

AUTONOMOUS AERIAL DRONES CONNECTING PUBLIC SAFETY: OPPORTUNITIES AND CHALLENGES FOR THE FUTURE

Raymond Sheh,¹ Donald Harriss,² and Karen Geappen³

Reliable voice and data communications are critical in supporting public safety responders during incidents and disasters and will become more important as ground robots and Internet-of-Things (IoT) sensors for public safety applications proliferate. Uncrewed Aircraft Systems (UAS), consisting of an Uncrewed Aerial Vehicle (UAV) and associated systems, are valuable tools for public safety personnel to reduce risk to the public and themselves. They improve the service provided to the community by delivering rapidly deployable “Eyes in the Sky.”

It is a natural progression for UAVs to be more closely integrated into the communications of public safety operations. For example, they can provide critical communication links in public safety operations and connect and provide information directly to nearby personnel, devices, and perhaps even the public. UAVs have the potential to fill a vital technology gap where existing infrastructure is damaged or non-existent and where structures or topography make traditional radios and satellite communications ineffective.

In this paper, we discuss the unique capabilities that the closer integration of UAS into public safety communications brings to current and future public safety missions, the challenges users may face in a UAS deployment, and the regulatory and ethical considerations that come with the use of UAS. This discussion dovetails with work that the U.S. National Institute of Standards and Technology (NIST) Public Safety Communications Research (PSCR) Division's Uncrewed Aircraft Systems (UAS) portfolio is undertaking to drive the development of UAS for public safety, with an emphasis on reliable communications, responsible cybersecurity, and artificial intelligence (AI) risk management. This paper will reference related data points from the recently released NIST Cybersecurity Framework (CSF) 2.0 and Artificial Intelligence (AI) Risk Management Framework (AI RMF).

¹ Uncrewed Aircraft Systems Research Lead, Public Safety Communications Research Division, National Institute of Standards and Technology, 100 Bureau Drive, Gaithersburg MD 20877, USA

² Uncrewed Aircraft Systems Technical Lead, Public Safety Communications Research Division, National Institute of Standards and Technology, 325 Broadway, Boulder, CO 80305, USA

³ Independent Researcher, PO Box 174, Bull Creek, Western Australia, Australia, 6149

INTRODUCTION

Public safety personnel rely on the integrity and availability of communications for their safety and that of the broader public. This reliance is especially true in dynamic response situations and unstructured environments, such as wilderness rescue, wildland and urban fires, and active shooter situations. Personnel in the field need accurate and relevant real-time information from various sources and locations to inform actions, such as a fire actively jumping a fire break. Incident commanders need to know how a situation evolves in real-time and need to be able to look to their personnel in the field and “Eyes in the Sky” to provide situational awareness of personnel, property, and other assets. There are many documented cases where communication delays and breakdowns in data delivery have cost lives, property, or unrecoverable damage to the local ecosystem.

Infrastructure disruptions can further complicate disaster management by preventing the use of cellular networks and other fixed radio infrastructure. Point-to-point systems, such as Land Mobile Radio (LMR), are limited by the chosen protocols, the laws of physics, and the laws of information theory pertaining to the quantity and quality of information LMR can transmit. These factors define an upper limit as to what is possible for a given number of people spread over a given area, in a given environment, and the use of a given amount of radio frequency spectrum.

Uncrewed Aircraft Systems (UAS), consisting of an Uncrewed Aerial Vehicle (UAV, often also called a drone) and associated systems and software⁴, are increasingly employed in public safety missions. UAS are becoming potential solutions to shortfalls in public safety communications. Their rapidly decreasing cost, increasing functionality, and ease of use make UAS valuable assets for many public safety agencies, whether as an “Eye in the Sky” or for more advanced operations⁵.

Today, most public safety UAS operations have two-way data communications only between the UAV and operator control unit (OCU, often a handheld controller). The OCU receives commands from the operator(s) and transmits these commands to the UAV. In turn, the OCU receives data from the UAV, such as imagery, videos, and telemetry (e.g., location and heading). The OCU may also receive other information from the UAV, depending on the UAV’s sensors.

While not usually considered part of communications, the UAV, and sometimes OCU, will receive information from systems such as the Global Positioning System (GPS) and broadcast information such as Remote Identification (Remote ID) to meet FAA and other regulatory requirements.

This standalone UAS model was suitable for situations where the UAS was only an “Eye in the Sky,” whereby the imagery had to be mediated through personnel at the OCU, such as the pilot or payload operator, and then perhaps the incident commander, before being available to other personnel in the field.

However, as the quality of data from UAS has improved, along with their ability to operate with increasing levels of autonomy, requiring information from the UAS to pass through the singular OCU can present a bottleneck to getting timely information to personnel in the field. Additionally, most UAS designs are closed or proprietary, making it difficult to communicate

⁴ In this document, we make a clear distinction between the UAV, the flying vehicle itself, and the UAS, which includes the UAV and its associated support equipment and systems, such as antennas, operator control unit (OCU), communications infrastructure, and so-on.

⁵ While this document primarily focuses on Small Uncrewed Aircraft Systems (sUAS) that are in common deployment across public safety agencies, the principles are largely applicable to UAS in similar roles, regardless of size.

with other endpoints in the field. For example, personnel in the field usually cannot receive information directly from the UAV. The UAV also usually cannot receive data from IoT sensors in the field.

Furthermore, a UAV, particularly one capable of loitering for more extended periods, is often in a much better position to communicate with personnel in the field than fixed or mobile ground-based stations. It is quite possible that the only equipment that is in the right place to transmit critical information to a particular person in the field might be the UAV.

Imagine if the UAV could integrate and participate more fully in the response team's communications ecosystem. Instead of only communicating operationally with one OCU, it behaves more like another team member in the communications network. Perhaps it could decide to directly transmit a subset of its data to personnel within range without needing to go through the OCU. Perhaps the UAV could act as a relay, using its extra height to link personnel to each other and to the incident commander, who might be on the other side of a ridge and be limited as to with whom they can directly communicate. Multiple UAVs could even form a mobile ad-hoc network (MANET)⁶, a type of wireless network that can reconfigure itself as the UAVs and other nodes move in and out of range of each other. The network dynamically routes information between nodes that cannot directly communicate with each other by finding intermediate nodes through which data can be relayed. Such a network can maintain communications across a changing operational environment as network nodes on vehicles, UAVs, and personnel move in and out of range of each other.

The principle is simple enough. The underlying technology exists and has been demonstrated to work. Indeed, with the increasing levels of autonomy of UAS, particularly those operating in clear sky and away from obstructions, a UAS can truly be a team player. It can take commands from different people at various times while performing autonomously as an “Eye in the Sky” and as a communications relay. However, as with all new technology, the details matter. It is not enough to prove that a system works. The system needs to be appropriately, reliably, safely, and ethically integrated into public safety operations.

This paper will first provide background in communications, UAS, cybersecurity, autonomy levels, and risk management. We will then describe how new, upcoming technologies relating to autonomous, highly connected UAS can benefit public safety and delve into the associated risks, potential regulatory interactions, ethics, and social acceptance.

BACKGROUND

Communications in Public Safety

First responder autonomous robots, like UAS, require secure, priority communication capabilities. This requirement presents technical challenges such as allocation of services, attention to cybersecurity, and use of AI data. Prior knowledge of public safety communications is required to understand the challenges presented by these new technologies. Public safety uses dedicated wireless interoperability frequencies for first responder events and communications. These frequencies include various channels in the 700MHz and 800MHz bands and interoperability channels in the Very High Frequency (VHF, 30 MHz - 300 MHz) and Ultra High Frequency (UHF, 300 MHz - 3 GHz) bands. These channels are divided into dedicated talk groups for specific purposes. Over time, many of these channels have moved from analog to digital communication methods for better bi-directional communication. However, switching

⁶ For the purpose of this discussion, we make the distinction between MANET and the broader category of Mesh Networks that include pre-defined, fixed networks.

from analog to digital has yet to solve scaling problems, especially where data-rich technologies require additional bandwidth. To better use these limited resources, a frequency spectrum segment from the 700MHz range has been allocated to deploying dedicated broadband cellular communications for first responders. This allocation includes enhancements for greater coverage, interagency collaboration, security, and data applications for situational awareness. New communication tools like UAS and data-rich methodologies, such as AI, require digital communication networks that go beyond traditional systems. Industrial, Scientific, and Medical (ISM) frequencies are used for most UAS communications. The ISM band has opened up new opportunities for innovation and accessibility in wireless communication, leading to the growth of various technologies that rely on wireless communication. However, there are not any dedicated spectrum offerings for UAS operations; first responder activities must compete with non-responder civilian traffic on the ISM spectrums. New deployment systems, such as Drone as First Responder (DFR)⁷ systems, which we will discuss in more detail, can take advantage of cellular broadband communications for prioritized, responder-focused networks.

