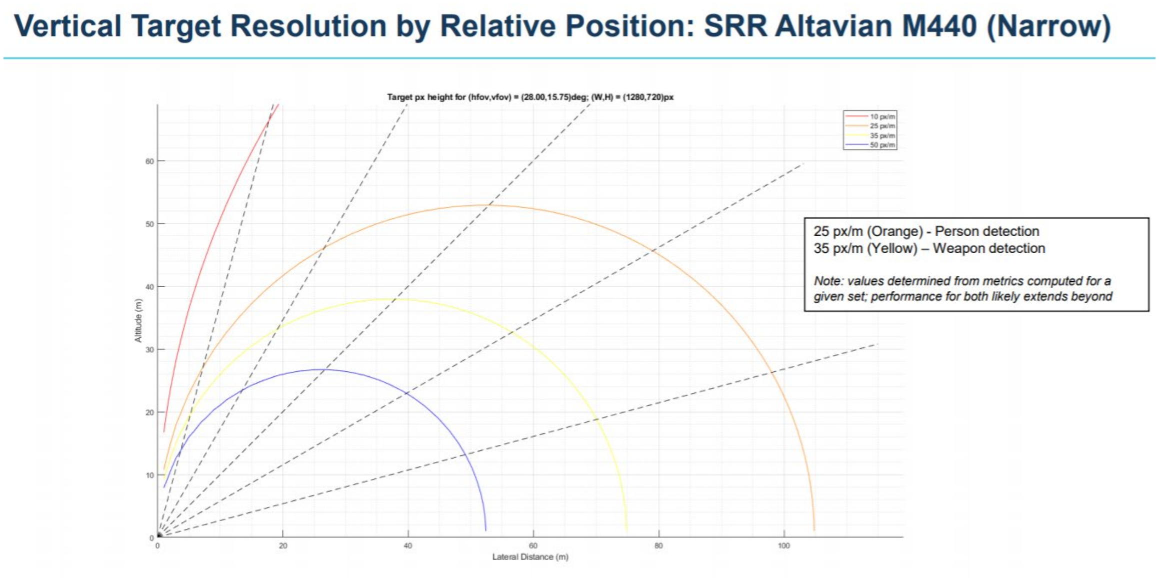

Figure 22. SPOTR Detection Range: Narrow-Camera77

B. OPERATIONAL IMPLICATIONS

This section provides an analysis of our three primary measurements: accuracy and recall, multiple target detection, and slant range. We consider the performance metrics achieved, the desired metrics, and the implications these will have for an SFOD-A in an operational context.

[76] Christoph Borel and S. Young, "Image Quality and Super Resolution Effects on Object Recognition Using Deep Neural Networks,," 7–8.

[77] Joe Kehoe, email message to author, May 7, 2021.

We approach this analysis with three primary assumptions: 1) that MDO will reflect the predominant operating environment of the future for SOF teams and that scenarios resembling our vignette in Chapter II, isolated and small teams requiring organic ISR support, will become increasingly common, 2) that MDO will limit the degree to which SOF teams will be able to rely on large, theater-level ISR platforms, and 3) operators will need to maneuver platforms within close proximity of potential targets for positive identification prior to taking lethal or non-lethal action. The following sections discuss our three categories with these in mind.

1. Accuracy and Recall

The implications of such a high degree of accuracy and recall are significant for SATLAS and SOF teams. The ability of an organic platform to accurately detect and categorize trained entities provides a unique situational awareness to a SOF team. Access to such a system will enable SOF teams to surveil areas of interest in their immediate proximity and detect and categorize moving entities in order to positively identify them. In the vignette, the ISIS militants maneuvered undetected within 75 meters of isolated SOF teams, ambushed, and killed U.S. Soldiers. Because the situational awareness of these SOF teams was degraded, small elements of enemy combatants maneuvered through the restrictive terrain and urban areas while taking advantage of security gaps. Operating in a MDO environment will only increase the impacts of diminished situational awareness on small and isolated SOF teams. Politically sensitive environments, reduced air superiority, and the use of clandestine techniques will exacerbate situational awareness shortfalls.[78] Object recognition-capable sUAS with a 90% accuracy would present a method to increase situational awareness for SOF teams in future MDO environments.

