Application of Dempster-Shafer Network to Real-Time Unmanned Systems Risk Analysis Framework

Joel Dunham *

Georgia Tech Research Institute, Atlanta, GA, 30318

Dr. Eric Johnson[†]

Pennsylvania State University, City, Province, Zip Code, Country

Dr. Eric Feron[‡]

King Abdullah University of Science and Technology, City, Province, Zip Code, Country

Dr. Brian German[§]

Georgia Institute of Technology, Atlanta, GA, 30318

Unmanned Aerial Systems (UASs) are continuing to proliferate rapidly [1]. Quantitative risk assessment for UAS operations, both *a-priori* and during the operation, are necessary for governing authorities and insurance companies to understand the risks and properly approve operations and assign insurance premiums, respectively. This paper reports the results of the 2018 UAS Safety Symposium held at the Georgia Institute of Technology, which was conducted as part of this research. This symposium was comprised of experts in UAS law, insurance, regulations, operations, and research discussing the direction of the UAS industry and the necessary steps to move the industry forward. Those results motivate the remainder of this research including the novel application of Dempster-Shafer networks using auto-updating transition matrices to the problem of UAS risk analysis and decision-making. This research trains a DS network based on simulated operation data, tests the capabilities of the trained network to make real-time decisions on a small UAS against a baseline system in a representative mission, and explores how this system would extend to the full UAS ecosystem as discussed in the 2018 UAS safety symposium. Conclusions are drawn with respect to the research performed, and additional research tangents are proposed.

Nomenclature

DS = Dempster-Shafer

JARUS = Joint Authority for Rulemaking of Unmanned Systems

*Senior Research Engineer, AIAA Lifetime Member

[†]Professor, Aerospace Engineering, and AIAA Member Grade (if any) for second author.

[‡]Professor, Department Name, and AIAA Member Grade (if any) for third author.

[§]Professor, Aerospace Engineering, and AIAA Member Grade (if any) for fourth author (etc.).

SORA = Specific Operations Risk Assessment

UAS = Unmanned Aerial System

I. Introduction

Unmanned Aerial Systems (UASs) are continuing to proliferate rapidly [1]. Consumer products are becoming pushbutton systems which provide a capability to operators without the operators requiring understanding of how the system works or potential failures. The DJI Mavic Pro [2] represents a near-entry-level consumer product (approximately \$900) with long flight times (approximately 28min), pushbutton operation, automatic handling of common issues such as near seamless switching between GPS navigation and optical flow navigation in the event that the GPS fails, smooth take-offs and landings that do not require user knowledge of vehicle limits, and geo-fencing with knowledge of Temporary Flight Restrictions (TFRs) [2], drastically reducing the requirements for operator knowledge of both current flight conditions and regulations. With this reduction of required operator knowledge and disconnect from direct operator input, it becomes even more necessary for the system to appropriately handle contingencies without input from the operator.

Current efforts are focused on aligning provable safety systems with related regulations to integrate unmanned systems with manned systems in the National Air Space (NAS), allowing them to operate in the same airspaces. In recent years — recent months for many of these advancements — great strides have been made both on the regulatory side and also on the safety technology side. Several risk analysis frameworks have been proposed that have backing from one or more regulatory bodies. Joint Authorities for Rulemaking of Unmanned Systems (JARUS) guidelines on Specific Operations Risk Assessment (SORA) [3] develops a framework in which Unmanned Aerial Systems (UAS) can fly specific operations in a variety of airspaces once preliminary risk assessment and, if necessary, risk mitigations have occurred and been approved by the governing authority. This framework is flexible and enables a UAS operator to tailor the approach of risk-based operation approval to almost any situation. However, while flexible, many gaps still remain, precluding access to many specific operations since the risk reduction requirements are too demanding for current technologies or methods of risk reduction to meet. Multiple Federal Aviation Administration (FAA) frameworks [4] [5] have been proposed, but these are all-encompassing frameworks which do not provide the low-level details necessary to make a risk assessment for particular operations. Together, these frameworks and philosophies form a risk management-based approach to aviation safety which could provide the regulatory backing to enable UAS integration into NAS.

Technological solutions to safety issues must be consistent with the regulations used to enforce them. Previous methodologies focused on specific technology requirements to access airspace. For example, sense-and-avoid was seen as a must-have capability for UASs intent on flying in the NAS [6]. After 20+ years of research with no accepted solution,

the focus has shifted to analyzing the risks associated with the integration of various technologies and determining whether the risk is sufficiently low to perform the given operation or flight plan, such as in the case of SORA [3] or the current FAA focus [7]. The challenge partially shifts to the technologies providing the necessary data on whether they are reliable. For example, fault detection systems are a large field in aerospace research. Concepts such as neural networks for detecting faults in IMUs [8] enable fault detection and handling for a sensor of which failure can cause complete loss of control on a multi-rotor. Sufficient tests of this system result in statistics for the reliability of detection and handling, which provide the necessary data for quantitative risk assessment. Likewise, research on emergency path planning for UAS also exists [9]. These designs provide solutions to the problem of how to respond once the emergency has been identified.

Recent research has enabled offboard real-time risk analysis for UAS, based on current data, Bayes belief networks, and most-likely-hypotheses [10]. Further, consistent separation analysis using avoidance volumes for aircraft based on the precision of navigation systems [11] provide methods of analytically predicting problems and dealing with them as they arise. These advances serve to underscore the importance of technological solutions to safety systems on the UAS that are consistent with the regulations that govern use of the UAS. Limitations in each of these methods provide part of the basis for the research in this proposal. Specifically, a complete onboard real-time system sufficient for small UAS is currently not operational. Further, quantification of risk through a Dempster-Shafer (DS) analysis [12] [13] provides decision criteria beyond the criteria provided by a Bayesian analysis, which better align with risk-based decisions. A scenario has been identified that covers a broad range of possible system failures and degradations ranging from vehicle issues, such as control or navigation system degradations, to sensor degradations due to external effects. The goal of this research is to develop a risk assessment model which connects risk evidence through long-term operation results analysis with real-time predictive actions to mitigate unacceptable risk during UAS operations. A further requirement of this research is that the chosen risk assessment methodology must be flexible to incorporate constantly changing data about UAS risks while the results must also be explainable to governing authorities (i.e. there is a clear relationship between risk factors and resulting risk assessment that a decision-maker can understand and follow), enabling this analysis to be a potential basis for operational authorization.

This paper reports the results of the 2018 UAS Safety Symposium held at the Georgia Institute of Technology, which was conducted as part of this research. Those results motivate the remainder of this research including the novel application of Dempster-Shafer networks using auto-updating transition matrices to the problem of UAS risk analysis and decision-making. This research trains a DS network based on simulated operation data, tests the capabilities of the trained network to make real-time in-mission decisions on a small UAS against a baseline system, and explores how this system would extend to the full UAS ecosystem as discussed in the 2018 UAS safety symposium.

A. Organization

This document is organized as follows. Section II describes the recent state of UAS operations, regulations, and insurance based on expert information obtained through a UAS safety symposium hosted in 2018, along with corroborating literature. Section IV applies the theory extensions to UAS — a news multirotor scenario evaluating the risk associated with remaining on station. Section V details the results of testing the real-time risk assessment and decision system in the scenario described in Section IV. Section VI documents the flight demonstration performed on a small UAS to show feasibility of the system on current platforms, and Section VII details how this work can be extended to apply to the full UAS ecosystem. Section VIII concludes the discussion and provides trajectories of future research to be based on the outcome of this research. Finally, the included appendix IX include notes that are not central to the work, but are necessary for a full understanding.

II. Safety Symposium

In August 2018, a UAS safety symposium was held in Atlanta, Georgia, USA, which included experts from many facets of UAS operations in the United States including law, insurance, operations, research, and regulation. While some changes to the UAS landscape have changed since that symposium, much of the principles remain the same. This chapter details those principles, highlights some of the changes since then, and corroborates some of those expert opinions with additional literature, setting the stage for research which could fit into this ecosystem. Note that this symposium was primarily focused on intended legal operation of UAS; while some information and artifacts were discussed with regards to deliberate, malicious, illegal activities, representatives agreed that a different set of principles were required to handle these activities, and those principles were outside the domain expertise of the representatives in attendance. Appendix IX includes the raw notes from the safety symposium. Sections II.A, II.B, II.C, II.D, and II.E are primarily written based on expert opinions from the safety symposium as documented in Appendix IX. As such, those notes are not constantly referenced throughout the following sections. Corroborating research is cited in these sections.

A. Summary of Safety Symposium

By far, the principle take-away was the lack of information concerning use of UAS legally and within regulations. Interestingly, this area is one in which there was significant divergence between regulators and the rest of the industry. Operator, insurance, and law representatives agreed that there was lack of information concerning UAS regulations and laws, yet the regulatory agency representative provided multiple means of obtaining information about regulations including through Part 107 [14]. Since then, the two ends of the spectrum have grown closer, particularly as the FAA closes towards regulations that will require UAS to identify themselves and transmit location information while operating [7] [15]. Second to this take-away was the consensus that the industry is in wait-and-see mode, especially on the law and regulations side, until precedents are set, usually through a major incident. Again, regulations have been pushing

forward as mandates have been created for the FAA to develop regulations for UAS operations [16]. All representatives agreed that enforcing regulations is difficult at best until a reliable means of identifying UAS is available. As mentioned previously, FAA regulations are going into force that will require operating UAS to transmit ID and location information [7]. Finally, there was agreement that operator education is currently lacking. Improvements have been made for Part 107 operation education [14], but there is limited education beyond that point.

B. United States Separation of Operations

Unmanned systems currently fly under accommodation practices — relying on operational segregation to avoid issues with manned traffic [17]. As such, they are restricted in use by 14 CFR Part 107, which is applicable for commercial use of small UAS under 55 lbs [14]. Note that UAS regulation is a purely federal affair, since the federal government regulations airspace from the tips of the blades of grass exposed to the outside up as high as the airspace extends. This jurisdiction was solidified as of *Boggs versus Meredith* in 2015. Part 107 regulation restricts flight to less than 400 ft AGL unless within 400 ft of a higher structure. Operations up to class B airspace [18] are permitted via Part 107 with Air Traffic Control (ATC) permission or in class G airspace with no permission required. All operations under Part 107 must be within line-of-sight of the certified operator who must retain visual sense-and-avoid capabilities over the aircraft, further limiting the available flight times to daylight or civil twilight and to one pilot per unmanned system. Aircraft flying under Part 107 are not allowed to fly over people not directly involved in the flight of the vehicle. These derived restrictions all come from one basic concept: operational segregation with human decision making as the immediate and final authority for all UAS. These regulations were finalized as of August 29, 2016. While these regulations allow small UAS to operate in national airspace in principle, the practical result is far from the fully integrated vision of researchers, manufacturers, and the FAA alike. Subsequent to this set of regulations, updates have been made, in large part spurred by UAS operations and planned operations that have forced regulations to move forward [16]. Likewise, efforts are being made in partnerships with the FAA and various UAS community stakeholders to develop UAS Traffic Management (UTM) National Campaign II, enabling access to low-altitude airspace for UAS [19]. Various efforts have been made to develop requirements for this airspace integration [20].

For larger UAS — above 55lbs — the only available regulation under which to fly, outside of FAA-designated test sites, is Section 333 [21]. Section 333 deals with waivers/exemptions, otherwise known as Certificate of Waiver or Authorization (COA). This regulation more clearly exemplifies the operational segregation by stating that any unmanned system can be flown in the area under the COA with advanced notice to ATC.

While allowing unmanned research to continue with the goal to develop fully integrated systems, these regulations clearly keep unmanned systems separate from fully integrated national airspace. Due to this segregation, however, unmanned systems do not require any type of certification. Only pre-flight checks by the operator are required under Part 107 [14]. Because many small UAS are never intended for fully integrated flight in national airspace, these regulations

enable low cost entry into the small UAS market, without the need for costly certification processes or pilot training. Conversely, all UAS are limited to these operational restrictions due to the inability to certify systems for flight in fully integrated airspace, precluding manufacturers and operators from developing and utilizing UAS for many operations.