First responders' operational demands are growing quickly, and traditional networks cannot provide reliable communications. Modern facilities are moving towards connected dispatch systems using AI-driven applications to orchestrate remote command centers and autonomous assets while monitoring boots-on-the-ground operations. To support these diversified, interoperable systems, networks will require more bandwidth and be able to handle multiple traffic types.

Flexible, scalable, and seamlessly integrated deployable broadband networks can meet the increasing demand for reliable, high-performance connectivity. These networks leverage AI-driven, advanced routing, software-defined networks and meshing protocols to ensure uninterrupted connectivity despite network failures or congestion. This level of service allows for robust applications that use AI at the user network edge in near real-time. It is good practice for organizations to achieve reliable communications by establishing end-to-end agreements among stakeholders to ensure the safe and secure use of networked systems and AI-driven applications. Building a reliable, sustainable ecosystem requires collaboration across industries and adopting common standards and protocols to enable interoperability and seamless integration of different first responder applications.

Current Use of UAS in Public Safety

Public safety personnel predominantly use UAS to obtain visual information, usually through conventional (visible light) cameras but also increasingly through thermal imagery and 3D mapping. UAVs are not only used in lieu of crewed aircraft but also in new applications, such as flying low under trees or into buildings, where it is impossible to use crewed aircraft. While the situations vary between different public safety applications, in general, the goal is to put eyes onto a situation from a perspective where it is difficult, dangerous, or impossible to otherwise gather that information.

In the general case, the UAS is deployed by public safety personnel at, or close to, the response location. The deployment may be by a dedicated UAS response team or by personnel who have other roles but also maintain proficiency and equipment for deploying UAS as necessary. In such use cases, the UAV communicates directly with an OCU, generally in an ISM band, although some communicate in specific public safety bands. The OCU sometimes has a

⁷ King, S., Major, s., Maccollum, M. "Drone as First Responder Programs: A New Paradigm in Policing," *MITRE*, 2023-08-11, <https://www.mitre.org/news-insights/publication/drone-first-responder-programs-new-paradigm-policing>

communications link, usually via a cable, wifi, or cellular service, to allow the UAV to communicate with other systems. These systems are often collaboration tools⁸ like Team Awareness Kit (TAK)⁹, DroneSense¹⁰, or DroneDeploy¹¹. These tools allow this information to be shared with other team members. Some of these services also provide the ability to control the UAV through the OCU.

We are increasingly seeing the use of DFR, which refers to a class of UAS that aims to have the UAV deploy to the response location independent of the responding personnel. The UAV is usually pre-stationed in weatherproof pods on the roofs of buildings in the service area, from which they launch and where they land without physical human intervention. They are generally controlled through a control center that is neither in proximity to the UAV nor the response location. Due to current regulatory requirements, there may be a person in the proximity of the UAV pod whose role is to maintain situational awareness of the airspace around which the UAV flies. With the advent of automated airspace monitoring technologies, this restriction is loosening.

Decoupling the flight of the UAV from the responding personnel requires a new paradigm in communication. The UAV is no longer near a dedicated OCU with which it communicates. Instead, systems use a deployment pod, which may have colocated antennas for communication with the UAV. The UAV is also in a dedicated service area, which allows the UAV to communicate with known communications infrastructure, such as 4G and 5G cell phone towers. We are starting to see UAVs that communicate with this infrastructure. For example, several DFR-focused UAVs have one or more cellular SIM cards and the appropriate cellular modems to connect to 4G and 5G infrastructure.

We are also starting to see UAVs for public safety use that can participate in Mobile Ad-hoc Networks (MANET). This allows the UAV to dynamically discover and communicate with other public safety entities, such as vehicles and people equipped with compatible MANET equipment, as they move in and out of range.

These new communications paradigms bring immense opportunities for greater reliability and safety, as well as new use cases. However, they also bring considerable technical, regulatory, safety, and ethical challenges to consider. The purpose of this document is to present some of these opportunities and challenges, to give stakeholders greater visibility, and to allow them to make informed choices regarding the development and adoption of such technologies.

Cybersecurity

For the purpose of this paper, we use the definition of cybersecurity used in NIST Special Publication 800-53 Rev. 5¹²:

“Prevention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and nonrepudiation.”

⁸ Any mention of commercial products is for information only; it does not imply recommendation or endorsement by NIST.

⁹ Team Awareness Kit Product Center, “TAK Product Lines,” accessed 2024-04-03, <https://tak.gov/products>

¹⁰ DroneSense, “DroneSense Home Page,” accessed 2024-04-03, <https://www.dronesense.com/>

¹¹ DroneDeploy, “DroneDeploy Home Page,” accessed 2024-04-03, <https://www.dronedeploy.com/>

¹² Joint Task Force, “NIST Special Publication 800-53 Revision 5 Security and Privacy Controls for Information Systems and Organizations,” *National Institute of Standards and Technology (NIST)*, 2020-12-10, <https://doi.org/10.6028/NIST.SP.800-53r5>

This is a broad definition that deliberately encompasses both a wide variety of assets and a wide variety of impacts on an organization.

Particularly important to note, when it comes to increasingly connected UAS, is that cybersecurity is not just limited to people breaking into the UAS to steal information. As technology has become more accessible, a wider variety of adversaries have proliferated, with a correspondingly wide variety of motivations to cause damage.

Cybersecurity risks have also increased in both variety and severity as critical systems become more interconnected. Attackers can enter through a seemingly unrelated system and find ways to jump between networks, an action known as “moving laterally.” For example, several well known hacks of big corporate networks started with the attackers compromising systems of related organizations, such as contractors or suppliers, and then using the access that those organizations had to break into their actual target. In this context, all organizations in the network have a stake in protecting each other.

While the wording of this definition focuses on equipment, the human factor is also critical to consider in cybersecurity, both in terms of attacks coming from inside an organization, as well as attackers stealing an individual’s credentials. This may be as simple as guessing things like passwords or password reset details or as complex as crafting specific, directed “phishing” emails that trick the individual into handing over important information. Two-factor or multi-factor authentication is often touted as a solution to this problem by requiring the user to prove in multiple ways that they are who they say they are. For example, the system might ask for both a password and a one-time code that is emailed to them. Of course, in the context of connected UAS, multi-factor authentication has to be chosen carefully. If the UAS is being used by public safety as part of an effort to restore communications, requiring cell phone access to receive a code to launch the UAV might not be a good solution.

There is much guidance available in the cybersecurity space. This includes:

- The NIST Privacy and Risk Management Frameworks.
- NIST SP 800-53 Rev. 5, Security and Privacy Controls for Information Systems and Organizations.
- NIST SP 800-171 Rev. 2, Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations¹³.
- The Criminal Justice Information Services (CJIS) Security Policy¹⁴.
- International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC) ISO/IEC 27001:2022 Information security, cybersecurity, and privacy protection¹⁵.

However, there is very little in the way of comprehensive guidance available in the UAS space. As we illustrated above and will discuss further later in this document, there are

¹³ Ross, R., Pillitteri, V., Dempsey, K., Riddle, M., Guissanie, G. “NIST SP 800-171 Rev. 2 Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations,” *National Institute of Standards and Technology (NIST)*, updated 2021-01-28, <https://doi.org/10.6028/NIST.SP.800-171r2>

¹⁴ CJIS Information Security Officer, “CJISD-ITS-DOC-08140-5.9 Criminal Justice Information Services (CJIS) Security Policy Version 5.9,” U.S. Department of Justice, Federal Bureau of Investigation, Criminal Justice Information Services Division, 2020-06-01, https://www.fbi.gov/file-repository/cjis_security_policy_v5-9_20200601.pdf

¹⁵ International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), “ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection, Information security management systems, Requirements,” <https://www.iso.org/standard/27001>

considerable dangers to blindly applying cybersecurity guidance to UAS without considering the unique aspects of the application.

Specific guidance regarding the management of cybersecurity risk for Uncrewed Systems is still in its relative infancy. Organizations that have issued guidance that relates to Uncrewed Systems include:

- The Cybersecurity & Infrastructure Security Agency (CISA), part of the U.S. Department of Homeland Security (DHS), has issued guidance that focuses on the procurement and use of Uncrewed Aircraft Systems¹⁶ rather than the development, manufacture, or broader integration into wider organizational systems and processes.
- Both the Defense Innovation Unit (DIU) of the U.S. Department of Defense (DoD) and AUVSI have issued guidance on the procurement of UAS based on the National Defense Authorization Act (NDAA)¹⁷. DIU maintains a “Blue UAS” list¹⁸ while AUVSI maintains a “Green UAS” list¹⁹, which list the UAS that they have evaluated according to the NDAA criteria. This guidance is primarily focused on geopolitical, cybersecurity, organizational, and supply chain considerations of the UAS manufacturer rather than its integration into the organization of the end user’s broader systems and risk management.
- The NIST Information Technology Laboratory (ITL)²⁰ produces guidance on a wide variety of cybersecurity topics, including the aforementioned resources. While it does not have an Uncrewed Systems-specific focus, it does focus on applications such as Automated Vehicles²¹. Much of the guidance within the Cyberphysical Systems and Internet of Things (IoT) topics is also highly relevant to Uncrewed Systems.