2. Quantity

SPOTR was able to simultaneously detect and categorize enough entities in the simulation to suggest its potential value for SOF teams. To achieve this goal, SPOTR does not need to be able to detect and categorize every trained entity in a given environment.

[78] Department of the Army, *Army Futures Command Concept for Special Operations 2028,* 12–25.

Instead, the detection and categorization of 10 trained entities could provide sufficient warning to a SOF team while avoiding cognitive overload. In the vignette in Chapter II, it was uncommon to encounter enemy combatants maneuvering in groups that included more than 10 personnel. While this will vary depending on the tactical environment, a human-in-the-loop system that alerts the operator to detections and enables the user to assess the entity as a threat or not provides the most value. Requiring SPOTR to simultaneously detect and categorize more than 10 entities would not improve its operational value and risks overwhelming the user with alerts in this specific context.

Perhaps more important than the accurate detection and categorization rate is the ability to toggle between the different detection modes and tailor the alerts based on the situation. While not an initial research question nor a measurement collected, the Phase III experiment suggested the importance of being able to easily adjust detection modes between personnel/weapon and vehicle/personnel models. This capability would allow SOF teams to collect the desirable information based on the operational environment while avoiding unnecessary detections and categorizations. For example, vehicle detection would have proved useless in the vignette in Chapter II due to the rugged terrain and inaccessible routes and the personnel/weapon detection model would have been most useful. This requirement will likely change based on the scenario in which SOF teams find themselves. The ability to tailor the specific detection model to the environment will reduce unnecessary information overload on SOF teams and allow them to collect only the needed information.

Another factor that deserves future attention is the ability to adjust the sensitivity of SPOTR based on the tactical environment. This consideration was beyond our scope; however, we were able to adjust the detection sensitivity during the Phase I experiment to reduce the number of false positives or false negatives. Depending on the operational environment, SOF teams may prefer to select a setting that makes only correct detections but may miss some detections or makes detections at every available opportunity and results in some inaccurate detections or is somewhere in the middle. This would enable the operator to tailor the number of detections received along with the accuracy requirement and control the cognitive burden of the system.

3. Slant Range

SPOTR did not meet the slant range distance requirements of RPUAS or SATLAS and should be the focus of future experiments or software adjustments. The current RPUAS and SATLAS requirements call for a system capable of detecting a weapon at 200 meters, personnel at 300 meters, and a vehicle at 400 meters during the day. Our simulated and live experiments incorporated a camera with resolution equivalent to a wide-angle GoPro Hero 4. Wide-angle cameras perform well to scan large areas; however, the lack of a narrow-angle camera with zoom capabilities required the sUAS to get within extremely close proximity to the entities to enable detection and categorization. However, this does not mitigate the inadequate slant range detection capabilities. In an operational setting, placing a platform within 60 meters of a vehicle, 40 meters of a person, or 15 meters from a weapon incur unacceptable risk to forces and would alert the targets that they are being observed.

However, we assess that the hardware and software solutions available to improve the slant range distance of detection for SPOTR also introduce technical implications worth noting. Improved camera systems could increase the payload weight of the sUAS which may impact the flight mechanics of the sUAS and the battery load. An increased battery load will deplete batteries quicker and could require solar panels, hand cranks, or generators to recharge the batteries. Adding super-resolution algorithms could increase the processing load of the microprocessor and limit the processing power available for SPOTR. These technical implications warrant further research into if and to what degree they may impact the current sUAS requirements and potentially reconsideration of the SATLAS Pillar 1.