C. Laws Applicable to UAS Operations

The United States operates under common law [22]. This means that the body of law arises from precedents having been derived from judicial decisions of courts and/or similar entities [22]. As such, US laws always run behind technology, waiting for a situation in which they will be interpreted to set the precedent for how the laws apply to that technology. Note that laws do not chase the technology; they chase the underlying issues which apply to the technology. Those issues and laws are interpreted with respect to the specific technology. As mentioned previously, *Boggs vs. Meredith* in 2015 set the precedent for establishing that the FAA has jurisdiction down to the blades of grass if exposed to the outside. Beyond that legal precedent, much of the law concerning UAS falls under nuisance laws such as Peeping Tom laws. For example, video requires consent from both parties — the party recording the video and the party being videoed. Since most UAS carry onboard cameras, UAS operators have to be careful of the field of view of the camera since video footage captured by those cameras are subject to two-party consent.

Since major incidents with UAS have not yet set a precedent, discussions beyond the precedents described above are based on hypothetical situations. Given a hypothetical situation in which a UAS causes the crash of a passenger jet by destroying engines, causing loss of life, the general consensus of the experts at the safety symposium was that everyone gets sued in a civil case — the airlines, the manufacturers, the UAS operator, etc. The burden of proximate cause is placed on the courts to figure out. In reality, this situation likely means that the defense is tendered to the insurance companies, who likely settle. The effect of the insurance will be discussed in Section II.D. The most difficult part of this process will be finding enough parts of the UAS to identify the operator. Required UAS ID and tracking — currently in the process of becoming regulations — will significantly simplify this process. However, this does assume all systems are functional, the operator is attempting to fly legally, and position updates between UAS in near vicinity of each other are sufficiently unambiguous along with recovered parts to conclusively point to the system at fault. In many cases, such as UAS crashing when flying near crowds or UAS operating in the way of emergency systems, unsafe operations are often traced through social media postings. Even when proof is found, proper chain of custody procedures must be followed. Often, UAS violations are called in to local law enforcement as first responders. However, since local law enforcement does not have jurisdiction over the airspace, this path is a dead end, usually resulting in the UAS and operator being long gone before law enforcement officials arrive. In summary, likely most UAS violations — other than clear criminal cases — will result in lawsuits that must be allocated by the court system and handled through insurance, assuming that sufficient evidence is available to track down the at-fault parties and the correct authorities are involved.

As can be clearly seen in the preceding paragraphs, there is still much ambiguity in the realm of UAS law and the

effects of lawsuits. As such, many operators, especially commercial operators, have resorted to obtaining insurance to handle cases in which lawsuits are the primary recourse by offended parties.

D. UAS Operation Insurance

UAS insurance is a legal issue since insurability brings in a host of concerns including violation of FAA rules, physical damage and bodily injury, nuisance laws, trespass laws, invasion of privacy, stalking and harassment, and wiretap laws. Of those, trespass laws and invasion of privacy laws are two of the major reasons for UAS violations. While there are multiple insurance policies that can be used to cover UAS operations, that is beyond the scope of this paper. Refer to [23] for more details concerning methods to insure UAS. Rather, of interest is high level breakdowns in methods of insurance and how they map to UAS operations. Insurance tends to be evaluated in one of two ways:

- It falls into a “normal” bucket. In this case, it is often passed along to re-insurers who insure/price it based on standard rates. Typically, this is fully automated/computerized with human oversight.
- Some parameters of the insurance request are outside the norm. In this case, it is typically evaluated by a human in the primary insurance companies who helps to determine the risk model and pricing.

Re-insurance, such as Swiss Re, accounts for 65% of recoverables from non-US companies. While often less well-known than primary insurers, re-insurance is a method of spreading risk — a way for primary insurers to insure policies with well-understood risks. To create these buckets of insurance policies that can be passed on to re-insurers, some insurers use exclusions to keep operators within the bounds they specify. However, this method often leads to unintended consequences. Global Aerospace — a primary UAS insurer — removed exclusions because crashes invariably break at least one exclusion regardless of the operation, making the insurance useless if the exclusions were in place. Without exclusions, pricing insurance premiums becomes much more dependent on information from the operators. In fact, the current insurance model is supposed to be on a per-flight basis, but it is often not executed this way.

In summary, the insurance industry already has a model of providing insurance which applies to, or at least overlaps with, UAS. The more pressing issue is that UAS risks are not well understood yet, resulting in standard models still adapting to the UAS model. Further, aligning the risk models used by UAS operators and the FAA with the insurance risk models will make the UAS ecosystem integration easier.

E. UAS Operation Safety

Operational risk and, conversely, safety closely tie to the question of insurance. In this section, the outcomes of the 2018 safety symposium will be discussed. Much research has also been performed for technology in this area, which will be discussed in Section III.A. Firstly, there are two classifications of safety that are of interest here: actual safety and perceived safety. Currently, the concept of actual safety appears to be safety from physical harm, applying primarily

to humans. Perceived safety has been shown to vary based on velocity, size, and distance of the UAS, with the primary determinant being velocity. The research into perceived safety has been done by interdisciplinary teams using measures of skin conductivity, head tilt, and heart rate, which seems to equate perceived safety with a physical fight or flight response. Perceived safety violations are the primary driver of reports to law enforcement and complaints to regulators versus actual safety violations. Perceived safety is affected more through education and marketing, such as the team referenced in Appendix IX that used assistance from a product design and art team to make UAS look more “friendly”.

This research is not focused on changes to perceived safety that are effected through visual product design. Rather, this research focuses on technology changes that can be made to enable improvements in both actual and perceived safety through risk analysis of the UAS ecosystem. The safety symposium experts suggested that additional divisions of weight classes for UAS could offer more refined risk categories for the ecosystem. Categories based on weight and/or max speed could be appropriate since weight and speed directly correlate to kinetic energy, which has a significant impact on whether the UAS can harm humans or property [10] [24] [25]. While these categories may have less effect on a numerical risk analysis, the experts also suggested additional segmentation for operators based on vehicle classes, beyond Part 107 [14] training, which could affect numerical analysis more strongly since operator capability to handle an emergency could be taken into account. The question raised here is whether different UAS (size and type) have vastly different flight and operational characteristics, or whether the mission and flight control software removes much of the differences in dynamics from the operator.

Regardless of how the risk and safety are determined, the panel believed that commercial use will drive regulations. Losses will come first, which will drive public opinion, in turn driving new regulations. In the meantime, the industry is in a wait-and-see mode. Additionally, there are many redundancy and reliability issues that must be addressed for vehicles and software. Regulation at the manufacturer level may be required to handle safety issues. Unfortunately, software reliability drives cost. Thus, increased reliability requirements may drive costs too high for many UAS manufacturers. An alternate view is that product liability will drive increases in safety and decreases in risk. Consider, for example, operations only being authorized if the risk (determined for each operation) is sufficiently low. Further, consider that even if the operation is authorized, the operators must transfer the risk (through insurance) to adequately protect themselves. Now, the risk is determined both by the authorizing agency and the insurance company. Too high of risk means no authorization to operate. Likewise, even with authorization, too high of risk means the risk transference cost is too high. Consequently, operators will require more reliable systems or alternate methods of operation which lower the risk, thereby driving a reduction in risk without regulation at the manufacturer level. By driving risk down through this method, manufacturers who are unable to meet software reliability requirements for lower risk may be able to decrease risk via other means, perhaps through operational requirements for greater separation from humans and property. Likely, a combination of the above methods will lead to sufficiently reduced risk to enable operations.

F. Quantitative Framework Characteristics

Based on the results from the safety symposium, three primary characteristics were derived for a quantitative risk-analysis framework for UAS decision-making:

- **Explainable:** the risk analysis framework will be used by both regulatory agencies for authorizing operations and by insurers for insuring operations. As such, any quantitative results from a risk analysis framework must be explainable to humans overseeing the operational decisions. This characteristic rules out purely data-driven methods such as neural networks [26] which are flexible but difficult to explain.
- **Flexible:** as noted in Section II.A, there is significant unknown with regards to UAS. Coupled with the fact that new UAS are being manufactured and introduced to the ecosystem on a regular basis, the risk analysis framework must be flexible to handle an ever-changing analysis. This characteristic rules out more cognitive systems such as expert systems and case-based reasoning [27] that are easy to explain but inflexible.
- **Initialize from unknown:** Unlike manned aviation, UAS introduced to the ecosystem are typically flown before risks are quantitatively established — i.e. there are no priors for initializing quantitative frameworks. As such, a framework that can incorporate these UAS quickly must handle lack of *a-priori* information.

III. Literature Review

This section details current research based on the motivation provided by the Safety Symposium results from Section II. Both research into risk analysis frameworks and research/technology that would feed data into a quantitative framework are discussed.

A. Current Risk Frameworks and Technology

For some specific operations, such as ones that are already authorized under current UAS rules, some current UAS technologies are sufficiently provable for authorization under proposed risk assessment frameworks. Others, however, do not provide the assurances required to operate in more complex and higher risk environments. In manned aviation, there is always the fall-back position to the human pilot and the airframe. The airframe, which includes critical vehicle systems such as required flight controls, is proven to reliability standards. If all non-critical functions such as sensor failure, the human pilot is believed to be able to recognize the failure and replace the failed system with their own capabilities, thereby providing the ultimate backup system. In helicopters, human-controlled auto-rotation is available. Manned quadrotors are not typically flown, and multirotors with more rotors (such as 18 rotor systems which have been recently demonstrated [28]) can lose at least one motor and still be brought safely to the ground.

For unmanned systems, the technology systems become more crucial since a ground-based pilot might rely on sensors to provide critical flight feedback, the system may rely on decision-making systems to make appropriate recommendations or decision, or, in many cases, the system may rely on flight control systems to control the vehicle even

if the operator takes as direct control of the vehicle as possible. For most small UAS, the development of these flight control and sensor systems are not regulated or verified nearly as rigorously as for manned aviation. The question then arises: will UAS be cost effective while proving that they will remain within safety constraints? For some operations, the answer is “yes” since the revenue of those operations outweighs the development and operating costs. For other operations, the answer is likely “no”.

A second approach taken by several innovative systems such as Xavion [29] is to create a highly reliable backup system that takes over in the event of failure of the primary systems. In this case, a single, well-developed backup system could be cost effective by spreading the cost across the myriads of UAS which would use it. However, this method too has its drawbacks. Either the system must be sufficiently self-contained such that all sensors, controls, actuators, etc. are highly proven and are carried as a second system onboard the vehicle (which is untenable for small UAS), or certain parts of the original system must be trusted. Herein lies the deeper issue: either the original system must accurately report when it is failing, or the backup system must have sufficient technology to detect and handle the failure.

Current research provides means to detect and handle certain failures. In particular, fault detection systems [8] [30] provide detection of failures with some degree of reliability. Further, risk assessment systems [10] provide means to predict the results of failing subsystems and provide the operator with an assessment of the implications. In both cases, some level of *a-priori* knowledge is required. The reliability of the fault detection systems must be characterized in order to feed the priors of the Bayes reasoning-based risk assessment systems [10]. Therein lies the issue: there is insufficient flight time and testing performed on these systems to provide accurate priors, leading groups such as JARUS to fall back to qualitative assessment methodologies in SORA [3]. Subsequent sections discuss research into risk analysis frameworks and underlying technology used to feed the data required for the risk analysis frameworks to function.

1. Risk Framework Research

Much research has already been performed on risk frameworks, with various papers focusing on individual operations to overall frameworks designed to incorporate technical risks analysis in various aspects of UAS operations. Several papers based on research performed by Roland E. Weibel [31] [32] [33] are focused on risk analysis and mitigated methods for integrating UAS into the US National Airspace System (NAS). While these papers deal with the topic of risk assessment and mitigation for UAS, these papers primarily form what could be seen as the basis of the SORA concept since these papers preceded SORA. These papers do incorporate a different focus in that they are interested in US NAS while SORA is focused on the European airspace [3]. However, their primary concern is still a holistic risk model. As such, grounding risk analysis research in this type of framework applies to US and European airspaces, making this a solid foundation for research.