As we will discuss later, we have established a working²² group that aims to adapt these resources to the UAS space to better fill the gaps in available guidance to stakeholders, including end users, manufacturers, developers, researchers, and regulators.

Artificial Intelligence

For this discussion, we will use the ISO/IEC 22989:2022²³ definition of AI:

“... a technical and scientific field devoted to the engineered system that generates outputs such as content, forecasts, recommendations or decisions for a given set of human-defined objectives.”

Within this broad definition, we will also discuss the following:

¹⁶ Cybersecurity & Infrastructure Security Agency (CISA), “Be Air Aware,” accessed 2024-04-03, <https://www.cisa.gov/topics/physical-security/unmanned-aircraft-systems>

¹⁷ House Armed Services Committee, “NDAA - National Defense Authorization Act,” accessed 2024-04-03, <https://armedservices.house.gov/ndaa>

¹⁸ Defense Innovation Unit, “About Blue UAS”, accessed 2024-04-03, <https://www.diu.mil/blue-uas>

¹⁹ Association for Uncrewed Vehicle Systems International (AUVSI), “Green UAS,” accessed 2024-04-03, <https://www.auvsi.org/green-uas>

²⁰ National Institute of Standards and Technology (NIST), “Information Technology,” accessed 2024-04-03, <https://www.nist.gov/information-technology>

²¹ National Institute of Standards and Technology (NIST), “NIST Automated Vehicles Program,” accessed 2024-04-03, <https://www.nist.gov/programs-projects/nist-automated-vehicles-program>

²² National Institute of Standards and Technology (NIST), “PSCR UAS Working Group, Cybersecurity and AI Risk Management for Uncrewed Aircraft Systems (UAS) in Public Safety,” updated 2024-02-27, <https://www.nist.gov/ctl/pscr/pscr-uas-working-group>

²³ International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), “ISO/IEC 22989:2022 Information technology, Artificial intelligence, Artificial intelligence concepts and terminology,” accessed 2024-04-03, <https://www.iso.org/standard/74296.html>

- Machine Learning (ML) describes AI systems whose behavior depends also on observed input data.
- Statistical ML describes ML systems whose behavior depends specifically on the statistical properties of the observed input data.
- Deep Learning (DL) describes Statistical ML systems whose behavior depends on fitting highly complex, randomly initialized sets of equations to the observed input data.

It is essential to understand that various types of AI have different properties. DL is very topical right now and has some very well-known properties, such as requiring immense amounts of data. However, other forms of ML, such as regression techniques, use task-related data differently and may require less of it. Some DL and ML systems use foundational models and transfer learning, whereby the system first learns to make decisions relating to a different task and then uses that experience to help it learn to make the desired decisions more quickly.

There are other forms of AI that are not ML and have yet other properties. Classical planning techniques do not rely on the statistical properties of data and instead solve problems expressed in much more mathematical terms. Reasoning systems, in contrast, make use of large amounts of background knowledge, such as knowledge of physics, combined perhaps with a limited amount of task-related data, to make their decisions.

This distinction is important because different forms of AI expose stakeholders, including users, the organization, and the community, to different risks. However, these may only be obvious when detailed information about the underlying forms of AI is known. These distinctions can significantly affect how risks associated with the deployment and use of the system should be managed. As UAVs become more autonomous, with less human oversight, and as these new communications paradigms become more common, the management of these risks becomes even more critical.

For example, if a user disagrees with a decision that an AI system makes, their ability to find out who was correct, and why the decision was made, varies considerably depending on the underlying form of AI used in the system. In some cases, it is obvious that the AI system made a mistake, such as the autonomous UAV clearly not reacting to an obstacle that appears in its sensors. In other cases, it is less obvious, such as when the UAV takes an unexpected route that still reaches its destination. Disagreements like this are not necessarily a sign of failure in the system. After all, one of the primary reasons for using AI is for the system to perform better than humans, so it should not be surprising that the AI system might sometimes make decisions that humans might disagree with in the moment.

Some AI systems are relatively transparent and able to explain their decisions in a way that suitably trained humans can understand in an actionable manner²⁴. If the system was correct in its decision, it might help users understand how it came to its decision, and the factors it considered. Systems based on classical planning and, to some extent, regression often have such properties.

However, many AI systems, particularly ones based on statistical ML, are not so transparent. This lack of transparency may be because the system makes its decisions in a manner that is too complex, or otherwise too difficult, to be meaningfully interpreted, as is the case in DL systems. In other cases, it may be because the system has not arrived at its decision in a way that has a meaningful interpretation, such as systems that develop their decision-making models through random trial and error.

²⁴ Sheh, R., Monteath, I., “Introspectively Assessing Failures through Explainable Artificial Intelligence,” *IROS-17 Workshop on Introspective Methods for Reliable Autonomy (IMRA)*, 2017

Along similar lines is the ability to correct errors that are encountered. The new applications enabled by highly connected UAVs often run with reduced human oversight. This makes human judgment less effective as a control on systematic errors and increases the importance of correcting detected errors in a verifiable manner. Once again, AI systems vary considerably in their ability to be repaired and to check that the repair fixes related errors without causing new errors.

Levels of Autonomy

When discussing opportunities and risks associated with autonomous UAS, it is useful to distinguish the various levels at which decisions are made within the UAS, if those decisions represent autonomy or automation, and if those decisions are made on the UAV or with the assistance of other parts of the UAS such as the OCU. In this section, we will describe some illustrative examples from currently deployed UAS, loosely guided by the NIST Autonomy Levels for Unmanned Systems (ALFUS)²⁵. As a new generation of connected UAS applications run with reduced human oversight, it is imperative to understand which decision-making processes, at the various levels of autonomy, were previously supervised and where systems are now expected to be trusted to “do the right thing.”

Control Autonomy. The vast majority of UAVs in use by public safety are autonomous to some degree at the control level. The operator will command a direction in which to fly, rotate, ascend, or descend. An autonomous controller in the UAV takes this command, combines it with sensor data, and turns it into motor commands that fly the UAV in the desired direction, accounting for the wind, variations in the UAV’s behavior due to age, wear and tear, damage, and any added accessories. This type of autonomous controller deals with very simple tasks in a moderately complex environment of wind and varying system characteristics. The human provides the commanded direction but does not provide the UAV with assistance in determining how to spin the motors to achieve that commanded direction.

This contrasts with other UAVs, such as traditional model aircraft or acrobatic drones, which rely almost entirely on the operator to deal with such variations. These UAVs use very simple analytical controllers and simple sensor data and processing to determine how to spin the motors in response to the operator’s commands.

More recent public safety UAVs have an additional autonomy layer that makes use of AI, often including at least some DL, to detect obstacles around the UAV and determine if the UAV should stop or fly around the obstacle, overriding the commanded direction of flight from the operator. These UAVs deal with the same simple task of “fly in a given direction” but now account for a more complex environment where there are obstacles with which to collide.

On top of these lower-level autonomy layers, most deployed UAVs receive their commanded flight direction from a human operator. Examples of these flight-level commands might be “fly forward at a moderate speed” or “turn left.” Some are controlled by an OCU that is much like that of a model aircraft, with two joysticks that the operator uses to command the UAV to move in a particular direction and at a particular speed. It is assumed that the operator is handling these higher levels of task and mission complexity while the lower levels handle, and abstract away, some of the environmental and system complexities.

²⁵ Ad Hoc ALFUS Working Group Participants, “NIST Special Publication 1011-II-1.0, Autonomy Levels for Unmanned Systems (ALFUS) Framework Volume II: Framework Models Version 1.0,” *National Institute of Standards and Technology (NIST)*, 2007-12-28, <https://doi.org/10.6028/NIST.sp.1011-II-1.0>

Task Autonomy. Most UAVs deployed in public safety applications also have some basic automation features to which the operator can delegate simple flying tasks, such as flying orbits or a grid pattern. Automation of these task-level behaviors is distinguished from autonomy in that automation is generally a set of pre-defined or pre-computed commands and does not respond to unexpected situations with the intent of continuing its task. They can only handle static and usually relatively simple environments. Most orbit or grid pattern features on UAS, be they running directly on the UAV or elsewhere in the UAS, such as the OCU, will generally stop if they encounter high winds or an unexpected obstacle (that the lower-level autonomy cannot handle transparently).

Fly-to-Point and Return-to-Home task-level behaviors can vary in their level of complexity. Early systems were purely automated, whereby the UAV would ascend to a given height and then fly directly towards a defined location, such as the denoted fly-to point, takeoff point or OCU location. Other Fly-to-Point and Return-to-Home features can be elaborate autonomous behaviors, which can fly around obstacles based on a 3D map, either provided or generated by the UAV or other parts of the UAS. DFR systems typically exhibit autonomous behaviors at this level, with the operator, usually in a command center, providing the fly-to points.