There are potential operational implications that could affect SOF teams. Increasing battery load could decrease sUAS loiter time and reduce ISR coverage time for SOF teams. SOF teams could then be required to carry additional batteries to account for reduced loiter times which would increase their truckload requirement. Better resolution often requires increasingly narrow fields of view and zoom capabilities. While this improves detection accuracy and range, it could limit the team's overall situational awareness. Additionally, larger camera payloads may increase the overall size of the sUAS and impact the ability of SOF team members to easily transport the system in a rucksack.

V. CONCLUSION

The purpose of this research was to manage the development and integration of SPOTR object recognition software into a sUAS and evaluate its performance. Our objective was to advance the SATLAS project Pillar 2, integrating object recognition capability into the sUAS with a focus on technical performance and operational implications. We used simulations and field experimentation to measure the capabilities of SPOTR software, determine whether it meets SATLAS and RPUAS requirements, and identify capability shortfalls to the developer.

We collaborated with an industry leading software company to conduct a series of experiments to measure the technical requirements for accuracy, number of entities capable of detection, and range of detection. Although we experienced challenges with integration timelines and testing constraints, we were able to use a simulated environment and a surrogate platform that proved sufficient in answering our research questions. We assess that at its current stage of development, the SPOTR software can achieve the SATLAS technical requirements for accuracy and for the quantity of entities detected. SPOTR does not currently achieve the technical requirements for slant range distance of detection for SATLAS.

We assess that, while additional software design and integration remain, SPOTR has the potential to significantly improve situational awareness for SOF teams. We advanced Pillar 2 of the SATLAS project by employing a sUAS that is being used by U.S. forces. The accuracy and number of entities SPOTR can detect present opportunities to increase situational awareness and survivability of deployed SOF teams. Also, we identified the slant range distance shortfall and methods for improvement that subsequent student teams can improve upon to further the overall objective of SATLAS. Overall, we provided a theory analysis and method that helps to evident SATLAS-type systems.

A. SUMMARY

Our conclusions are organized by technical findings, operational requirements, and integration process.

1. Technical Implications

We assess that SPOTR can detect at over 90% accuracy and categorize entities in a simulated environment and temperate field environment while operating at the optimal slant range distance of 20 meters. We found that SPOTR was able to detect up to 13 entities and categorize eight when operating with limited processing power equivalent to that which is aboard the test drone. We also discovered that within the confines of our experiments and with the available camera, SPOTR was only able to detect entities out to approximately 40–45 meters slant range distance. Currently, SPOTR meets the technical requirements of SATLAS with respect to the accuracy and number of entities detected; however, it does not meet these requirements at the required range. There are hardware and software options to improve this capability, but they may introduce implications for the platform.

SPOTR's performance is dependent on pixel size and processing power to ensure accurate detection. Integrating a more capable and narrow view camera payload may be an approach to improve detection range yet this may increase the weight and battery load. Additionally, super-resolution algorithms offer another avenue to improve detection range, but this may require additional processing power and limit that available to SPOTR.

2. Operational Implications

The demand for SOF teams to employ small, organic sUAS in future operational environments will likely increase in the future. As risk to force and risk to mission factors increase in these scenarios, object recognition-enabled sUAS can be one tool to increase situational awareness and mitigate risk. We have assessed that accurate object recognition software maybe capable of simultaneously detecting the required number of entities aboard sUAS within the context of our experiment.

The current inability of SPOTR to achieve the range of detection required by SATLAS has important operational implications. In order for SPOTR to be useful to an SFOD-A, its detection range must increase. Increased payload weight can decrease loiter time and require SOF teams to carry additional batteries to compensate while narrow field of view cameras can reduce overall situational awareness.

By defining the SATLAS technical requirements and evaluating the accuracy and number of entities capable of being detected, our study advances the SATLAS project. This research serves as an additional step towards integrating object recognition and sUAS that can improve the situational awareness and survivability of SOF teams. Our operational requirements are based on our operational experience and informed by Army requirements and MDO doctrine.