More recent work develops risk assessment tools designed to be trusted as the basis for FAA assessment and operator evaluation [34]. This work focuses on the numerical risk assessment of air and ground collisions, using historical data

from manned aircraft to validate the model. Further, this work relates the risk to costs to insure the operation in order to make the results more meaningful. Similarly, recent work in characterizing the consequences of UAS collisions in the NAS [35] provide an improved understanding of the risks, paving the way for integration into a comprehensive risk analysis framework. Other recent risk assessment frameworks also seek to fulfill the FAA's need for a consistent measure of risk assessment [36] that takes into account the effect of UAS collisions with the ground.

Other recent works seek to evaluate the operational risk of particular operations, such as urban cargo delivery by small UAS [37], which focuses on the entire risk assessment for that class of operations. "A ConOps derived UAS safety risk model" develops a risk assessment model derived from concept of operations (ConOps) that uses Bayesian Belief Networks (BBN), causal narrative, and Huggins engine to determine rolled-up probability of failure for specific scenarios [38]. Likewise, "UAS (Unmanned Aerial System) Safety Analysis Model (USAM)" uses a BBN for developing a data-driven, integrated safety analysis model [39]. Similar work uses expert opinions to derive fault trees and associated risk analysis [40]. Combined with risk analysis frameworks discussed previously, there is fairly extensive literature for developing a risk framework and quantitative analysis for UAS. The research in this paper is not intended to replace the risk analysis frameworks developed previously, since using a framework such as SORA [3] or those developed by Weibel, et al. provides an adequate basis for grounding this research. Instead, this research focuses on numerical frameworks that can provide a basis for generically calculating the risks, which deals with the shortcomings of the technologies underpinning numerical frameworks already in use.

2. *SORA*

Since the SORA [3] framework or a similar framework is used as a basis for grounding this research, a short background of SORA is provided. More details of the SORA framework along with gap analysis can be found in [41]. SORA is derived from the JARUS guidelines document [3]. SORA [3] proposes a methodology to assess risk required to support an application for authorization to operate a UAS within the specific category. In many cases, UAS operators only desire or need to operate the UAS within a limited or restricted manner. In such cases, full design approval, airworthiness certification, type certificate, and vehicle certification consistent with a pilot's license (assuming the vehicle systems make decisions while flying) is unnecessary and restricts the ability of the UAS to perform the desired tasks. The SORA methodology is based on a bottom-up, total system safety risk assessment model that evaluates risks for a specific operation. This analysis includes all threats for the operation, the relevant design, and any mitigations to determine the boundaries for safe operation. The definition of risk used is the combination of the frequency of an occurrence and its associated level of severity. The consequence of each occurrence is a harm. While there can be many levels of harms, multiple studies have shown that the energy associated with a crash is consistently well above the low energy levels required for a human fatality [3]. Further, human fatalities are well-defined and, in most countries, well-known by authorities. Therefore, under SORA, only human fatalities are considered as harms due to ground

collisions or catastrophic mid-air collisions. As a result, the best measure of quantifying risk is the number of deaths in a given time interval or per special circumstance (such as per take-offs). The SORA bow-tie model in Figure 1 shows the flow from threats to harms and options for mitigations to reduce the risk.

Fig. 1 [3] The bow-tie model shows the flow from threat to harm in any specific operation. Using this model, mitigation can be employed at each step to reduce the risk that a specific threat will result in a specific harm.

The three categories of harm defined in SORA are fatal injuries to third parties on the ground, fatal injuries to third parties in the air, and damage to critical infrastructure. It is the operator’s responsibility to ensure that no other categories of harm arise in their specific operation. Degradation of mission is not considered a harm in SORA. The only hazard specifically defined in SORA is “UAS operation out of control” [3]. Out of control means that the operation being conducted is beyond of the approved operation, which is significantly wider than simply loss of control of the UAS. For example, the UAS entering an unauthorized airspace, even when under full control by the operator and within visual line of sight, would be considered out of control since that airspace was not authorized for the specific operation and could lead to mid-air collision harm.

The five categories of threats identified by SORA are the following:

- Technical issue with the UAS
- Human error
- Aircraft on collision course
- Adverse operating conditions
- Deterioration of external systems supporting the UAS operation

SORA presents a framework for systematically analyzing each of these threat categories, determining the appropriate threats in each for the given specific operation, and determining the paths from those threats to the specific hazard and the three specific harms defined in SORA. These paths can be analyzed through quantitative probabilities, which are controlled by the limits in Figure 1, derived from an equivalence to manned aviation.

The inherent difficulties associated with quantitative analysis of probabilities in complex systems, however, limit the

Table 1 [3] The probability limits of fatalities per flight hour for SORA are equal to the limits for manned aviation, since use of those limits provides a consistent basis of acceptable risk levels in aviation. These limits are further broken down to enable limiting risks depending on the category of operations.

	Number of fatal injuries to third parties on ground per flight hour	Number of hazards per flight hour	Number of persons struck per flight hour	Probability that person suffers a fatal injury
Certified Category	1^{-6}	1^{-6} to 1^{-4}	1^{-2} to > 1	1
Specific Category	1^{-6}	1^{-6} to 1	1^{-5} to > 1	0.01 to 1
Open Category	1^{-6}	1^{-2} to 1	1^{-5} to 1^{-2}	0 (harmless) or 0.01 to 1

utility of a numerical analysis. Completeness uncertainties, due to inadequacies of the model, modeling uncertainties due to lack of knowledge of how to model complex phenomena, and parameter value uncertainties due to lack of test data to provide those values are just three of the difficulties with quantitative analyses that limit the utility of the analysis results. Therefore, SORA uses a qualitative analysis with a range of numerical levels that include risks, mitigations, and their levels of robustness. This analysis begins with an initial risk assessment. The goal through the SORA process is to ensure that the risk is commensurate with the proposed Concept of Operations (ConOps) through strategic and tactical risk mitigation/reduction strategies with associated robustness requirements [3].

3. Separation Standard

In the paper “Establishing a Risk-Based Separation Standard for Unmanned Aircraft Self Separation” [42], Roland Weibel, et al. propose quantifying the “well-clear” aircraft separation standard based on a time to closest point of approach analysis. Further, per the safety symposium notes in Appendix II.A, the University of Central Florida (UCF) estimated the well-clear distance for UAS and helicopters, which went into the FAA guidance and has been corroborated via other research [43]. Quantifying standards enables further quantifiable risk analysis, enabling a broader numerical risk analysis. The limitation, however, is that sufficient knowledge of all other air traffic that could collide with the UAS is required. Since the FAA regulations that have been proposed in 2020 [7] will require transponders on all UAS, this information will be available, although datalink reliability, update rates, and other factors will still need to be handled in the risk analysis.

4. Failure Impact Analysis and Crowd Modeling

Analyzing the risk of UAS operating over populations is an essential component of integrating UAS into NAS and has been the subject of various research endeavors [36] [44] [45] [46] [47]. Several recent research endeavors have focused on predicting the impact of a failure. For example, analysis of population centers [25] [48] [49] is used to

predict the likelihood of loss of life due to a failure and resulting ground collision. Decisions based on risk analysis are comprised of two primary components: likelihood of failure and impact of failure. Thus, this analysis is essential to risk analysis and decision-making for UAS. Failure modeling consists of understanding the crowd or population dynamics and the vehicle dynamics during impact.

Current population modeling uses estimates based on number of people living in population centers and various models for macro-movement versus individual movement versus hybrids [50] [51]. Recommendations are made for incorporating cell phone statistics into the measures for a more real-time estimate. Since multi-lateration (triangulation of a transmitter position through measurements obtained via multiple receivers) is already used extensively in the cell phone industry and has been researched for developing population maps [52], centralized, anonymous statistics could be used to evaluate the safe zones for landing in the event of degradation of systems essential to a UAS [53]. While this concept works for populated areas in which there is sufficient cell phone use to anonymize the statistics, regulations prevent using this data in more rural areas in which the likelihood of identifying the cell users is too high. An alternative in rural areas is to plan flight paths via zip codes, using population statistics for the zip codes. A third method for capturing crowds in real-time is to use geo-located social media feeds. For example, Twitter geo-located tweets can be used for analysis of crowd formation and movement [54]. Using video feeds from UAS, research has also shown that neural networks can be used to improve the crowd modeling, thereby supplementing data sources with direct observations [55].

While the previous research focused on population and crowd modeling while including basic dynamics of the vehicle upon impact, other research has been done to model the dynamics of the vehicle including bounce effects to better understand the potentially affected area [56].

5. Path Planning

A significant portion of handling off-nominal conditions is path planning and estimation, both before and after events occur. Path planning before the occurrence of an event that causes off-nominal conditions has the ability to provide a safety net of options in the case of such an event. This type of path planning, based on nominal vehicle dynamics, is well researched, with further ongoing investigations. These planners typically perform optimizations with parameters that allow the designer to choose the balance between safety (related to particular situations) and mission efficiency. A good example of this type of path planning is detailed in the work by Vian, et al. [57]. Once an event has occurred that causes off-nominal conditions, the flight trajectory prediction is based on the new vehicle dynamics. Since the off-nominal vehicle dynamics are not known *a-priori*, they must be estimated either based on known failures or through an online learning algorithm. For example, the neural network capabilities in the GUST software [58] enable learning the new vehicle dynamics online, allowing path planning and estimation to be based on the updated model of the vehicle dynamics [59] [60] [61].

6. Fault Detection, Identification, and Accommodation (FDIA)

For small UAS, fault detection is often difficult as there are not always backup or monitoring systems available to detect or handle subsystem failures. Redundancies and detection for many aspects of small UAS, including IMU, can be provided if designed properly [62]. In “Fault Detection, Identification and Accommodation Techniques for Unmanned Airborne Vehicle”, Lennon R. Cork, et al. [8] proposed using a neural network to detect abnormal operation of the IMU and to replace the IMU data with output from the neural network should such a detection occur. While neural networks are limited in application for flight approval since it is difficult to prove that neural networks have repeatable responses, this concept demonstrates that even the most central of sensors for multirotor stabilization can be provided with both a backup and observer. These methods provide an information input to a risk analysis model — both as a safety alternative and the probability of the alternative fulfilling its role should the primary system fail.

Fixed wing aircraft provide easier backup systems in the event of motor failure since they can glide safely to a landing, should a landing zone be available. Options provided by Pedro Fernando Almeida Di Donato enable potentially safe landing options for UAS [9]. Further, autonomous auto-rotation for rotary wing aircraft, such as detailed by Thanan Yomchinda, et al. enables safe descent methodologies for the third major class of UAS (rotary wing versus plane or multirotor) [63]. These technologies enable online risk analysis with backup options. However, there still remains the risk associated with whether the FDIA system correctly handles these faults, and more importantly for risk analysis, whether that risk is quantifiable. This question goes back to the heart of system design and development with a focus on the approval process. If FDIA systems cannot adequately quantify their ability to detect and handle faults, then the risk analysis must ultimately still be handled by an *a-priori* qualitative risk analysis, which limits the ability of these systems to make decisions in real-time.

A recent development regarding FDIA is the use of a backtracking algorithm to determine the critical failures that can cause a specific degradation or failure to occur [64]. In addition to FDIA pathways, this method also provides risk assessment of the critical failures and probability of the downstream failure to occur, enabling this technology to be easily included in risk assessment frameworks that use probabilities as the basis of the risk assessment.

7. Online Risk Analysis

There is another option for a backup system to handle contingencies: predict the effect of off-nominal behavior on the trajectory of the UAS. There is a subtle difference between this option (online risk analysis) and the previous option (FDIA). In the previous option, the backup system must detect a failure (fault identification) and handle it (fault handling) in order to return the system to a near-nominal condition. Instead, online risk analysis stochastically predicts the new trajectory of the UAS [10], enabling either an operator or an intelligent system to respond appropriately. This option can only be performed online since the cause of an off-nominal trajectory is not known until it occurs. Online risk analysis provides several advantages, but the primary one is the ability to delay decisions until they become necessary.