Some UAVs have a degree of landing autonomy that allows the UAV to vary its landing location depending on the suitability of the terrain. Such systems make extensive use of AI to characterize the terrain and assess the risk of landing on a bad surface versus the risk of running out of power. Some of these assessments can be quite involved, particularly if they also account for other agents, such as animals that may appear at the intended landing site. Such systems can potentially override the human operator. This capability is crucial for systems where there is a significant latency between the operator and the UAV, such as in the case of novel communications technologies that might improve reliability and range but at the expense of increased latency. In such cases, this additional latency may make it impossible for the operator to quickly intervene to address unexpected events.

There are autonomous goal-directed behaviors that are increasingly seen on UAS deployed by public safety. In such systems, the location that the UAV should fly to is determined on the fly by (usually) onboard sensing, with the human operator providing some specific mission information. An example of such an advanced task-level behavior is a UAV that autonomously tracks and follows an operator-designated object of interest that is moving relative to it through the use of its onboard sensors, such as its cameras. Such a system may need to deal with substantial mission and environmental complexity with minimal human intervention. This is especially the case if it must handle the object of interest being temporarily hidden behind other objects and, thus, needs to decide how to re-acquire the object.

Mission Autonomy. Finally, there are systems that can autonomously generate whole mission-level, complex, multiple-stage, dynamically adaptive flight plans based on data collected in real time. For example, some systems can autonomously fly an initial simple grid pattern over a complex environment, such as a vehicle accident, to build a rough 3D map. While the UAV is still in the air, it can analyze the rough map to understand the location of various objects in the environment and then autonomously develop a flight plan and fly that plan to gather more detailed data to generate a comprehensive 3D map. As it does this, it might continue to analyze its map and adjust its mission plan accordingly, perhaps even scheduling time to return to its base to replace its battery. In such a case, the human operator might only tell the system the boundaries of the location to map and the desired level of fidelity; the system makes all subsequent decisions by itself.

As we will discuss later in this document, the integration of additional communications roles into autonomous UAS will require additional levels of autonomy that will build on those mentioned above.

SAE Levels of Autonomy. It is important to distinguish between the levels of autonomy mentioned above, which divide autonomy into levels of control abstraction and complexity, and levels of autonomy that describe the relative need for human intervention in the system as a whole. While there is often some correlation between the two, they represent different concepts. Perhaps the most commonly cited is outlined by the Society of Automobile Engineers (SAE) International in Taxonomy and Definitions of Terms Relating to Driving Automation Systems for On-Road Motor Vehicles J3016_202104²⁶ (J3016). This reference is often called the “SAE Levels of Autonomy.” Although J3016 was originally developed in the context of self-driving cars, these levels are often adapted for other forms of autonomous systems, including UAVs.

The J3016 levels range from Level 0, meaning no automation, to Level 5, meaning full automation. Of course, a Level 0 vehicle can still have autonomy at the lower levels, as we refer to them in this document. For instance, according to J3016, a Level 0 vehicle might still have autonomous systems that sense obstacles and apply the brakes or adjust the steering to avoid an accident.

It is also essential to consider the complexity of the application and environment when adapting J3016 to other applications. For example, a Level 5 autonomous vehicle, according to J3016, is specific to vehicles driving in a “conventional” manner on publicly accessible roadways. The Level 5 vehicle may not know how to drive forward along a four-wheel-drive trail, nor is it expected to know how to back up a large trailer on a roadway (even if it may be capable of doing so).

When applying SAE Levels of Autonomy to describe, for example, UAVs used as communications relays or as systems that otherwise provide information to multiple personnel on the ground, it is vital to do so in the context of a well-defined mission and environment. A UAV that only needs to fly to and hover at a fixed location during the day, in good weather, with a clear path between its launch/land point and that fixed location in the sky, could claim Level 5 autonomy with fairly modest capabilities. In contrast, a UAV that must maintain relay communications to multiple mobile ground stations in a mountainous wildfire scene with challenging terrain, smoke, and wind would need very advanced capabilities to be considered Level 5 Autonomous. In this scenario, it would need sensing and AI capabilities to determine the best location to be in as the ground stations moved, avoid terrain and other obstacles, and respond to changes in local weather due to the fire itself.

Management of Risk

NIST SP800-39 defines risk as follows:

“A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence.”

Risk management is further defined as:

“The process of identifying, assessing, and responding to risk.”

²⁶ Society of Automobile Engineers (SAE) International, “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_202104,” 2021-04-30, https://www.sae.org/standards/content/j3016_202104/

Risk management is not necessarily (or even usually) about purely minimizing a set of risks associated with a given activity. Rather, it is more nuanced and includes such tasks as enumerating and characterizing the different risks associated with the activity relative to similar alternatives, comparing them to the risk tolerance of the given application in consideration of the benefits, including possible reduction in other risks, determining controls that are worth applying, and deciding if the residual risk of that activity is acceptable.

The public safety community is more aware than most of the need for risk management. Public safety personnel routinely make on-the-spot decisions as to the risks to take in the moment, accounting for a wide variety of unusual situations, be it the decision to enter a burning building or to drive at high speed the wrong way down a busy street. Coupled with this is an innate sense of factors that manage this risk, many of which are pre-prepared and some of which are societal. For example, when public safety personnel decide to get to the burning building quickly by driving the wrong way down a street, society participates in this risk management by getting out of the way when we see the flashing lights and hear the sirens.

As already defined, risk management requires identifying, assessing, and responding to risk. These three tasks are usually well-defined for well-understood domains, such as driving at high speed the wrong way down a busy street. Even with the advent of vehicles with increasingly autonomous features, such as adaptive cruise control on public safety vehicles, regulations regarding minimum levels of performance usually allow traditional risk assessments to continue to be relevant²⁷, in large part because the human is still responsible and maintaining situational awareness. There are second-order risks associated with humans beginning to rely (perhaps too much) on some of these assistive autonomous features; however, thus far, these have not risen to a significant enough level to induce a major change in risk management culture.

The different levels of autonomy and the associated implementations of AI significantly influence the options available for managing risk. This is particularly the case in new domains where there are not commonly accepted standards of performance, requirements, or even a common language with which to express requirements and capabilities. This results in an inability to appropriately identify the risks, assess their impacts, and respond to them accordingly. Where risks are underestimated, personnel are likely to unknowingly take excessive risks. In contrast, over-estimating risks can result in not deploying capabilities or deploying other capabilities that may be less appropriate, shifting risk to other parts of the operation.

In the process of increasingly integrating UAS into public safety communications infrastructure, it is critical that public safety is able to appropriately identify the associated risks, assess their impacts, and respond to them. While these new technologies provide new capabilities and reduce other risks, they can impart new risks to existing communications infrastructure.

Governance for Uncrewed Systems

Governance is the collection of systems of processes, mechanisms, and rules by which an organization is controlled and operates. It ensures that the people within the organization are held accountable when moving towards organizational objectives. Appropriate governance bridges across the organization as well as up and down its hierarchy. It is a traceable collection of artifacts linking organizational objectives with processes, procedures, and technology. It also needs to flow in both directions, with approvals informing teams and metrics reported to determine effectiveness. The focus on organizational objectives amalgamates what is important to the

²⁷ The behavior of new technologies in other vehicles in the environment is a different question, as evidenced by recent problems with public safety interactions with autonomous vehicles and electric vehicle fires.

organization, such as social considerations, regulatory requirements, technology, and organizational skills and capabilities.

Governance is a key pillar when it comes to risk management including the use of UAS for public safety response. It is necessary to have in place governance that defines and helps personnel responding to incidents to make the decision to deploy UAS, the actions they should take when using UAS, and when to recall UAS. When an organization is tasked with operating in a real-time environment where there are human safety aspects, governance is critical in ensuring that all personnel operate knowing what others are doing, even with degraded capabilities. For example, in a corporate setting, not having a continuity plan for the loss of network connectivity can be inconvenient, whereas in an emergency response situation, not having network connectivity can be the difference between life and death.

By creating and operationalizing appropriate governance, an organization achieves risk management in two aspects: organizational and operational. First, the organizational aspect is with knowing that the UAS is used in accordance with organizational objectives. It will support and enhance and not inadvertently hinder a mission. Defining what it can and can't be used for buffers the organization from aspects related to inappropriate deployment and use. It allows for monitoring of the UAS as an appropriate tool in the organization's arsenal and fine-tuning of its capabilities, configuration and deployment within operation sets.

The second is to manage the risks associated with the deployment and use of the UAS in the field by establishing pre-determined and understood operating bounds, expectations and processes for all personnel involved. It removes ambiguity and assumptions between operators who may not be able to directly communicate with each other. It allows incident coordinators and commanders to have clear expectations when directing and coordinating between teams. By having such aspects already documented, precious time is saved in briefings and response set-ups. It also saves precious time and avoids confusion during an incident response when likely disruptive events (such as loss of communications) can occur as there are pre-defined and understood continuity and work-around processes.