3. Integration Process

While managing SPOTR integration, we simultaneously tested SPOTR capabilities for further SATLAS research. We were unable to complete Pillar 2 analysis due to the complex nature of integrating systems from two different commercial industries. To mitigate the impact of delayed development, we employed simulations to measure object recognition performance. We propose that SPOTR performance in the simulated environment was comparable to real-world performance when camera settings were adjusted accordingly. This should provide an avenue for future SPOTR testing without being dependent on resolving commercial industry requirements to collaborate between competitors. Since the SATLAS prototype design currently in development incorporates a more capable camera than the Nibbler, adjusting the camera test specifications will enable future studies to accurately measure performance. The ability to integrate NPS student projects and commercial industry to develop solutions to real-world problems does present significant opportunities for SATLAS to achieve our primary objective.

B. LIMITATIONS

This thesis was limited in several ways. We intentionally limited its scope to Pillar 2 of the SATLAS project while leaving Pillars 3 and 4 for future studies. An 18-month academic research cycle and prototype development delays did not allow us to test our intended prototype. Instead, we relied on simulations and a surrogate platform to examine the SPOTR software. The reader should keep these limitations in mind when considering our conclusions and recommendations.

C. FUTURE WORK

This section provides recommendations for future work to develop an organic, object recognition-enabled sUAS for SOF teams. Our study focused on SPOTR's detection and categorization performance in a simulated environment and onboard a surrogate platform. We recommend that future work focus on SPOTR's capability onboard the SATLAS Altavian M440. Given the dependence of object recognition software on the available camera, we expect the M440s to perform better than the Nibbler.

Future work that is immediately pursuable is adjusting the simulation camera settings to replicate the more capable camera onboard the M440. We believe this would facilitate measuring the comparable accuracy and range of the desired prototype. Future efforts will rely on continued management of the integration onto the M440 in collaboration with AeroVironment. Employing the simulation as a testing environment will allow initial measurements to be gathered while integration efforts continue.

Additionally, we recommend that future studies examine the use of super resolution algorithms as a software technique to increase detection range. We found detection range to be the only shortcoming of the SPOTR software. Employing a more capable camera will undoubtedly increase detection range but it still may not meet the specifications required for SOF teams.

Finally, we recommend that planning begins to integrate Pillars 3 and 4 into testing. Autonodyne is a California-based company that has begun development on a sUAS common control station that integrates several autonomous features. Given that the overall goal of SATLAS is to integrate all four pillars into one complete system, Autonodyne may offer opportunities for collaboration on the final two pillars.

LIST OF REFERENCES

Alberts, David and Hayes, Richard, *Code of Best Practice for Experimentation*, CCRP Publication Series (Washington, DC: DOD Command and Control Research Program, 2002). http://dodccrp.org/files/Alberts_Experimentation.pdf

Borel, Christoph and Young, S. "Image Quality and Super Resolution Effects on Object Recognition Using Deep Neural Networks," in 2019 SPIE *Conference for Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications.* (SPIE, 2019), https://doi.org/10.1117/12.2518524

Collins, Liam and Morgan, Harrison "Affordable, Abundant, and Autonomous: The Future of Ground Warfare," *War on the Rocks* (April 21, 2020): https://warontherocks.com/2020/04/affordable-abundant-and-autonomous-the-future-of-ground-warfare/.

Conley, Heather et al., *The Kremlin Playbook: Understanding Russian Influence in Central and Eastern Europe*, (Washington, DC: Center for Strategic and International Studies, 2016), https://www.csis.org/analysis/kremlin-playbook.

Clausewitz, Carl V. *On War*, Translated by Michael Howard and Peter Paret, (Princeton University Press, 1976), Book 1, Chapter 7.

Daniel Midgett, email message to author, September 4, 2020.