For example, a system without any online analysis must restrict its flight to ensure all contingencies do not cause out of control operations. Online risk analysis provides the means to allow a broader flight envelope until the risk of out of control operations is too great. Flight trajectory predictions provide the operator with a means to steer the trajectory to minimize the impact of the off-nominal conditions. Further, these evaluations do not need to occur in real-time. Close to real-time is sufficient, provided that the time delay of computations is incorporated into the trajectory predictions.

A couple papers have recently been published which begin to realize this option. “Real-time Risk Assessment Framework for Unmanned Aircraft System (UAS) Traffic Management (UTM)” by Ancel, et al. [10] provides the first real-time risk assessment implementation in literature for a UAS. This is a significant achievement towards the realization of real-time risk analysis and adaptation for UAS, but there are still many gaps to be addressed. The primary gaps in this research are the use of offline resources for computations, which become unusable in the event of loss of communications, and the restriction to Bayesian Belief Networks, which typically require a reasonable set of *a-priori* data to provide the priors. Likewise, recent developments have been made in online guidance updates based on collision risk assessment [65]. This research uses decision trees and probability of collision to determine whether to maneuver and, if so, which policy choices to follow based on offline learning.

Analysis of air vehicle avoidance volumes based on navigation errors has also been performed [11]. Specifically, this approach both provides the avoidance volumes for a given system and allows the computations of the necessary avionics to achieve a desired avoidance volume. This work, together with the former risk assessment framework, provides a mechanism for real-time evaluation of a vehicle’s trajectory and the areas it might impact. Further, combining these analyses with the failure impact analysis in Section III.A.4 provides a complete means to evaluate the air or ground collision risk of a UAS, albeit with some shortcomings.

8. Dempster-Shafer Theory

General, quantitative risk analysis frameworks require an underlying technology to calculate the risk, usually in the form of probabilities. The frameworks discussed in Section III.A.1 use various methods; Bayesian Belief Networks is used several times as a structured data method that allows reasoning over stochastic variables. A less well-known theory which applies to this problem is Dempster-Shafer Theory, which many view as a generalization of Bayesian reasoning [13]. However, unlike Bayesian reasoning, Dempster-Shafer Theory does not require knowledge of priors, meeting the third characteristic defined in Section II.F. Furthermore, Dempster-Shafer is based on the concept of probabilities over sets of sets, including the complete set, which is equivalent to unknown information. Thus, Dempster-Shafer explicitly defines and calculates what portion of the probability distribution is unknown, allowing the theory to start from fully unknown data (i.e. no priors) as well as make different decisions when faced with lack of information versus balanced probabilities between various options. Primarily used in sensor fusion and risk analysis [66] and typically requiring higher computational resources, Dempster-Shafer Theory provides useful properties that overcome some of the issues

facing the risk calculation technologies underpinning quantitative, general risk analysis frameworks discussed in Section III.A.1. Extensions to DS networks were developed as a prerequisite to this work [67], enabling the UAS risk analysis performed in Section IV.

IV. Scenario and Test Setup

This section focuses on applying Dempster-Shafer networks with auto-updating joint conditional probability matrices [67] to a UAS scenario — a hovering multirotor making real-time decisions on whether to land, and, if so, in which area to land. The scenario is shown in Figure 2.

Fig. 2 The layout for the UAS news multirotor scenario. An area of operations, which limits the risk of the UAS flight to lives not involved in the operation, is defined and shown. The goal of the news multirotor is to maintain the best visual coverage of the area of interest while maintaining an ability to land if issues arise, in order to keep the operation risk manageable. For this operation, two safe zones are specified as areas in which the multirotor could land without risk to lives. Additionally, the scenario assumes that some monitoring method for these safe zones are available, which could be as simple as an operator actively monitoring the zones and notifying the UAS if the zones are becoming unsafe for landing. Since the goal of the UAS is to maintain visual coverage of the area of interest, the multirotor hovers over the primary safe zone, but will move to the secondary safe zone if the primary zone is compromised. Additionally, the UAS will land if the risk becomes too high. This scenario encompasses many facets of UAS risk analysis including a mechanism to assess risk, multiple options/hypotheses, and decision criteria associated with the risk analysis.

Previous research in real-time health assessment of UAS [68] developed a system that provides good, warning, and failure indications of various UAS subsystems, which results in a comprehensive health diagnostic for the UAS. Three known limitations of this system are as follows:

- All subsystems must report a common basis of health diagnostic (e.g. good, warning, or failure).
- The comprehensive health diagnostic selects the worst case subsystem health diagnostic, which ensures a worst-case response only and does not provide a full understanding of the risks faced by the system.
- Configurable delays help to handle conditions causing rapid changes between instantaneous health diagnostic states (e.g. switching rapidly between good and warning), but these delays deal with this problem by introducing hysteresis mechanisms without capturing the underlying distribution between the states.

Integrating this real-time health analysis system with a DS risk assessment overcomes these limitations by translating reported health diagnostics to a common basis through joint conditional probability matrices, providing a full distribution including unknown and ambiguous components of the risk assessment. Further, the DS analysis easily incorporates additional information external to the UAS. Figure 3 shows a DS network that can be used for UAS real-time risk analysis.

Fig. 3 Dempster-Shafer network for UAS risk analysis. Each node includes three θ_s or individual options being evaluated: low risk, medium risk, and high risk. This network is more appropriate for a small, lower-cost UAS that will not respond differently to risks in each internal subsystem. Power and flight systems are still separated since power system warnings and failures are more common issues for multirotors and have pre-planned responses. Likewise, the operator is a separate node since the capabilities of the operator (whether Part 107 [14] certified, etc.) play a strong role in the overall operation risk. Multiple risks can be assessed for the environment including weather, terrain, crowds, etc. Assuming that a Part 107 operator is correctly following rules and flying in appropriate weather for the UAS, the environment risk analysis is simplified to focus on safe zones, which are known, monitored landing zones for the UAS. This network includes multiple hypotheses for the safe zones, which are not depicted in this figure.

A. Decision Criteria Design

The UAS DS network includes two decisions. The first is the operation risk decision that determines whether the UAS should land or continue the operation. The second determines which safety zone should be used by the multi-rotor. If both safety zone hypotheses have the same result when compared to the decision criteria, then the primary hypothesis is used. Otherwise, the better hypothesis is chosen. This criteria is set more restrictively than the operation risk criteria such that the vehicle can pre-position itself before making the decision whether to continue the mission. Table 2 shows the decision criteria values based on the decision types defined in Figure 4.

B. Evaluation Metrics

Evaluation of this real-time risk analysis method was performed in comparison against the baseline real-time health assessment architecture implementation in GUST [58] [68]. For this comparison, the following metrics were chosen:

- Response time: the time in seconds between the first notification of a degradation in the system (e.g. an instantaneous health status changes from good to warning) and when the system takes contingency action in

Fig. 4 Decision criteria defined for Dempster-Shafer analysis. **Precise Understanding** means that the decision-maker requires little-to-no unknown. **Flexibility** means that the decision-maker requires significant ambiguity or unknown in the system. **Low Known, Unknown Risk** means that the decision-maker requires both the probability of a particular risk as well as the maximum possibility of that risk to be low. **High Possibilities, Belief** means that the decision-maker requires that there is a strong belief that the value under consideration is true, and the possibility of that value being true is very high; this case could be applied to the stock market. **Low Possibility of Risk** is used when the decision-maker is only concerned with the maximum possibility of a risk; this limit could be applied when the impact of the risk being realized is too high to accept, so the possibility must be minimized. **High Belief** is used when only the probability of the situation is important to the decision-maker; this choice is a typical Bayesian approach. **High Possibilities** is used when only a the possibility of a situation is of interest to the decision-maker; this criteria could likely be used in gambling situations. Finally, **Low Belief** is used when only the probability of the situation is important to the decision-maker; this limit is, again, a typical Bayesian approach.

response to the change

- False positives: the number of times the system takes contingency action when there is no degradation in the system
- False negatives: the number of times the system does not take contingency action when there is a degradation of the system
- Processability: whether the system is capable of processing the full autopilot update including the DS network in the 100Hz update loop

For each system (the baseline and the DS network — i.e. the system under test), trade-offs between the first three metrics are handled through configuration parameters. These parameters were chosen to provide a similar balance between each of the first three metrics before comparisons were made between the systems. The final metric is a requirement for running on small UAS, which is previously demonstrated on the baseline system [68] and must be evaluated on the test system. Since the baseline system did not include the concept of safe zones and multiple hypotheses,

Table 2 The numerical decision criteria for the UAS Dempster-Shafer network. For real-time flight system implementation, the decision criteria is reversed such that when the decision criteria is met, the system executes the contingency action. As such, these criteria are consistent with Figure 4. The Safety Zone Risk criteria are more restrictive than the operational risk, enabling the hypotheses and the chosen safety zone to be switched before deciding to end the mission, thus providing an opportunity to continue the mission with a lower risk safety zone.

Criteria Set	Limit Application	Low	Medium	High
Operation Risk	Belief	≤ 0.5		
	Plausibility	≤ 0.9	≥ 0.4	≥ 0.3
Safety Zone Risk	Belief	≤ 0.6		
	Plausibility	≤ 0.9	≥ 0.3	≥ 0.2

that capability is evaluated separately by testing whether the UAS can choose the less risky safe zone (relative to the decision criteria defined for safe zone choice) and delay aborting the mission by lowering the risk through the safe zone choice. Further, the baseline system does not provide a way to include health or risk information external to the system, such as operator risk. This capability will also be tested separately by evaluating whether the UAS chooses the defined contingency action when risks from external systems become too high or whether the UAS continues the mission with no change.

C. Network Training

Training for the DS network was performed offline via simulation. While online training is feasible through this system, current methodology [67] uses the results of each operation to learn the transition potentials — the relationships — of the network. Specifically, this method connects the long-term and real-time in-operation timelines through the transition potentials. Since DS combinations are used to combine the evidence per episode for the transition potentials, the captured relationships can change over time as the evidence suggests shifts. Offline training was chosen for this application due to the number of operations that were necessary to capture reasonable relationships for testing against the baseline system. To clarify, evaluating only a few test cases results in significantly higher levels of uncertainty mapping between nodes, which does not provide a similar response to the baseline system. Further, the approximate relationships were known *a-priori*, enabling rapid confirmation or disconfirmation through offline training.

Failures were modeled as Poisson Point Processes [69], which assumes stochastic independence between each failure occurrence. More accurate failure models are likely available on a manufacturer-by-manufacturer basis such as by Velos Rotors [70], which provides a maintenance schedule for their UAS parts. The simulation used for this research includes a replaceable model, and the Poisson Point Process was chosen since the flight demonstration platform, described in Section VI, is experimental with too few flight hours of data to provided a more accurate model. Offline training was performed by setting the probability of failure and warning (failsafe) per hour for each of the nodes without parents (flight systems, power systems, operator, and safety zone). The probabilities of failure and warning were set to zero

and 50% each during training cases. While 50% failure probability is high, it gives a view into the relationships when components are at high risk of failure. Simulation lengths were randomly chosen between the minimum of 0.1 hrs and maximum of 30 minutes, assuming a commercial multirotor such as a DJI Mavic Pro [2]. 250 operations were run per learning case.

Conversion functions from the subsystems’ instantaneous reports of good, warning, or failure to evidence input for the network has a substantial impact on the training and how quickly the transition potentials converge. Table 3 shows the chosen mapping from result to evidence input. Higher ambiguities and unknown results in slower, more consistent convergence. Moreover, the ambiguities, while necessary for smoothly combining the evidence, also have real-world meaning. A successful operation does not mean that the risk was necessarily low. It merely means that the risk did not manifest itself into a failure or failsafe. Therefore, the (success, failsafe) ambiguities and the (success, failsafe, failure) unknown mass capture this possibility. Higher ambiguities and unknown masses highlight greater uncertainty between the operation results and the risks in the operation, allowing more operations to reduce the uncertainty through evidence combination.