To achieve good eventual governance of UAS, the governance process must start when debating whether and what systems to acquire. A clear goal must be the alignment of usage with the organization's key objectives in public safety. Capabilities and applications described in this paper and others that are not able to be covered that are relevant to an organization should be considered when assessing the bounds an organization sets for UAS, which then feed into acquisition requirements. Once selected and procured, governance artifacts such as policies, plans and procedures are needed to support and implement the governance. By assessing capabilities, application and alignment with objectives prior to acquisition and commissioning, it ensures the UAS is fit for purpose and the personnel can consistently make decisions regarding the use of UAS, all personnel have pre-defined knowledge and expectations of the UAS, and everyone is clear on when to use and when to disengage UAS from the field.

OPPORTUNITIES AND CHALLENGES

As introduced earlier, we will consider two mutually compatible forms of integration of UAS into public safety communications. The first is enabling the UAV to communicate information it has gathered from its sensors, perhaps with some processing, directly with personnel in its vicinity, rather than needing to first transmit this through its OCU. The second is using the UAV as a communications relay, receiving information and transmitting it onward.

UAV-to-Personnel Communications

Some existing UAVs are able to communicate with personnel in the field (and, in the future, other devices such as robots and IoT sensors). Direct UAV-to-personnel communications enable a variety of novel public safety use cases. As previously discussed, at its simplest, this allows personnel in the vicinity of the UAV to observe what the UAV is observing without those personnel needing to communicate with the OCU. This is often the case for analog video transmissions, where any appropriately configured device can also observe the video stream.

Direct UAV-to-personnel communications also enable more complex use cases. For example, in a fast-moving response situation, it may be advantageous for personnel on the ground to call for an autonomous UAV to make an observation, such as looking in a window or providing an overhead view. This use case involves the ad-hoc, dynamic assignment of a UAV to personnel requesting particular data. While the person making the request might be considered the UAV operator, they are providing relatively abstract, mission-level instructions (such as “look in that window” or “give me an overview of this area”) rather than low-level flight commands (such as “fly forward” or “turn left”). The person is also only in command of the UAV for a short part of the flight.

One could imagine a fleet of UAVs stationed or in flight in the response area, waiting to service such requests, somewhat like a temporary, on-scene DFR system. While waiting, they might even perform background tasks, such as acting as communications relays or building more detailed 3D maps of the environment.

However, existing widespread methods pose some challenges to more general adoption. In this section, we will describe how such communications usually happen with most currently available UAS. We will then outline how this is being achieved with some of the latest systems and systems of the future.

Existing UAV to Personnel Communications. Having the UAV communicate with other personnel is perhaps the simplest use case, and, in fact, many early public safety UAVs and even some current smaller and lower cost UAVs have such a capability. These UAVs broadcast their video feeds, similar to how a radio or television station might broadcast to anyone with a receiver. The UAV’s imagery can be received by any device that is using the same frequency and, if the broadcast is scrambled, has the descrambling key. The video feeds through such systems tended to be lower in resolution and were subject to interference that degraded image quality. However, they did have the advantage of gradually degrading image quality as signal quality decreased rather than suddenly cutting out, as was the case with legacy digital systems. Nowadays, such analog systems are limited to smaller and lower-cost UAVs, particularly ones for indoor use, as these systems tend to be smaller and lighter. The sudden variations in radio link quality in indoor environments due to obstructions can also make the gradual degradation in image quality an advantage. For example, a skilled operator may still be able to make enough sense of a severely degraded image to fly the UAV back to where there is a good signal.

Most modern, larger outdoor UAVs are fully digital systems that communicate to their paired OCUs through a data connection. Using such a data connection is more like receiving video through an on-demand Internet streaming service, where the communication is directed at one receiving entity. In some cases, this is based on conventional 802.11 (WiFi) protocols. Others, such as DJI Ocusync, use their own proprietary methods. Unlike analog systems, which broadcast a signal without any feedback confirming that the signal was received, these digital systems generally establish a two-way data connection between the UAV and its dedicated OCU and often use the same connection for control and transmission of other information, such as telemetry.

Digital systems are more secure and allow for the reliable transmission of greater amounts of data. They can take advantage of advanced underlying radio systems that can account for errors due to the environment, such as interference from other devices in the same radio band and radio reflections from buildings and other structures. They also provide ways to authenticate the identity of the receiver and log who has access to, and control of, the UAV. These data links can be properly encrypted, using techniques that, unlike analog scrambling, require significantly greater resources to compromise.

Some newer and more advanced digital systems also have the advantage of being able to dynamically adapt the amount of data used according to both the needs of the mission and the quality of the communications link. For example, when the UAV is stationary and has a good signal, the picture changes slowly so the UAV can transmit a higher-resolution picture to the OCU. As the signal quality degrades, the UAV can adapt and, depending on the situation, either transmit lower resolution imagery that is still sufficient for control or further reduce the frame rate if the priority is resolution. This brings back some of the advantages that analog systems have, by degrading the image as the signal becomes too poor to support high quality imagery, rather than cutting out completely.

However, by switching from a broadcast to a connection-based protocol, most modern UAVs lose the ability to send the data to multiple receivers with one transmission. Instead, it is typical for the UAV to communicate directly with the OCU, which then transmits the data onwards. The OCU is usually bandwidth-constrained, as it communicates via a cellular modem. It usually transmits the data onwards to a streaming server on the Internet, which then sends the stream to multiple receivers. When an incident commander and other personnel at the site of the response receives a video feed from a UAV on their mobile device, that video is likely to have taken the following path, even if the UAV and operator are right next to them.

- The UAV transmits the video to the OCU via its dedicated radio link.
- The OCU transmits that video via the cellular network to a streaming server hosted by a cloud services company.
- The incident commander's personal device, along with the personal devices of other personnel, make separate connections to the streaming server via a cellular network to download and display the video.

Multiple data transmissions are not a major issue in places with good signal coverage. However, where cellular infrastructure is already strained, or where there are only a few temporary cellular towers present, multiple transmissions present a real challenge, both for the UAS and for other users who need to communicate over the same airwaves. Furthermore, this also means that the UAV can only communicate with users who are either within range of the OCU or the cellular network. In a response situation, it is possible that public safety personnel on the ground have no signal from the cellular network or the OCU and yet are in proximity to the UAV.

UAV to Cellular Network Communications. UAVs are starting to become available with cellular modems that can connect to one or more existing cellular network providers. Data from the UAV is sent directly to the cellular network rather than through a single OCU. In many urban applications, and particularly for applications like DFR, the service area has reliable 4G or 5G data coverage. This makes cellular networks an attractive option for use cases where it is safe to assume that a reliable network is present, enabling Beyond Visual Line of Sight (BVLOS) operation anywhere in the covered area. This is particularly advantageous in areas with taller buildings that would block the signal between the UAV and a single OCU.

Even in situations where the network may be stressed, direct-to-cell network communications can still have an advantage over routing communications via the OCU. The primary advantage is avoiding the need to maintain a radio connection with the OCU, which may be impossible, or require much higher levels of power due to distance or buildings. For DFR applications in urban environments, where there may be many small base stations serving a large population, it is quite possible that the closest cellular tower may be closer than the OCU. Thus, the UAV may require less power to connect to the cellular network than to connect directly to the OCU.

A second-order advantage of having the UAV communicate directly with the cellular network is that it can relieve some of the congestion of the airwaves in the vicinity of the OCU. This may be co-located with other emergency response personnel who also require significant quantities of bandwidth, both on the cellular network and in the ISM or dedicated public safety frequencies within which the UAVs operate. Having the UAV communicate directly to the cell towers, and potentially different towers than the ones near the OCU, can reduce contention for this limited resource, especially if the connection through the cellular network back to the OCU is shared among personnel in the area.

One of the primary disadvantages of having the cellular network as the primary communications channel for the UAV is that of latency. As the imagery must now travel from the UAV to the cellular network, to a server somewhere on the Internet, and then to the OCU, most likely via another cellular connection, significant delays can occur between an event occurring in the view of the UAV and the operator seeing it on their screen. Similarly, the operator's commands will also be delayed in being received by the UAV. Worse yet, this delay can vary over time, sometimes rapidly, due to network congestion, particularly in networks where public safety traffic does not have priority.

This delay, particularly if it varies randomly, greatly increases the difficulty and risk involved in having the operator perform tasks that require a timely response, such as avoiding obstacles or reacting to wind and other disturbances. Therefore, such UAVs will need to rely more heavily on autonomy to operate safely. In particular, they will need to avoid or adapt to fast-changing situations that require actions that are infeasible for an operator to address through a high-latency connection. Understanding the implications of such communications systems on AI risk is an important part of assessing the trade-offs of adopting AI technology.

The use of cellular networks also changes the failure modes associated with the communications between the UAV and OCU. On the one hand, the range of the UAV is likely to be increased for a similar amount of power required, as the UAV can connect to towers that are closer to it than the OCU or to towers that are not obstructed by buildings or other structures. On the other hand, the involvement of additional systems in the communication between the UAV and OCU increases the probability of communication dropouts that personnel on the ground have very little control or visibility over. Addressing this risk likewise requires a UAV that is capable of reliably and autonomously handling fast-changing situations while the operator is unable to communicate with it.