Defense Advanced Research Projects Agency,*With Squad X, Dismounted Units Partner with AI to Dominate Battlespace*, (Washington, DC: Defense Advanced Research Projects Agency, 2019), . https://www.darpa.mil/news-events/2019-07-12

Department of Defense, *Broad Agency Announcement For Technology Development and Advanced Technology Development* (Washington, DC: Department of Defense, 2020).

Department of Defense, "Capability Production Document For Rucksack Portable Unmanned Aircraft System (RPUAS) Increment 2," (Washington, DC; Department of Defense, 2013).

Department of Defense*, Defense Innovation Unit Announces sUAS Product Availability to Provide Secure, Capable Small Unmanned Aerial Systems for Critical Uses Across the Government*, Washington, DC: Department of Defense, 2020. https://www.defense.gov/Newsroom/Releases/Release/Article/2318799/defense-innovation-unit-announces-suas-product-availability-to-provide-secure-c/

Department of Defense, *Eyes of the Army: U.S. Army Roadmap for UAS 2010–2035*, (Washington, DC: Department of Defense, 2010). https://fas.org/irp/program/collect/uas-army.pdf

Department of Defense, *RQ-11B Raven Small Unmanned Aircraft System (SUAS)* (Washington, DC: Department of Defense, 2016), https://asc.army.mil/web/portfolio-item/aviation_raven-suas/

Department of Defense, *Summary of the 2018 National Defense Strategy of the United States of America.* Washington, DC: Department of Defense 2018. National Defense Strategy, Department of Defense, 2018. https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf

Department of the Army, *Army Futures Command Concept for Special Operations 2028,* Army Futures Command Pamphlet 71-20-4: (Washington, DC: Department of the Army, 2020), https://api.army.mil/e2/c/downloads/2021/01/05/bdd61c44/20200918-afc-pam-71-20-4-afc-concept-for-special-operations-2028-final.pdf

Department of the Army, *Army Special Operations,* ADP 3–05 (Washington, DC: Department of the Army, 2019), https://armypubs.army.mil/epubs/DR_pubs/DR_a/pdf/web/ARN18909_ADP%203-05%20C1%20FINAL%20WEB(2).pdf

Department of the Army, *The U.S. Army in Multi-Domain Operations 2028,* U.S. Army TRADOC Pamphlet 525-3-1 (Washington, DC: Department of the Army, 2018), December 6, 2018, https://api.army.mil/e2/c/downloads/2021/02/26/b45372c1/20181206-tp525-3-1-the-us-army-in-mdo-2028-final.pdf

Eckstein, Megan "Marines' 3D-Printed 'Nibbler' Drone Creating Lessons Learned on Logistics, Counter-UAS," *USNI News*, (September 2017), https://news.usni.org/2017/09/27/marines-3d-printed-nibbler-drone-creating-lessons-learned-logistics-counter-uas.

Eidson, Bill "Nibbler Drone Is an Advanced Manufacturing 'Flagship' for Marines," *MITRE*, (January 2019), https://www.mitre.org/publications/project-stories/nibbler-drone-is-an-advanced-manufacturing-flagship-for-marines.

Everingham, Mark et al., "The Pascal Visual Object Classes Challenge: A Retrospective," *International Journal of Computer Vision* 111, no. 1 (January 1, 2015): 98–136, https://doi.org/10.1007/s11263-014-0733-5.

Fraga-Lamas, Paula et. al., "A Review on IoT Deep Learning UAV Systems for Autonomous Obstacle Detection and Collision Avoidance," *Remote Sensing* 11, no. 18 (January 2019): 2144, https://doi.org/10.3390/rs11182144.

Friend, Alice *DOD's Report on the Investigation into the 2017 Ambush in Niger*, (Washington, DC: Center for Strategic and International Studies, 2018), https://www.csis.org/analysis/dods-report-investigation-2017-ambush-niger.