Table 3 The mapping from results to maximum risk evidence inputs for each element of the power set. Actual evidence inputs for each operation are randomly selected up to the maximum values in the table, with any leftover mass being assigned to the complete set (unknown) to ensure the evidence masses always add up to 1.0. This partially random input simulates evidence from risk assessments performed after the flight operations for each specific operation. In practice, the high number of runs along with Murphy’s combination rule [71] for averaging inputs results in a similar outcome to fixed values. The operation risk has a higher degree of uncertainty per outcome to simulate the uncertainty associated with the risk of the overall operation. Individual subsystem risks are more precise, simulating more detailed fault analysis on the subsystems.

Power Set	Operation Risk			Individual Risks		
	Success	Failsafe	Failure	Success	Failsafe	Failure
Low	0.4	0.0	0.0	0.7	0.0	0.0
Low, Medium	0.3	0.15	0.0	0.3	0.3	0.0
Medium	0.0	0.4	0.0	0.0	0.7	0.0
Medium, High	0.0	0.22	0.3	0.0	0.4	0.3
High	0.0	0.0	0.4	0.0	0.0	0.7
Low, Medium, High	1.0	1.0	1.0	1.0	1.0	1.0

There is the need to capture overlapping training episodes for the transition updates. In Figure 3, it can be seen that all nodes affect the operation risk, creating a situation in which no one node correlates to all of the possible episodes. Figure 5 shows the relationships between downstream nodes and nodes that are directly controlled in training episodes. For example, an episode of training that is running operations with a high likelihood of failure in the flight systems will teach the network about the transitions between the “Flight Systems Risk” node and the “Internal Risk” node and between the “Internal Risk” node and the “Operation Risk” node. However, the transition from the “Operator Risk” node to the “Operation Risk” node is not affected significantly, and including the effects of the episode on that transition

is detrimental to the training of that transition.

A third and final restriction was added to training transitions of multi-parent nodes. In the two previous methods, even if there was only a significant change in the distribution of the “Power Risk” node (i.e. the “Flight Systems Risk” node had no effect on the change in risk of the “Internal Risk” node), the two transitions are still trained together. This is a direct result of the simplification in the transition updated methodology [67], in which the result of the parent nodes multiplied by the transition potentials are assumed to be the same when provided as evidence to the combination method in the child node. This third method still uses that simplification, but then only trains the transition to which the episode applies.

Fig. 5 UAS risk analysis Dempster-Shafer network showing which nodes are primarily affected through changes to the evidence inputs. For example, changes to the Flight System Risk primarily will affect the Internal Risk and Operation Risk, leading to those three nodes being trained as part of the same episodes.

V. Test and Evaluation

Given the network and metrics defined previously for this scenario, multiple tests were created to evaluate the systems relative to the metrics in Section IV.B. 50 trials per test case were then conducted.

- Simple switch: The instantaneous health notification of a subsystem step changes from good to warning without switching back.
- Continuous switch: The instantaneous health notification of a subsystem continually changes on each update between good and warning.
- Stochastic switch: The instantaneous health notification of a subsystem changes between good and warning with some defined probability. A goal of this test is to determine whether there is a point at which the baseline and DS systems clearly differ in their response.

Testing capabilities were built into the GUST flight control to enable both simulation and in-flight tests while minimizing risk to the UAS. Each subsystem in GUST for which there is health analysis provides an instantaneous health status of good, warning, or failure to the health subsystem, which uses those statuses to evaluate the overall status

of the system [68]. A low risk method of testing the health system response while not degrading critical flight control systems is to inject a test system that reports an instantaneous good, warning, or failure status into the health subsystem. Test code was added to compute contingency response time based on the test subsystem injection and to check whether the health subsystem caused a contingency response without a degradation being injected by the test subsystem.

As with the network training in Section IV.C, subsystem instantaneous reports of good, warning, and failure had to be converted into evidence inputs for the DS network. This conversion is simpler than the training case because these are short term evidential views that are equal in weight and constantly updating. Thus, the simple mapping in Table 4 is used.

Table 4 The mapping of instantaneous subsystem health state to evidence input into the Dempster-Shafer network used by the health subsystem. Since each of the instantaneous health states are a form of evidence, a simple mapping is used to add ambiguity and unknown, allowing the Dempster-Shafer combination algorithms to work effectively and providing a slower, less stark response to changes in state.

Power Set	Instantaneous Status		
	Good	Warn	Fail
Low	0.6	0.0	0.0
Low, Medium	0.2	0.2	0.0
Medium	0.0	0.4	0.0
Medium, High	0.0	0.2	0.2
High	0.0	0.0	0.6
Low, Medium, High	0.2	0.2	0.2

Results are grouped to show the trends in the baseline system, the DS network system, and comparisons between the two. All test cases did not exhibit any frame overruns, meaning that both health analysis methods completed all executions within the 100Hz frame update rate in the simulator. The results of the baseline system is shown in Figure 6. The baseline system is a deterministic system with a one second delay before taking contingency action. The system is designed to ignore short, intermittent failures, but captures sustained or significant system degradations. However, as seen from Figure 6, the deterministic system is barely able to capture a highly degraded system. The inverse of the system would behave similarly: if the system is biased towards safety, then only slight degradations would result in the system taking contingency action.

The results from the health subsystem using the DS network are shown in Figure 7. The results act similar to a low pass filter in the sense that the network more slowly catches failures as sufficient percentages of the instantaneous health reports are “warning” or “failure” over a window of time. In this case, the update rate is 0.2 seconds, and the window is 10 sets of evidence, resulting in a two second window for determining failure conditions. For these evaluations, Zhang’s combination rule [72] is used, since this provides nice properties for eliminating outliers and provides a faster response to substantial degradations. With these settings, this method captures failures nearly as quickly as the baseline system for simple failures and can capture most failures down to 25% probability of reporting an instantaneous warning or

Fig. 6 UAS baseline response to subsystem degradation injection into the health subsystem. The system is graded on response time. Numbers in parenthesis below the trials indicate the number of false negatives (uncaptured degradations) in the 50 trials. No false positives were captured. Note that this system is biased away from false positives to avoid safety maneuvers during flight tests. A safety pilot is assumed to be present during flight tests since this is an experimental aircraft. Simple failure is the same as 100% probability of instantaneously reporting failure (i.e. the system simply fails and continually reports a failure). Percentage failures are the probability that the subsystem will instantaneously report failure (i.e. the subsystem is degrading, but not fully failed). Lower percentages than 85% are not shown since no failures were captured at 80% or below.

failure status, which is well beyond what the baseline system can capture.

The results in Figures 6 and 7 are compared directly in Figure 8 for the cases in which both methods can capture failures. While the baseline method responds faster for simple failures, the DS network clearly captures more failure cases and provides consistent responses to those failures. These two systems create a Pareto frontier. However, given that the DS network has a similar response time to the baseline system for simple failures and out-performs the baseline system for other failures, overall the DS network clearly performs better than the baseline system.

As stated in the metrics definition in Section IV.B, further capabilities are provided by the DS network method that are beyond the scope of the baseline system. These capabilities are evaluated separately by determining if the system responds correctly to higher risk in the operator or safety zone analyses. Three test cases were run with 50 test runs per case. The first test case increased the medium risk evidence of the operator. The second case increased the medium risk evidence of the primary safety zone, leaving alone the secondary safety zone. The final test case increased the medium risk evidence of both safety zones together. These risk evidence inputs were pseudo-random. Limits were provided for each of the subsets within the risk evidence, and the evidence was pseudo-randomly chosen up to those limits. Any remaining mass was placed in the unknown set. This method simulates evidence based on observations since these

UAS Degradation Response - Dempster-Shafer Network

Fig. 7 UAS Dempster-Shafer network response to subsystem degradation injection into the health subsystem. The system is graded on response time. Numbers in parenthesis below the trials indicate the number of false negatives (uncaptured degradations) in the 50 trials. No false positives were captured. In order to have comparable results to the baseline system, this health subsystem was also biased away from false positives, meaning that significant deviations from the “good” distribution were required to trigger a contingency action. Two noteworthy points arise from these results: (1) all contingency response times have a distribution — even the simple failure case — since the response is no longer deterministic. (2) this method captures down to 25% failure, albeit with some false negatives and significantly longer times to capture the failure. Moreover, this system gracefully degrades in the sense that the tail of the distribution extends consistently as the failure rate lowers. Test case meanings are the same as in Figure 6. The final test case — switching — is a case in which the subsystem alternates reporting good and failure on every update. This is a case that is impossible for the deterministic baseline to capture, but the Dempster-Shafer network captures this quickly.

inputs were not driven by actual subsystem calculations, unlike the previous tests. As seen in Figure 9, degradations are captured successfully for both cases. The response time to the safety zone degradations is longer than the response to the operator degradation, which aligns with first checking whether the alternate hypothesis is a better choice. The cases in which only the primary safety zone had increased risk are not shown since, in these cases, the health subsystem correctly switched hypotheses and did not take any contingency action.

VI. Demonstration

This research was grounded in the application of risk analysis for unmanned systems. Further, a key assumption was that this research could be applied to the full UAS ecosystem — small systems through large systems. To show that this assumption holds, a flight demonstration was performed using the vehicle in Figure 10 and shown in flight in Figure 11. This vehicle uses a Raspberry Pi 3B+ embedded computer with an Emlid Navio2 [73] autopilot sensor suite running the GUST [58] flight control system. This vehicle was previously used to demonstrate the GUST health monitoring architecture [68], providing a baseline for expanding to the DS network. For this flight demonstration, the

UAS Degradation Response - Dempster-Shafer Network to Baseline Comparison

Fig. 8 UAS health subsystem response comparison between the baseline method and the Dempster-Shafer network method. Only cases in which both methods can capture failures are shown. The baseline system clearly reacts faster for simple failures, but the Dempster-Shafer network method has a consistent, albeit slower, reaction for both the simple failure and lower failure reporting probabilities. The Dempster-Shafer network model clearly captures significantly more failure cases while not slowing the response time substantially.

Operator and Safety Zone Response - Dempster-Shafer Network

Fig. 9 UAS Dempster-Shafer network health subsystem responses to increased operator and safety zone risks. Only two cases are shown — the operator risk increase and the dual safety zone risk increase. All risk increases were captured. The dual safety zone risk increase has a longer response time as the two hypotheses (the dual safety zones) are first considered to determine whether there is an alternate option before deciding to take contingency action. The single safety zone risk increase test is not shown since the UAS never took contingency action in this case. Instead, the UAS chose the secondary safety zone for landing when necessary.

trained network from Section IV.C was loaded onto the UAS, mimicking the testing performed in Section V. Once the UAS was airborne, two evidence injections were used from Section V:

- Higher risk at the primary safety zone, which should cause the UAS to move to the secondary safety zone
- Higher operator risk, which should cause the UAS to land

Fig. 10 UAS research platform used for the flight demonstration. Flight control and onboard computing is provided by a Raspberry Pi 3B+ embedded computer with an Emlid Navio2 [73] autopilot sensor suite. UAS frame size is 400mm. A small platform and basic embedded flight computer was chosen to demonstrate applicability to the full range of UAS sizes, as larger platforms can carry more powerful computers.

Fig. 11 UAS research platform in flight during the flight demonstration. The flight demonstration was kept to a constrained area for personnel safety. All onboard health decisions were performed through the Dempster-Shafer risk analysis network.

The primary question of whether the DS network can run efficiently on the embedded flight control system is answered in Figure 12. Computations including the DS network stay within the time allotted for the 100Hz fixed frame

update rate, demonstrating the DS-based system can run on small embedded computers.

Fig. 12 Analysis of frame overruns for the two flight demonstrations. Frame overruns are defined as each time the computing cycle takes longer than the time allotted in the 100Hz fixed frame update. As seen in the plot, there were zero frame overruns during both flight demonstrations.