Finally, there are potential problems associated with the fact that current cellular networks are simply not expecting to connect to user equipment (UE) that is flying in the sky. Cellular networks are highly optimized to cover specific areas of service, generally on the ground and in buildings. This optimization ranges from the antennas, which can be highly directional to targeted at these service areas, to the algorithms used to hand connections off between towers based on when a particular piece of UE is likely able to see another tower, and the calculations that the network makes to determine how much power to use to communicate with a given piece of UE in a given location.

While UAVs flying in the service area can communicate with the towers, the fact that they are flying in the sky, in places that the cellular network does not expect them and was not optimized for, can cause problems. For example:

- Cellular networks closely manage power levels used to communicate with UE based on specific assumptions to avoid unnecessary interference and optimize bandwidth. These calculations can be incorrect for UAVs that appear in locations that the network does not expect.
- The antennas on cell towers can be highly directional and are typically pointed downward slightly as it expects UE to be at ground level. This can result in poorer-than-expected reception for a UAV that is flying higher than the tower.

UAV Direct Communications. There are also increasing options for UAVs with digital communications to directly communicate with personnel and devices in the vicinity. In its simplest case, multiple devices, be it multiple OCUs or a combination of OCUs and other devices, can communicate directly with the UAV. In applications that require only a low amount of bandwidth, such as the UAV gathering data from IoT sensors in the field or receiving commands from users, this is generally straightforward.

Where there is a requirement for higher bandwidth, such as streaming video, there are several options available. The simplest is to have the UAV make a direct connection to each recipient of the video. This allows each recipient to have a unique stream but requires significant bandwidth and computation power on the part of the UAV, as it must send the data multiple times. Alternatively, it may require that the image quality be degraded due to limited bandwidth and computation.

More advanced systems may be able to broadcast or multicast the same video stream to several recipients in the field. This means that each recipient receives the same video, but it also means that the UAV needs to send the video stream only once.

UAV as Communications Relays

In addition to the UAV communicating directly to other entities for the purpose of its own mission, UAVs can be used to relay communications from other entities. At its simplest, the UAV could be a fixed communications repeater, with connections to the two people or devices that need to communicate. This setup takes advantage of the UAV's elevated position to "bounce" the signal around an obstacle or over a distance that would otherwise prevent the two people from directly communicating. In such an application, the UAV would simply need to station itself at a fixed location and perhaps return to its base occasionally to replace batteries or refuel. Two UAVs can also "take turns" so that one is always in the correct location while the other is returning for a battery change or refueling. Alternatively, a tethered UAV could stay in one location for an extended period of time. Such a paradigm can involve relatively little autonomy, especially if a human operator is closely monitoring the mission.

A significantly more complex example would be the use of multiple UAVs, each equipped with MANET radios, that communicate with each other and with multiple radios on the ground, be it personnel or devices. Such a system can be an advantage, both in increasing the possible communications and coverage range and also in increasing available bandwidth. This can be an advantage in a congested radio environment, even if the people and devices could otherwise communicate with each other. This is because if a radio needs to run at high power to reach the desired receiver, it takes up that frequency across a large area of space, much like a person shouting across a room interferes with all of the conversations in that room. In contrast, a radio running at lower power, talking to a nearby UAV, that then relays that data onwards to the next

closest UAV and so on to the receiver, allows other nearby communications to occur on the same frequency, much like how many people can talk quietly in a room without interfering with each other, even if some of them are passing messages along.

More advanced use cases include using a combination of technologies. For example, in a location that has no cellular network coverage due to remoteness or damage to infrastructure, it is possible to place a cellular network antenna on a large UAV to which conventional cell phones can connect. The UAV can then connect to a MANET or other type of network that provides connectivity to the outside world²⁸. The UAV becomes a “flying cell tower”, providing coverage to public safety in the critical initial period of a response before additional temporary structure is deployed.

While such a system can be effective with fixed locations for the UAVs, in the future, tight integration between advanced autonomy on the UAVs and the radio systems can also allow the UAVs to dynamically station themselves in the locations where they can be most effective given the changing communications needs of the network, the movement of the radios on the ground, and the need for the UAVs to take turns returning home for battery changes.

Applications and Benefits. To analyze the benefits and risks of using UAS as part of a MANET or as a data ferry, let’s give an example using a high-altitude mountainous recreation area. Many of these locations are often remote and lack the infrastructure to provide data or voice communications. Typically, these are purposefully designated and protected areas, such as national parks and wilderness areas, oriented to preserve the natural environment. As such, the buildout of any infrastructure is limited or prohibited. Not only are broadband cellular communications unavailable, but first responder LMR is often limited or non-existent in these areas. As such, responders typically depend on VHF wireless communications to relay messages and information.

First responders operating in wilderness areas contend with a variety of incidents, such as search and rescue missions, drownings, exposure to elements, medical emergencies, criminal activities, and natural or human-made events. Because of this, the first responders require specialized skills, training, and communications tools, such as UAS. UAS aid in search and rescue missions and can also relay communications in difficult terrain. In-flight UAV with data ferry and/or MANET capabilities can enhance information exchange for responders and incident commanders in extreme environments. With enough resources, MANET provides network redundancy, load balancing, reliability, and quality communications. It is able to do so dynamically, as nodes, such as UAVs, vehicles, and personnel with radios, move in and out of range of each other.

The interconnectedness provided by MANET allows responders to share resources that one node can access across all nodes in the network. For example, a UAV with a MANET node and a cellular modem can use its extra height to reach a cellular network tower and obtain internet access. In that case, cellular access can then be provided to personnel, UAVs, and other equipment in the vicinity who may be unable to access that tower. This allows personnel to communicate more widely and access crucial resources such as databases, weather reports, and Common Operating Picture (COP) tools. It also enables more advanced AI features on UAVs and other equipment that benefit from access to the internet, be it to access additional data and computation or share the data with other personnel not in the immediate vicinity.

²⁸ Knight, R. “Restoring Communication After a Disaster,” *Inside Unmanned Systems*, 2023-12-13, <https://insideunmannedsystems.com/restoring-communication-after-a-disaster/>

MANET also brings with it new risks beyond those of traditional wireless networking systems. For example, a variety of Denial of Service attacks are possible on MANET that can inhibit all network communications. This is particularly problematic for attacks that target the dynamic reconfiguration capabilities of the network. Latency can also be an issue, as data must be passed from one node to the next, each adding some delay. This can be exacerbated with the use of UAVs, where a UAV leaving the network to recharge might mean that the network needs to re-route data through longer paths. Worse yet, if the UAV happens to have been the single path between two sections of the network, it can result in the network splitting in two, at least until other nodes can span the resulting gap. To an external user, it may not be obvious that a single node might be the only path between two parts of the network and that it should be protected. Due to its additional workload, this single node may also deplete its battery more quickly.

In situations where MANET is not feasible or practical, such as where there are not enough UAVs or other nodes to maintain continuous coverage, a UAV can be used as a data ferry to collect information by periodically flying over an area and relaying data back to mission commanders. Although this approach can still accomplish most of the aforementioned applications, it may not provide real-time services.

A data ferry is a good solution for applications, such as periodically gathering data from IoT sensors, that may not be time sensitive but require regular updates. The periodic nature also means that low powered sensors need only connect and broadcast data to the UAV when it is nearby, reducing its overall power use. In a public safety context, this can be particularly valuable for a future generation of deployable sensors in places such as unstable buildings or wildland fires. The lack of a continuous connection may also be an advantage in an adversarial situation, such as in law enforcement, as it becomes much more difficult for an attacker to use the periodic connectivity to enter and attack other systems.

Data ferry solutions also carry risks that include single-point-of-failure of the UAV, the lack of real-time connection, and the lack of bandwidth. For applications where the data is sensitive, the potential for the UAV to fall into malicious hands is also a concern. Due to the significantly larger quantities of data that might be carried by a data ferry UAV, particular care needs to be taken to encrypt data at rest (when stored on-board the UAV) as well as in flight (during transmission), and to properly authenticate legitimate devices and users.

Managing AI Risks

To fully and appropriately take advantage of these novel communications paradigms, it is necessary for the UAVs to exhibit at least task, if not mission, levels of autonomy. For example, even if the mission is to fly to and stay in one location, increased communications delays between the UAV and OCU due to the use of MANET or cellular networks means that the UAV must be able to autonomously handle fast-changing situations and respond to situations like other flying objects in the area (including birds), sudden gusts of wind, or unexpected obstacles in the path. More advanced use cases require one or more UAVs to decide, within their mission area, the best locations where they can be most useful, perhaps making decisions that even human operators would not have the information to make.