Garamone, Jim "Milley Makes Case for U.S. Military Keeping Up With Global, Technology Changes," U.S. Department of Defense, (December 2, 2020): https://www.defense.gov/Explore/News/Article/Article/2432855/milley-makes-case-for-us-military-keeping-up-with-global-technology-changes/.

GoPro, *GoPro Hero 4 Silver User Manual,* 2014, https://gopro.com/content/dam/help/hero4-silver/manuals/UM_H4Silver_ENG_REVA_WEB.pdf

Hambling, David "U.S. To Equip MQ-9 Reaper Drones with Artificial Intelligence," *Forbes*, (December 2020), https://www.forbes.com/sites/davidhambling/2020/12/11/new-project-will-give-us-mq-9-reaper-drones-artificial-intelligence/?sh=37723b3c7a8e.

Kofman, Michael et al., *Lessons from Russia's Operations in Crimea and Eastern Ukraine*, RR 1498 (Santa Monica, CA: RAND, 2017), https://www.rand.org/pubs/research_reports/RR1498.html

LeCun, Yann et al., "Deep Learning," *Nature,* no. 7553 (May 2015): 436–44, https://doi.org/10.1038/nature14539.

Li Liu et al., "Deep Learning for Generic Object Detection: A Survey," *International Journal of Computer Vision* 128, no. 2 (February 2020): 261–318, https://doi.org/10.1007/s11263-019-01247-4.

McArdle, Jennifer "Simulating War: Three Enduring Lessons from the Louisiana Maneuvers," *War on the Rocks* (March 2021), https://warontherocks.com/2021/03/simulating-war-three-enduring-lessons-from-the-louisiana-maneuvers/.

Mohey El-Din, Doaa et. al. "An Automatic Detection of Military Objects and Terrorism Classification System Based on Deep Transfer Learning," in *Proceedings of the International Conference on Artificial Intelligence and Computer Vision (AICV2020)*, ed. Aboul-Ella Hassanien et al., Advances in Intelligent Systems and Computing (Cham: Springer International Publishing, 2020), 594–603, https://doi.org/10.1007/978-3-030-44289-7_56.

Midgett, Daniel et. al., "Semiautonomous Threat Learning Alert System," (master's thesis, Naval Postgraduate School, 2020), 46–52, December 2020.

Miller, Jack "Strategic Significance of Drone Operations for Warfare," *E-International Relations*, (August 2013), https://www.e-ir.info/2013/08/19/strategic-significance-of-drone-operations-for-warfare/.

Office of the Director of National Intelligence, *Annual Threat Assessment of the U.S. Intelligence Community*, (Washington, DC: Office of the Director of National Intelligence, 2021), https://www.dni.gov/files/ODNI/documents/assessments/ATA-2021-Unclassified-Report.pdf

Seffers, George “Army Tactical Assault Kit Always Adapting for New Era,” *Signal,* October 28, 2020. www.afcea.org/content/army-tactical-assault-kit-always-adapting-new-era

Statnikov, R. et. al., “Multicriteria Analysis Tools in Real-Life Problems,” *Computers & Mathematics with Applications* 52 (2006) 1–32. https://doi.org/10.1016/j.camwa.2006.08.002

Sterio, Milena “The United States’ Use of Drones in the War on Terror: The Legality of Targeted Killings under International Law,” *Case Western Reserve Journal of International Law* 45 (2012) https://scholarlycommons.law.case.edu/cgi/viewcontent.cgi?article=1072&context=jil.

Suchman, Lucy “Algorithmic Warfare and the Reinvention of Accuracy,” *Critical Studies on Security* 8, no. 2 (May 3, 2020): 175–87, https://doi.org/10.1080/21624887.2020.1760587.

Yang, Zhi et al., “Deep Transfer Learning for Military Object Recognition under Small Training Set Condition,” *Neural Computing and Applications*, no. 10 (October 1, 2019): 6469–78, https://doi.org/10.1007/s00521-018-3468-3.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California