The injected risk evidence and contingency responses for the two flight demonstrations are shown in Figures 13 and 14. Recall that the injected evidence is the same as in Section V. As such, only the low risk belief value is shown in the figures, as this provides a surrogate for when the updated evidence is injected. In both demonstrations, the recorded data starts with the UAS hovering over the primary safety zone. When updated evidence for the primary safety zone (SZ1) is injected, the operation risk shows a momentary change in risk analysis, but recovers quickly as the secondary safety zone hypothesis takes over. Concurrently, the UAS exits the hovering state and flies to the secondary safety zone, proceeding to hover there. When the higher risk operator evidence is injected, the overall operation risk increases, triggering the second flight plan, which is an immediate landing. Upon landing, the UAS returns to waiting for the next command. Both flight demonstrations show a consistent response to the inputs, and both flight demonstrations are consistent with the testing results seen in Section V.

VII. Extensions

The real-time decision application in this chapter assumes sufficient evidence through operation results to learn the transitions between nodes in the network. Per the discussion in Chapter II, risk analysis needs to be applied to UAS operations before hundreds or thousands of flight hour results are available. Through the environment discussed in Chapter II, this data is available. Figure 15 provides a more complex risk analysis DS network which can leverage common data across various systems and environments. However, data flows and evidence data retention operate differently in a multi-vehicle context with multiple simultaneous operations compared to a single-vehicle context. Previous data retention methods discussed for training in Section IV.C uses Murphy's rule [71] and weighting rules previously developed [67] to capture episodic data, maintaining information across extended periods. This method

Fig. 13 Flight demonstration one of Dempster-Shafer network risk evaluation with real-time decision-making onboard a small UAS. The reduction in low risk for the primary safety zone (SZ1), which signifies an increase in medium/high risk for that safety zone, results in the UAS deciding to switch safety zones to the secondary safety zone. During this maneuver, the UAS continues the mission since the resulting operation risk is sufficiently low. The reduction in low risk for the operator, which signifies an increase in medium/high risk for the operator, results in the UAS deciding to land since the operation risk is too high to continue the mission.

Fig. 14 Flight demonstration two of Dempster-Shafer network risk evaluation with real-time decision-making onboard a small UAS. The UAS responses in this demonstration are consistent with demonstration one in Figure 13, showing that the overall system is repeatable in its responses.

assumes, across similar episodes, all previous data is captured in a single, consistent network, allowing nodes such as the Operation Risk node to retain information. For example, the combined data retained in the Operation Risk node is still equal to the combined data in all other nodes, such as the Navigation Risk node, multiplied by the transition matrices and combined as evidence into the Operation Risk node. The environment discussed in Chapter II breaks this model. Figure 16 shows one potential model of the UAS environment. Compiled data can be stored in many datastores within this environment. For example, insurance agencies would likely compile data within their risk models of flight

and/or pilot/operator behaviors. Regardless of where the data is stored, data is compiled at different rates as operation hours are built up for UAS, operators, and various environments.

In order to evaluate the risk for a particular operation, the appropriate data must be pulled from each of the datastores, entered into the risk network, such as in Figure 15, and calculated to determine the overall risk of the operation. Decisions are then made regarding risk transference and maximum allowable risk. Once the operation is complete, the data is used to improve the relationships in the network. Herein lies the difference from the single system model previously used. The operation risk is no longer a history of data captured in Murphy's Rule [71]. Instead, it is a combined set of data calculated through the network based on multiple datastores with difference amounts of retained evidence, making Murphy's Rule [71] unusable. In order to get past this issue, the Evidential Reasoning Rule [74] is used, which, like Dempster's Rule [12], does not require the evidence history to add new evidence, but still enables weighting the new evidence based on the hours of operation associated with the new evidence. The evidence combination result will not be retained directly, but will be used to update the appropriate episodes to improve the relationships learned in the Figure 15 network. As such, the results will be retained through the transition relationships and the datastores for retained data such as the various environment datasets, UAS datasets, etc.

Fig. 15 Dempster-Shafer network for the UAS ecosystem risk analysis. This network is similar to the network in Figure 3, but it is more complex to include various features which could be considered common across operations. For example, the risk of hitting the ground (Ground Risk) in a given area of operations is likely to be common across operations in that area and could leverage previous research to estimate the effects of ground impact [10]. Likewise, the same DJI [2] platform models could leverage common data in the Internal Risk node while common autopilot navigation systems could leverage common Navigation Risk information. Decisions for operations are shown in orange as the acceptable risk transference (a question of insurance) and the acceptable risk level (a question for the regulatory agency).

VIII. Conclusion

Regulatory agencies are moving towards risk-based approval frameworks for UAS access to air space [7]. Governments have mandated that these agencies find a way to safely integrate UAS into the air space [16]. How that integration is done

Fig. 16 A model of the UAS environment, including many of the factors that would impact risk evaluation of operations and the relationships among the various systems in the environment. This model incorporates five major systems: the UAS which perform operations, operators which execute operations with the UAS, insurance agencies which insure the UAS operations, a regulatory agency which ensures safe UAS operations, and the environment in which the UAS operate. Each of these systems interact in multiple ways, and the data flows depicted in the model enable the risk analysis, which each of the systems — other than the environment model — perform. The insurance agencies use models of the operator and flight risk and reward to determine operation premiums. The regulatory agency uses models of the operator and flight risk to determine whether the risk is within maximum acceptable risks. Operators use risk and reward models to determine whether they are willing to pay the insurance premiums required to operate. The UAS uses operation risk models to determine real-time risks of various operation profiles to inform the operator or make automatic decisions.

is still a work in progress as the necessary infrastructure, such as remote ID [7], is being developed, per the UAS Safety Symposium notes in Appendix IX. The safety symposium held in August 2018 as part of this research provided a strong basis of understanding for this environment. While risk assessment is currently qualitative, this model is unsustainable as the proliferation of UAS means there will be too many approval requests to evaluate qualitatively. Thus, a quantitative mechanism is required which supports and integrates with a risk analysis framework. Much effort has been put into research for these quantitative mechanisms (Section III.A), yet most suffer from either a limited operational focus or a mechanism that requires too much *a-priori* data, such as with the Bayesian Belief Networks. Dempster-Shafer Theory [12] [13] is commonly applied to risk analysis yet suffers from high computational requirements, often exceeding what is available on small UAS. Valuation networks [75], of which Dempster-Shafer networks are a subset, offer an option which could be computationally feasible for small UAS but require *a-priori* information for the joint mass distributions.

This research pulls together those various threads by applying auto-updating Dempster-Shafer network transition matrices to the UAS risk analysis problem in Section IV. The DS network significantly improved performance over

the baseline system for health monitoring. Contingency response to full failures were similar to the baseline system, and the DS network detected intermittent failures far beyond what the baseline system could detect, with a graceful degradation in response time. Much of this improvement was realized by using a stochastic system, which would only require a single DS node (i.e. not require the update extensions for transitions [67]). The network added capabilities by enabling the health subsystem to include operator and safe zone risk analysis. These values can either be updated in real-time or can use pre-determined values (e.g. a risk level for the operator given current skill level). Further, multiple hypotheses were demonstrated for the safety zone analysis, enabling the UAS to choose the lower risk zone in a manner consistent with the rest of the risk analysis. This analysis was underpinned by the network training based on the previously developed extensions [67], which allows the network to be adapted as new information becomes available. This mechanism is critical for a UAS ecosystem which provides the necessary information to understand the risk relationships. Finally, the full network can run on a flight computer that powers small UAS, as demonstrated in this chapter, thereby applying to the full UAS ecosystem. This application demonstrates the capability to provide a quantitative risk analysis that supports risk analysis frameworks like SORA [3] currently being pursued by regulatory agencies. Further, since the risk analysis network is easily expandable to greater details, the network has the capability to utilize data from algorithms previously developed, in development, and under future development as detailed in Chapter II. This capability provides a new baseline, and Section VII discussed the extensions to the full ecosystem.

In summary, a quantitative risk analysis framework capable of capturing long-term data for understanding risk relationships and applying that to real-time decision making on small to large UAS has been developed and demonstrated through this work. This quantitative risk analysis framework could serve as the numerical mechanization of the risk-based approval frameworks being developed for regulatory agencies to integrate UAS into air space and for insurance agencies to inform their risk models.

A. Recommendations

The following research trajectories are suggested to extend this work:

- Use the UAS ecosystem to train the Dempster-Shafer network, as described in Section VII.
- Integrate current research and algorithms cited in Section III.A into this risk analysis model. This step includes taking work such as the ground collision models that provide impact estimates and integrating these risk models into the quantitative risk model developed in this research to have a fully modeled and data-driven risk analysis.

The addition of these research trajectories would result in a system capable of updating on its own with human supervision, but no human interaction required during the updates. Furthermore, the resulting system would use cutting edge, data-driven models to feed the evidence inputs for this risk analysis network, enabling accurate analysis of the risks, which combine to form the full operation risk analysis. Missing or unproven models have significantly higher unknown masses, resulting in high possibilities (plausibilities in Dempster-Shafer terms) for each of the high, medium,

and low operational risk estimates, thus making operation approval significantly less likely. The beauty of this system is that it gracefully handles refinement of these analysis models while using the same approval criteria. Finally, the UAS ecosystem would enable the risk analysis model to much more quickly incorporate data from similar or identical UAS platforms, operation scenarios, etc., resulting in significantly faster network training.

Appendix

IX. UAS Safety Symposium Raw Notes

Date: 7/23/2018

A. Summary/Major Takeaways

- There is a decided lack of information in this space, but there are several areas and methods of information outreach available by a variety of different enterprises.
 - www.faa.gov/uas
 - www.faa.gov/go/waiver
 - <https://faadronezone.faa.gov> portal
 - 14 CFR Part 1.1 and 107
 - 49 USC 40102
 - FAA Advisory Circular 107-2 Small Unmanned Aircraft Systems
 - Social media platforms (Facebook, Instagram, YouTube, Twitter, and others)
 - AMA
 - UAST (Unmanned Aircraft Safety Team)
 - Drone Advisory Committee
 - Commercial Drone Alliance
 - AUVSI
 - Waiver Safety Educational Guidelines
 - Public Law 112-95 and subsequent reauthorization bills for FAA
 - Pathfinder, Public Safety Partnership, Integrated Pilot Project and other industry-FAA coordinated agreements
 - FAASafetyTeam: FAASafety.gov
 - uashelpdesk
 - FAA AFS-800 Policy Branch
 - Variety of conferences, shows, symposiums, conventions and meetings
- Much of the industry, especially on the regulations and law side, are in a wait-and-see mode for the precedents to

be set, usually when forced by a major incident

- Out of a million registered users, there are 100,000 commercial operators with a sUAS Remote Pilot Certificate registered to operate for compensation or hire, as well as Public Aircraft Operations (government function)
- The biggest issue with enforcing regulations and prosecuting offenders is identification of the UAS and operator: agree, but working on a rule that will provide policy.
- Operator education is essential and is currently severely lacking: agree

B. Notes from Presentation 1 — The Rise of Drones and Insurance

- Reinsurance handles 65% of recoverables from non-US companies – a method of spreading risk
- Farm mutuals rent drones to operators for surveying land, soil, blight, etc. – typically fixed wing craft are used.
- Real Estate is another application
- Legal issues could be faced / how insure UAS operations for those?
- Everything is new with drone – drones, operators, and the law
- Law is always behind tech, but is catching up: My comment at all my presentations is that lawyers and the insurance companies will drive the industry more than the FAA, unless there is blood spilled, in which case Congress will get involved.
- FAA – finalized the current drone rules as of 8/29/16
- True, and have added Remote Pilot certificate testing and recurrent testing, as well as an Advisory Circular, Inspector Guidance, Test Sites, academic coordination with Center of Excellence (ASSURE), and a variety of pilot programs for evaluating operations of drones and risk analysis
- Drones are a federal affair, because the feds are the primary regulators of national airspace
- There have been drone incidents, often involving manned aircraft
- The rules from the FAA lowered the barriers to entry for operators
 - Previous rules: required a pilot's license: a certificate is required for all pilots of all aircraft
 - Current rules: 14 CFR Part 107 requires a "license" and applies to operators flying for compensation or hire. (FAA still does require one for drones: the Remote Pilot Certificate, for those who want to operate for compensation or hire under Part 107. Recreational pilots who do not subscribe to a community based set of guidelines fall under jurisdiction and enforcement of 107, although – no pilot's license ("certificate") is necessary if the drone meets the requirements of:
 - * Must fly for hobby/recreational use only
 - * Operate under a community based set of safety guidelines
 - * $\geq .55\text{lbs}$ and $\leq 55\text{lbs}$