The NIST AI Risk Management²⁹ Framework (AI RMF)³⁰ and Playbook³¹ are tools that can be used to help manage the risks associated with transferring more of these responsibilities onto AI systems. While profiles that directly relate the AI RMF to the UAS for Public Safety application are the topic of future work, the core playbook is still a useful resource to guide the risk management process. In particular, it highlights where potential risks lie and facilitates more informative discussions with vendors by focusing on the key aspects of risk associated with these systems.

Adopting the AI RMF can also provide organizations with renewed visibility into other risks that may exist elsewhere in their operations with which they may not be aware. Much of the AI RMF is not specific to AI, particularly when it comes to governance.

The AI RMF divides the management of AI risk into four functions: Govern, Map, Measure, and Manage.

The Govern function focuses attention on cultivating an appropriate culture of risk management across the development, deployment, evaluation, and acquisition of systems that include AI. It also outlines processes, policies, procedures, and documents for the management of these risks and provides structure that helps align AI risk management functions with those of the organization. This includes human factors such as training, legal issues, and consideration for the unexpected, such as incident detection, continuity, response and recovery when an AI or similar system is embedded into a process. Public safety organizations who wish to deploy increasingly autonomous UAS may not have had to integrate AI concerns into their existing organizational risk management policies. The Govern function helps to provide structure to assist with this integration.

The Map function contextualizes the framing of risk relating to the AI system. It helps to enumerate the many different processes involved in the system, including both the UAS and systems that connect to it, and to determine where AI components may exist, their effect, and their scope. It also helps users to identify possible interactions between these systems, AI and otherwise, that may pose additional sources of risk.

The Measure function highlights the tools that are available to assess the various risks surfaced by the Map function. It also provides insight into what these measurements mean and where the state-of-the-art might be deficient.

Finally, the Manage function takes the identified and measured risks and helps to determine appropriate actions to take based on the outcomes of the Govern function. Public safety organizations are particularly well placed to adopt the Manage function, for many of their activities are generally considered too risky for the general public. It is generally a bad idea to run into a burning building or drive a large vehicle the wrong way down Main Street. Public safety personnel sometimes have to manage the risks associated with doing such activities to avoid a potentially bigger risk of not undertaking those activities. The Manage function, informed by the Map and Measure functions, helps to develop policies that enable this nuanced approach, whereby riskier activities are permitted in appropriate situations, with suitably controlled risk.

²⁹ National Institute of Standards and Technology (NIST), “NIST Trustworthy & Responsible AI Resource Center,” accessed 2024-04-03, <https://airc.nist.gov/>

³⁰ National Institute of Standards and Technology (NIST), “NIST AI 100-1 Artificial Intelligence Risk Management Framework (AI RMF 1.0),” 2023-01-26, <https://doi.org/10.6028/NIST.AI.100-1>

³¹ National Institute of Standards and Technology (NIST), “NIST AI RMF Playbook”, accessed 2024-04-03, https://airc.nist.gov/AI_RMFKnowledgeBase/Playbook

Managing Cybersecurity Risks

Like the AI RMF, the NIST Cybersecurity Framework (CSF) 2.0³² is designed to help organizations understand and manage their cybersecurity risks. As UAS become increasingly connected, and in novel ways, they expose organizations to new risks that they may not have been aware of. Furthermore, the increasing connectedness and advances in technology more generally also means that organizations are exposed to a much wider range of cybersecurity threats. Adversaries are no longer individuals or entities who are trying to break into a specific organization. Ransomware targets anyone who values their data or system more than their money. Hactivists target anyone who can provide visibility to their cause. Any number of anarchy-focused individuals and groups simply want to cause disruption regardless of target. Frameworks like CSF 2.0 highlight the space that organizations need to consider and provide structure to guide the management of these risks.

Unlike the preceding AI discussion, it is likely that organizations adopting highly connected, increasingly autonomous UAS will already have cybersecurity policies and regulations, such as the Criminal Justice Information Services (CJIS) Security Policy, NIST SP 800-53 Rev. 5 Security and Privacy Controls for Information Systems and Organizations, or NIST SP 800-171 Rev. 2 Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations. However, it is very likely that the existing policies do not account for the novel requirements, capabilities, and risks associated with highly connected, increasingly autonomous UAS (or, indeed, UAS or cyberphysical systems in general). Most cybersecurity risk management policies are geared towards IT systems in an office environment, perhaps with some limited exceptions for field use.

The CSF 2.0 is not another cybersecurity checklist. Rather, in the context of integrating these new technologies, it provides a systematic, structured way of analyzing an organization's existing and desired security postures, measures that might be considered to control the risks associated with these new technologies, and how security policies might be adapted to accommodate these while continuing to appropriately manage risk.

For example, highly connected UAS are increasingly being connected to critical systems, such as Computer Aided Dispatch (CAD), Geographical Information Systems (GIS), and Common Operating Picture (COP) systems, either directly or through intermediaries. Compromising critical systems can result in very significant disruption and risk to public safety organizations, even if the attackers do nothing beyond forcing the system offline³³.

They also connect to local systems that enable real-time streaming of videos and other information to the incident commander and other personnel. Such links are increasingly expected by personnel and management as the benefits of these technologies become clear. However, these systems often do not satisfy the cybersecurity policies of the networks to which they connect. Some of the unique aspects of highly connected UAS that often do not fit well with conventional cybersecurity policies include the following.

³² National Institute of Standards and Technology (NIST), "NIST Computer Security Resource Center," accessed 2024-04-03, <https://csrc.nist.gov/>

³³ Rector, K. "Hack of Baltimore's 911 dispatch system was ransomware attack, city officials say," *The Baltimore Sun*, updated 2019-08-07, <https://www.baltimoresun.com/2018/03/28/hack-of-baltimores-911-dispatch-system-was-ransomware-attack-city-officials-say/>

- Lack of traditional cyber confidentiality controls. An emphasis on the availability and real-time processing on potentially limited hardware resources, resulting in less focus or ability to implement traditional cyber controls such as encryption and authentication.
- Non-standard operating system setup. The computer systems in the UAV, OCU, and other components of these advanced UAS are often embedded computers, running operating systems based on Android or Linux. However, these are often heavily customized and proprietary.
- Proprietary network protocols or messages. Similar to customized and proprietary operating systems, many such devices utilize proprietary network communications to maximize performance across potentially unreliable connections. This inhibits the ability for traditional network controls, monitoring, and securing.
- Security updates. Due to the high level of customization and the need for soft or hard real-time operation, updates to the software and operating system require extensive testing. This means that security updates for the underlying operating system are likely to be delayed or unavailable from the manufacturer and difficult or impossible for the end user to apply themselves.
- Inability to run management software. These computer systems are often heavily locked down by the manufacturer and cannot run other software. Safe physical operation of the system relies on the manufacturer having certainty about the performance and reliability of the system, which precludes the running of software that the manufacturer has not tested.
- No access for organizational network security personnel and retention of manufacturer's default access tokens. For the same reason, administrative access to these computer systems is often not available to the organization's network security personnel to prevent changes that may affect the safe operation of the UAS. While the network security personnel may be IT and security experts, they may not necessarily be experts in real-time system control or the UAS's internal safety measures. The manufacturer may also retain a level of administrative access tokens on the UAS for maintenance in warranty or repair situations which can't be changed by the organization.
- No physical access. The purpose of a UAV is to go where personnel cannot; therefore, by definition, it will spend much of its operating life outside the direct physical possession of organization personnel.
- Security theater. The novelty of the field makes it easier for better known security risks, which themselves may or may not have some element of legitimacy, to become over-emphasized. This draws attention and resources away from other security concerns that may be less well understood. For example, security controls relating to country-of-origin can help address the issue of state-sponsored adversaries. However, there are still many cybersecurity concerns that affect all systems, regardless of country-of-origin. An over-emphasis on country-of-origin controls can result in these other risks being de-emphasized. At the very least, it is necessary to acknowledge that even if the well known security issues are controlled, there are still risks across the entire system that need to be otherwise mitigated or accepted.

Organizations with security policies that do not include the unique needs of UAS and other cyberphysical systems end up having these systems function under exemptions, formal or otherwise. This results in ad-hoc, or no, oversight or visibility into these systems by the organization's cybersecurity team.

Relaxing these security policies to such systems without increasing actual cybersecurity risk exposure can be a difficult process. For every control that is loosened, additional controls need to

be put in place to manage that risk. For example, adopting UAS may require additional segregation of networks to make it more difficult for a malicious actor to move laterally from a compromised UAS into other critical systems.

However, it is vital that this be done in such a way as to maintain the benefits of these advanced UAS and their novel communications capabilities; otherwise, the organization risks falling into a “worst of both worlds” situation. The CSF allows organizations to analyze how these novel systems can fit within their existing risk management. It provides assistance for the development of procedures and policies for their appropriately risk-managed use and deployment, helping to identify where these additional controls might be placed to manage risk while allowing the benefits.

Governance Issues for Connected UAS

Governance relates to the human processes and expectations around using UAS within the operating objectives for public safety response. The challenge is to be appropriate to the intended application of UAS and not miss any aspects brought on by the novel application and use of UAS. Governance concerning UAS, when done correctly, is something that is used as one of the many inputs into the system - so that when it is operational, the personnel using it know the system performs in accordance with their processes and procedures.

In terms of governance, a starting point is to think of an intelligent UAS as a “new team member,” who is unfamiliar with existing teams, processes, practices, and expectations. This is reasonable particularly if the new AI capabilities of the UAS allow it to take over decision making tasks that were previously the responsibility of a human. However, there are some significant differences to consider. For example, there are basic assumptions around common sense, ethical standards, and cultural values that may need to be made explicit. On the other hand, there are concerns that do not apply to AI systems, such as the need for organizational processes to be “fair” to the AI system.

It is important to create governance artifacts that document these governance decisions. This ensures that the UAS is appropriately specified and configured (in verification and validation testing) and that the UAS is provided with appropriate inputs and used in the appropriate manner, consistent with these governance decisions. These artifacts are also important for incident and activity review, where opaque AI may preclude easy root cause analysis without the ability to analyze the initial inputs and explicit instructions.

For example, a UAS may be directly communicating with personnel and making determinations on what information to disseminate from its potentially vast set of information. Governance would drive the policy that dictates how much, when and to whom information is disseminated. An inappropriate policy creates the risk that information is disseminated to the wrong people, or is irrelevant or incomplete for their situation.

Governance may be at a more abstract level and provide Incident Commanders the flexibility to determine how the UAS should disseminate its information dynamically. However, too much flexibility runs the risk that the UAS is asked to perform beyond its capabilities. From the personnel’s point of view, governance sets their expectations on the information available to them during an activity. If the UAS behaves in a manner inconsistent with governance, such as due to situations that were unforeseen in the process of developing the governance, user expectations may not be met.

Another example of a factor to consider when developing governance for autonomous systems that disseminate information is the need to consider the human effect of information. Responders may receive information of a distressing nature, such as news that their properties or loved ones

have been affected by the incident to which they are responding. Previously, such information might have been disseminated by a human, who could overlay judgment and take additional measures to address the distress that this may cause. For example, they may choose an appropriate time when support is available. Automated dissemination of such information, by default, has no such filter. An influx of information can also distract personnel or cause information overload. While it is possible to autonomously curate information, inappropriately curated information can also lead to incorrect decision making. Appropriate governance is necessary to develop consistent organizational (and, by extension, system) policies to address these issues.

The use of UAS as part of the communications infrastructure also presents new governance challenges. While UAS bring considerable advantages to a wide variety of public safety applications, they also bring significant risks for which the management is still immature. Governance decisions need to be made to determine when it is appropriate to use such systems and if the resources necessary to control the risks are appropriate. If the UAS is performing other activities in addition to being a communications relay, how is it going to decide which activity to continue performing if there is contention? For example, a UAS may be tasked with keeping watch on a moving firefront. To continue performing this task, it may need to move from its position and disrupt communications. It is sometimes hard enough for a human to make a decision where the consequences either way are either unknown or can be catastrophic.

At an organizational governance level, it is important that there are checks and balances to ensure underlying plans and procedures align with the overarching organizational policies and rules. This is not just for new technologies like highly connected UAS, but also for all forms of technology and is generally good governance practice anyway. Governance artifacts should be harmonized or, at least, compatible with each other and fit with other existing processes. Even without the introduction of new technology, governance artifacts can become a patchwork, written at different points in time and by different departments. The introduction of new and novel technology can exacerbate this patchwork. Two risks in particular arise from incompatible procedures across an organization.

The first results from confusion, where it is unclear which procedures might apply in a given situation. For example, it can be unclear what procedure might apply to a cellphone that is also used to control a UAS. It might be a personal or corporate managed device that is temporarily connected to the OCU for controlling the UAS. Alternatively, it may be a cellphone that is built into the OCU from the factory. The user may not even be aware that the OCU contains a cellphone.

The governance stack has to be, at the same time, consistent and yet also aware of the risks associated with the cellphone as part of the UAS and connected to the network. Problems can arise if the policies and procedures for cellphones were developed without visibility into the organization's UAS program and that cellphones used in the UAS program may have very different requirements to those used elsewhere in the organization. Worse yet, such policies will not account for additional risks due to these cellphones being connected to more systems than is normal. Consider the scenario where a cellphone is connected via the OCU to a UAV. The UAV might then be serving as a node in a MANET carrying other traffic. Even if the intent is for the cellphone to be logically separated from the traffic on the MANET, it is necessary for governance to be aware of the additional risks. This may include an adversary being able to exploit a vulnerability to move laterally between the two systems, an attack that has already been

demonstrated in motor vehicles³⁴. Alternatively, a configuration problem in one could affect the availability and performance of the other.

Secondly, the novel use of new technology could require deviations from well-established existing security processes. It may be impossible to perform important tasks with new technology while following existing policies and procedures. This incompatibility may be due to aspects of the underlying technology, design choices made by the developer, existing policies and procedures from a time when these technologies did not exist, and/or because of the way in which the policies and procedures were written.

Exceptions to existing policies and procedures may increase risk, particularly where they are made informally and not subject to proper risk management and controls. Exceptions can also affect organizational risk culture. The precedent might also cause users to make assumptions and generalizations that if the policy could be “broken” for UAS, they could also break it for all systems.

In the case where the incompatibility is due to the manufacturer’s design decisions, particularly ones that reduce security, it is preferable to write into specifications the need for the system to be compliant with appropriately secure policies and procedures. These specifications eliminate the need for the exception. For example, suppose a UAS might not satisfy an organization’s data encryption requirements for a given application. It would be reasonable to have the manufacturer add the necessary encryption, for the organization to move to a system that does have the requisite encryption, deliberately decide that the additional risk is acceptable (perhaps subject to additional controls), or to forgo the capability altogether.

There may be instances where exceptions are required, such as for a UAV that is particularly small or lower powered and where adding the necessary encryption affects its performance in an unacceptable manner. In this case, UAS specific governance must include a risk assessment process to justify its need, measures that could be taken to otherwise control the risk, approval at the appropriate level of authority to accept any residual risk, clear scope of when it is acceptable, and appropriate logging and monitoring. Periodic review of the exception and repeating the risk assessment and approval should also occur as part of overall risk management, including evaluating if technology has progressed to the point that it may be worth transitioning to a different system that does not require the exception.

CONCLUSION

As UAS become more connected with people and systems beyond just the operators, so too do the challenges of appropriately understanding and managing both the benefits and the risks of employing highly connected UAS. Cybersecurity and AI risks, in particular, tend to be poorly understood and managed. In this paper, we have discussed some of the background of communications and UAS in public safety. We also provided an overview of AI, autonomy, and governance as it applies to UAS and the public safety use case.

We then presented some of the opportunities and challenges in these new, highly connected, increasingly autonomous UAS. We focused on two categories of use cases. The first were use cases where the UAV communicates mission data, such as videos and commands, directly with personnel and other systems in its vicinity rather than going through the OCU. The second were use cases where the UAV serves to relay communications between other people and systems. While there are many other types of use cases, these two encompass a wide range of applications.

³⁴ Greenberg, A., “Hackers Remotely Kill a Jeep on the Highway—With Me in It”, *Wired*, 2015-07-21, <https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/>

We describe some of the specific applications within these categories, followed by specific challenges in cybersecurity and AI risk management, and how frameworks such as the NIST Cybersecurity Framework 2.0 and the NIST AI Risk Management Framework can be helpful tools. We also discuss governance issues that should be considered when conducting a risk-benefit analysis of the use of these systems.

Our efforts in this space are just beginning; much work will be done in the future. We have established a working group to develop guidance and other resources for managing cybersecurity and AI risks of UAS in public safety applications. This effort aims to help all stakeholders, including manufacturers, vendors, developers, researchers, trainers, regulators, insurers, lawyers, and politicians, to understand this space better.

The initial focus of this effort is on the development of resources to help fire chiefs, police chiefs, and other public safety managers ask questions of vendors and other technical people, including guides that help them understand the responses and the ensuing effects on risk management. Much of the information that end users need to make the decisions discussed in this document is not provided by default to end users before or after deployment. Indeed, much of this information is only currently known to those who are directly developing the software systems and may be several steps removed from the end users. Having end users better informed when asking questions of their vendors is an initial step in improving cybersecurity and AI risk management.

We are also developing corresponding resources for other stakeholders in the UAS for Public Safety ecosystem. We will particularly focus on manufacturers and developers to help them better understand the reasons why end users need this information and allow them to proactively surface this information from their organizations and software supply chains. This will also help inform the design of new systems and improve their ability to be appropriately and transparently risk-managed.

Finally, we are developing free, open-source resources for educators, with an initial focus on graduate and advanced undergraduate students in engineering and law. The goal is to make it easier for educators to incorporate cybersecurity and AI risk management into their teaching, with UAS in public safety as an example use case. The hope is to improve the future management of cybersecurity and AI risk further, not just in terms of highly connected UAS for public safety but also in society in general.