- * Class G airspace only ($\leq 400'$ AGL typically)
- * Maintains Visual Line of sight
- * Avoids manned aircraft
- * Does not endanger the National Airspace
- * Notifies airports and ATC when within 5 miles
- * Per NDAA 2018, must register as operator, using registration number for all owned aircraft over .55lbs and ≤ 55 lbs as recreational user or for each aircraft. If a Remote Pilot operating under 107, then single aircraft per registration number.
- * Request permission to operate in controlled airspace
- * Don't fly impaired
- * Also regulates recreational use of drones per Part 1.1 and 101.
- Under Part 107 as a model aircraft operator:
 - * Maximum airspeed of 100mph
 - * Cannot operate over people or moving vehicles or emergency response
 - * Minimum weather visibility of 3 miles
 - * No carriage of hazardous materials
 - * Preflight inspection required
 - * Daylight hours only unless waived, similar to other conditions requiring waiver
- Most common drone use is renting the drone and operator
- Waivers can be obtained to Part 107: under 14 CFR 107.200 and .205
- Other countries have drone regulations also. Canada and the US are fairly closely synced, but some other countries are more progressive in their regulations
- Drone insurance is a legal issue – Insurability brings in a host of issues:
 - Violation of FAA Rules
 - Physical Damage and Bodily Injury
 - Nuisance laws
 - Trespass laws (One of the primary reasons for drone violations)
 - Invasion of Privacy (The other major reason for drone violations)
 - * Private individuals
 - * Government use/searches
 - Stalking and harassment
 - Wiretap laws
- Boggs vs Meredith – 2015 – where does private property end and public airspace begin?

- With drones, FAA can regulate down to the blades of grass, if exposed to outside.
- If indoors, it's not airspace, and the FAA doesn't regulate it
- ISO drone endorsements
- CGL (Commercial General Liability = A standard insurance policy issued to business organizations to protect them against liability claims for bodily injury (BI) and property damage (PD) arising out of premises, operations, products, and completed operations; and advertising and personal injury (PI) liability.) policy – look to if there's a loss
- Drones as a service – drone highways
- Insurers have reasonable pricing models, but these are dependent on information from the operators.
- Some insurers use exclusions to keep operators within the bounds they specify
- Global Aerospace removed exclusions because crashes invariably break at least one exclusion regardless of the operation.
- The current model is supposed to be on a per flight basis (based on the forms). Often, it is not executed this way
- Insurance tends to be evaluated in one of two ways:
 - It falls into a “normal” bucket. In this case, it is often passed along to the re-insurers who insure/price it based on standard rates. Typically, this is fully automated/computerized with human oversight.
 - Some parameters of the insurance request are outside the norm. In this case, it is typically evaluated by a human in the primary insurance companies who helps to determine the risk model and pricing

C. Notes from Discussion One

1. Group 1

(Bullet points are statements made during a dialogue)

- What is there for regulation and insurance of re-certification?
- Probably if within the scope of work, covered under underwriting questions with input of engineer, probably coverage issued based on underwriting risk audit
- No standards of safety versus known vulnerabilities, but a risk management profile could be developed
- Too new – law hasn't caught up with tech. Drone may need certain characteristics. Geo fencing = no access for drones that are equipped with this feature
- DJI implements altitude restrictions, but can be bypassed
- Authorized first demo at airport – DJI “bricked” because the system shut down. It recognized that it was in restricted space, and there was no way to bypass that restriction. DJI has since updated its software.
- Risk mitigation – there are workarounds
- Business friend in Korea – when too close to the DMZ, the internal GPS will not start

- How can you run self-diagnostics before traveling to an authorized place to fly if the system won't start up outside of the authorized location?
- the definition of "safe" catches up to tech
- safety is piecemeal, developed based on experience and reaction/over reaction
- The UCF team has analyzed data on how people are using UAVs. FAA said not going to fund that work.
- Court case – FAA does not have the right to regulate hobbyists
- Parker – Pre 107, since superseded
- Hobbyists were "refunded" for UAS registration fees, although the money wasn't returned, and 107 was fixed
- FAA on 107 commercial side have authority to regulate, no ability to regulate "hobbyists" guidelines: this isn't true re: hobbyists (model aircraft operators). We regulate through Part 1.1 and Part 101 directly, and through Part 107 if modeler doesn't comply with a community-based set of guidelines
- There are flight restrictions that apply to UAS as well
- Helicopters and UAS operate in similar airspace, which creates potentials for issues
- Regulations are for safety. Pathfinder program – fly drones out of LOS and at night
- The airport had a 2 day safety risk assessment with airlines. Identify risks with operations
- Beyond LOS analysis: higher concern is UAS and helicopters because they occupy the same airspace. UCF analysis estimated the well-clear distance, which went into the FAA guidance. They got a 333 exemption for their flying.
- How can you stay "well-clear" when you don't know what's out there? How do you prove you know what's out there? Definition of "well-clear" is key.
- Assume worst case – they have a few years of helicopter data from Boston, built simulations. There are no flight plans for helicopters.
- Put transponders on UAS?
- It's going that way
- Look at new technology on cars
- Look at the UTM Systems
- Application information on FAA website. Regulation data there also.
- Blocked finding info about the impact of an impact. What happens in a drone strike? Likelihood of damage, shutdown, how shut down. Wind turbulence of wing hit another plane?
- 1st UAV strike this past year was into a helicopter
- NTSB investigated. The drone pilot had no idea the helicopter was there because drone was BVLOS at the time and operator walked away
- White house lawn drone. Operator had no idea where the drone went down – lost contact with it.
- Regulations cannot rely on operators having full control: not sure what this means, but disagree. We want

operators to have control. That is a main focus. Note: Joel Dunham discussed this with the FAA representative after his comments. The point was made with respect to degraded situations in which the UAS may have to take over some level of control and respond appropriately, given that many UAS now have a large disconnect between the operator and the low-level controls that fly the UAS. The FAA representative's response was that the FAA would likely mandate a level of training and understanding by the operator such that even in a degraded situation the operator would know how the UAS would likely respond (such as standardized return-to-home in the future) and be able to predict and plan for the UAS responses. In that way, the operator still has a level of control over what the UAS will do for the operation.

- Strike tests/simulations have been done. There is a fair amount of data. Ask Hartfield-Jackson International Airport representative: how do they enforce the law?
- Depends on who within the Atlanta Police Department answers. They rely more on peeping tom laws. No local ordinances/restrictions. "Master's" got restrictions. Most rely on nuisance laws. Video: consent needed from both people. Example: drone outside a high rise videoing – drone collected as evidence due to video
- Trespass suit? Violate federal law?
- Law cannot chase technology. It chases underlying issues. Cobb county wanted to protect Suntrust Park and sent a letter to the FAA. FAA said "stay in your lane. The air is not your lane".
- Devices interfere with drones
- Video and video collection – when does this violate laws? When is this useful for demonstrating regulation compliance?
- Depends on use. There is an expectation of privacy
- On UCF campus, cannot lift off drone from campus, but can take off from across the street and fly over campus
- The airport is the same. May get waiver, but not through airport screening. How do they enforce hobbyists? How do you control trespass in a restricted area?

2. Group 2: Subject: Safety

- There are two types of "Safety". Actual safety [which seemed to be understood as safety from physical harm] and perceived safety. Perceived safety has been shown to vary depending on the velocity/size/distance of the drone. The primary determinant is velocity.
- The research into perceived safety has been done by interdisciplinary teams and has been determined using measures of changes in skin conductivity, head tilt and heart rate. [this seems to equate perceptions of safety to a physical fear/flight response]
- The measures are specific to the individual as each person has differing physiological responses based on their life experiences.

- Other measures discussed for perceived safety included possible violations of presumed privacy and fear of unknown factors having to do with the drone operations (e.g. who is the operator and what is their intent)
- Perceived safety violations are more of a driver of reports to law enforcement and complaints to regulators than actual safety violations.
- Actual safety can be impacted by being designed into the product, perceived safety can also be designed into the product, the team discussed above used assistance from a product design and art team to make drones look more “friendly”.
- Actual safety would include improvements in system redundancy. This can be limited in application due to the payload and weight requirements.

3. Subject: Safety Regulatory Structure

- Current structure based on weight seems to be the most logical way to divide classes of drones. This makes the model similar to other regulatory classes (autos, aircraft).
- Perhaps additional divisions of the regulatory structure into more weight classes of drones could be useful as the categories are pretty broad right now with each weight class required to be manufactured with additional redundancy and safety characteristics to mitigate the risk.
- Additional segmentation for operators could influence regulation and insurance based on skill level and experience of the pilot, also composition and design of the drone (e.g. a frangible airframe)
- Question for further discussion, should drone operators be “type rated” as pilots are for large commercial aircraft? Do different drones and the control software have vastly different flight and operational characteristics?
- Belief of the panel: Commercial use will drive regulations. Regrettably, losses will come first and drive public opinion which will influence/prioritize new regulations. The industry is a wait and see mode until a loss experience occurs. Agree.
- Other comments:
 - There are a lot of different redundancy issues for the vehicles and software that need to be addressed in the safety arena.
 - It may require regulation at the manufacturer level to regulate some safety issues. [product liability?]
 - Is it possible to regulate drone operations through the software?

4. Some Assumptions

- An underlying assumption for panel discussions was that this applied to UAS operators who want to follow the law/do the right thing.
- The balance on regulating through software is between education/testing/restrictions and convenience. Because

we assumed that operators want to do the right thing, regulating through software could be effective. However, as with many of the checklists in software, too much results in lack of convenience and good actors sidestepping the rules.

- This is perhaps where operator certification could apply, as a certification number validated through the system could bypass training and tests.
- One of the big reliability issues with software is the cost – as manufacturing processes become more rigorous, cost increases quickly, building barriers to entry.
- Law expert's inputs and discussions included the assumption of malicious actors (as he has experience with South Carolina Corrections and the use of drone by prisoners). In this case, how to deal prevent and deal with incidents is significantly different.

5. Pathway for Operations beyond Part 107

Currently waivers applications are online. Responses to waiver requests usually take about 90 days. Once a critical mass of waivers is reached in a specific area (e.g. night flights) the FAA may change the rules or there may be some standardization of responses. Also airspace authorizations are key, with LAANC coming on line will help a lot.

D. Notes from Discussion 2

1. Cases

- 1 A UAS videoing a major foot race crashes, causing the death of some runners
- 2 UAS down a passenger jet by destroying engines, causing loss of the jet and loss of life
- 3 A UAS flying without authorization prevents a fire fighting tanker from entering the fire zone, causing millions of dollars of property damage in a California urban area

2. Group 1 discussion on Scenario 2

(Bullet points are statements made during a dialogue)

- There is a loss of communication. Should have been geofenced.
- Everyone gets sued – the airline, manufacturer, UAS operator
- Statutory liability for passengers. Everyone gets sued. Point of proximate cause. Tender defense to insurance companies – probably settle.
- Airport contractors require additional drone coverage.
- All drone operations should have liability coverage
- Heightened training requirements if near an airport, additional hours, contractual requirements. Experience requirements for hobbyists. There are waiver applications. This goes back to risk management – was the UAS

pilot properly experienced?

- How would they know the experience?
 - Ask at training (the software asks for the certification level of the operator, number of hours, etc)
 - DJI UAS log flight hours
 - City keep log of drone close calls
 - FAA has sightings database – most entries are bogus. Sightings are reported to police and perhaps to the FAA.
- Study of UAV pilots focused on weather. If not a pilot, don't focus on weather.
- Small UAS are too cheap to worry about crashing, so they don't care.
 - Behaviors better, but not trained
 - Congress said cannot regulate hobbyists and probably most drones out there.
 - As UAS get cheaper, more are buying them
 - More manufacturers also (like GoPro)
 - GoPro failed due to power loss issues – they had to work out the kinks on the first aircraft
 - More manufacturers mean cheaper UAS
 - What about the use of collision avoidance systems
 - Must be bigger and heavier
 - Put transponder in UAS for ADSB to alert planes/helicopters? Add collision avoidance? Technology will advance and likely become smaller
 - Sensor array to ID the pilot must be within the infrastructure
 - Question to Hur about avionics devices running
 - The LiPo batter was a big development for avionics
 - The money going into technology is enormous. Leads to breakthroughs in smaller and more powerful batteries.
 - Fuel cells demonstrated with fixed wing. Proof of concept shown years ago.
 - Syria: small UAS being used to drop grenades
 - Most legal issues with companies will settle due to reputational issues
 - Peachtree road race scenario – who carries insurance coverage that is enough?
 - Anyone involved in the race would be sued – City, PRR, UAS
 - Aerial Photography – Additional coverage for UAS?
 - Human risk managers review to ensure appropriate coverage for applications
 - Often churches, etc are running races. Are they permitting UAS? Checking coverage?
 - Probably not covered
 - Large loss with power transformers?
 - Frequency and severity issues – people don't follow the laws

3. Moving Cases through the Court System

- The hardest part of the process right now would be (for case #2) finding enough parts of the drone equipment to identify the operator.
- For cases #1 and #3, many of the unsafe operations are being traced by postings to social media.
- A third issue is finding the “proof” and being sure that the evidence follows proper chain of custody procedure. E.g. destroying the SIM card where flight data is stored, worries that the logged track could have been changed if not in safe custody.
- Often a drone “violation” (trespassing, unsafe operations) are called in to local law enforcement as first responders. As drone operations are federally regulated, this is usually a dead end. Due to law enforcement response times, the drone and the operator are long gone by the time they arrive.
- Other issues with pursuing a case into the court system, in addition to finding the drone, tracing the operator, and getting the correct authorities involved, once those parties are in place, if there are damages, often the recovery of damages must be allocated between the operator, the manufacturer and the victim. The victim may need to pursue recovery from their own insurance company for damages similar to being hit by an uninsured motorist.
- On the state level, depending on the state law, the percentages of responsibility will be allocated by a jury. (Negligence theory). Some states have a system that if the majority of the fault does not lie with the perpetrator then the victim does not recover.
- Can this be changed with regulations? Another issue raised was enforcement of the regulations.
- Question: Is insurance a driver or reactor to the drone operations environment? Yes, both.
- There are still many gaps in the regulations that will get decided on a case by case basis.
- Disconnection of direct physical control of the drone as an object impacts perceived safety.

4. Pathway beyond 107

(Bullet points are statements made during a dialogue)

- Folks working on beyond visual LOS. 1st person goggles (FPV = First person view) or spotter. They mapped UCF campus using 4-5 spotters. That’s no longer allowed. Could ADSB be put in? Not possible for many systems.
- Hobby associations will push back on an ADSB requirement
- Hobby stay in G airspace. Limit them to 400ft.
- Risks of wind gusts, GPS failures, must harden systems against inappropriate misuse. Too easy to leave/stray out of approved airspace
- US done over Iranian airspace was taken over electronically.
- Hacker controls it.
- Hijack collision avoid systems. Spoof system by transmitting fake vehicle reports

E. Presentation 2 by Dr. Vela – UAS Statistics

- Most flight data is acquired from flight tracking websites to which the operators post their flights.
- The websites are usually shut down soon after discovering that their data is being scraped.
- Amazon Web Services (AWS) generally has pretty unsecured data.
- The data presented covers probably 25-30% of all drone flights – which shows thousands of flights per day.
- Americans are, in general, pretty good about following regulations. Usually around 1%-1.5% of drone flights (max) fall outside the regulations. These calculations are based on out of visual range, too close to airports, too high of altitude, etc.
- Note that more of the flights are likely out of visual range due to obstructions in view (trees and such), but this can't be calculated effectively in this data set (yet).
- Altitude tests (how high can I go) are usually the first test attempted by new UAS operators, and usually go over 400ft in open areas, which is illegal.
- 5% of flights are 500+ feet and within 5nm of an airport

F. Presentation 3 - Malicious use of UAS

- There are not currently any good ways to stop malicious actors from using UAS. The biggest problem is still identification.
- A drone can be flown to a prison carrying many cell phones in a small container, grabbed by a prisoner from a window, and be gone before the guards arrive. The cell phones will already be distributed, and if the guards find the operator, it is typically a kid who was paid to fly it and has no idea where the money came from.
- There are desires to block cell phone and UAS control signals around prisons.
- Blocking UAS signals has been discussed around airports/non-prison settings as well. However, if the blocked signal forces a lost-comms scenario, there is no guarantees of what the UAS will do. Will it land (is it in a safe place to land)? Will it return to base?
- Taking control of a UAS by regulating authorities through a signal has also been discussed. This is similar to the flight baggage model in which all locks have a master lock that TSA can use to open and check. In the case of UAS, this has the potential issue of malicious actors gaining access to these “backdoors”.

G. Discussion 3

1. Goals of the Parties

- Desired advancements? Improved software, awareness of other air traffic for avoidance, issues with “line of sight” operations due to the environment.
- Big/difficult current issue is: What are the rules?

- Operators (recreational, light or entry-level commercial) may not know where to find out. Some education by manufacturers is in place, unknown how much should be their responsibility.
- Controllers (FAA towers, airspace controllers) may not communicate with operators, may not know their responsibility.
- Law enforcement (local police, sheriff, etc) may not know rules.
- Much more education in this space is needed.
- Updates to flight management software in this area might help.
- Huge need for uniformity of access and review of accident reports, incident reports.
- Need for uniformity and standardization in NOTAMS and TFRs to speed dissemination and comprehension by non-professionals.
- Question, should the minimum requirements for operator certification be increased? Will that lead to less compliance?
- Desired advancement in software and hardware: Lost signal guidance (what happens if. . .)
- Additional safety enhancement / damage mitigation planning for commercial operations.
- Ideas
 - “safe zones” for drones to pilot to in case of problems
 - trouble broadcast frequencies and codes similar to aviation’s “guard” frequency and emergency codes used in both aviation and marine
- Idea: Regulators maintaining a listing of authorized or certificated pilots, perhaps for only commercial operators. We already do.
- Much advancement is waiting for commercial operators to take up. Many improvements are not yet required but may come as there is consolidation. Much is waiting for industry maturity.
- Issue: General Aviation and helicopters are not tracked for common flight tracking when flying VFR if not requested. Flight following is available, but there is tracking if filed IFR.

2. Operators

- Other traffic data desired – currently can get ADSB from a website, but not immediately
- Must be careful where point their camera
- Educational issue for operators training – they don’t know what is illegal in many cases: Very true.
- Do operators understand privacy? Cognitive develop process and how it applies here? Do they just not know? Can educational videos help?
- Fun education training exercises for users – when power up a system the first time.
- But purchasers may not be the users

- Similar to the issues for lasers and planes
- Notifications need to be provide to users about warnings – real time versus delayed and GPS accuracy are issues
- What do commercial users need?
 - Single source for rules/laws: FAA provides regulations; lawmakers provide statutes/law.
 - Where are no fly zones, etc: Available on FAA App
 - Can they takeoff/land in national parks? No, due to Department of Interior jurisdiction
- Getting airspace authorization from towers – only have to alert, not get authorization now
- Language is not standardized and is not user friendly. Aviation language is standardized for air traffic and pilots, but new additions are part of the new world of UAS.
- Commercial users and/versus hobby users – what do they need to know?
- Database of operator history and experience – there are companies that are set up/looking into developing databases, but they are struggling in this. Due to privacy issues, but will become more transparent with FOIA
- The industry needs to mature.

H. Discussion Notes 4

1. Insurance and Lawyers

- Desired future of UAS.
- Primary: How to explain the issues, operations, etc to a jury.
- Methods for identification of the parties (operators, equipment).
- Protection from the destruction of data.
- Pilots providing data for research, possible economic incentives?
- Underwriting and claims are the two main touch points when insurance companies will get the best data to price their products.
- Unintended consequences of providing and receiving data and loss history, same for regulation. The need for a feedback loop to correct the unintended consequences.
- Impact of the hacking and mods culture getting around limitations. Encryption limiting the ability.
- Current regulations allow pilots to self-report to NASA/NTSB for accidents and incidents. Are drone pilots aware of this? What information is important to gather?
- Membership organizations: AOPA, AMA, UAS coalition, current major influencers of policy, lobbyists, RC clubs all influencing the space and might be avenues for dissemination of education.
- There might be resistance as some answers the organization's members don't want to receive.
- There is a high level of trust in the manned aviation community for accident reports leading to improvements in aviation instead of punishment. Could this work in the drone community?

- What problem are we trying to solve for: Human deaths seem to be the measure and driver of change, but more due to shock of “mass” casualties (aviation accidents) versus one-by-one casualties (auto).
- Mental separation between the action and results of the action the farther away from the event in space.
- CNN, research on frangible drones.
- Are innovations better driven from the commercial or manufacturer side?
- Will additional regulations (on manufacturers) drive up the cost of production and kill the market?
- Regulatory certainty allows more progress in the field.
- Uber model: Disregard the law and allow it to catch up to the “new reality”. Change management by building a critical mass of change.
- Does safety “sell”? Like in the auto market? Did it always, or is this a recent development? It is always assumed safety is the responsibility of the FAA, so it doesn’t sell. No air carrier advertises that they are safer than brand X, because of “minimum requirements” being met to be certified.
- Are the incentives different in the drone market?
- What uses are drones being put to and does this matter for regulations? Yes, due to the construct of the rules, i.e. agriculture, spraying, dropping of stores, carriage of external goods and property of another across state lines, urban taxis, etc.
- Data safety issues? Is someone stealing/hacking the data produced? Can that cause operator harm?
- Can the software be hacked in a major way to cause harm? (e.g. Were airline reservations systems hacked last year when each of the major airlines had “system crashes” that stranded large numbers of passengers all over the world?) What are the weak points (AWS)?

2. *Insurance and Law*

- Drone legal services – AOPA
- UAS Coalition
- Discussion is analogous to firearms, boats, and motorcycles/dirtbikes

I. Discussion Notes 5 – Where from Here?

- Issues for research: Remote ID (craft and operator), Key as it is the starting point for much of the enforcement environment.
- New licensing models: Equipment size and function.
- Education of operators on existing regulation now that drones are defined as aircraft. Drones have been defined as aircraft since 2005, not just since Part 107.
- Current education providers with classes for Part 107 licenses.

- Integration with STEM in current school curricula. Design and deployment of new courses.
- Atlanta airport developing systems for use and improvement of waivers for 107 ops at airport.

Funding Sources

This research received no external funding.

Acknowledgments

The authors are thankful to Olivia Jagiella-Lodise for her contributions in editing this paper.

References

- [1] Satell, M., “PHILLY BY AIR,” <https://www.phillybyair.com/blog/drone-stats/>, 2020.
- [2] “DJI,” <https://www.dji.com/>, 2018.
- [3] JARUS, “JARUS guidelines on Specific Operations Risk Assessment (SORA),” Report, Joint Authorities for Rulemaking of Unmanned Systems, 06 2017.
- [4] Federal Aviation Administration, U. S. D. o. T., “Safety Management System (SMS),” <https://www.faa.gov/about/initiatives/sms/>, 2018.
- [5] Federal Aviation Administration, U. S. D. o. T., “Compliance Program,” <https://www.faa.gov/about/initiatives/cp/>, 2018.
- [6] Jamoom, M. B., Joerger, M., and Pervan, B., “Unmanned Aircraft System Sense-and-Avoid Integrity and Continuity Risk,” *JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS*, Vol. 39, No. 3, 2016.
- [7] “Remote Identification of Unmanned Aircraft Systems,” <https://www.regulations.gov/document?D=FAA-2019-1100-0001>, 2020.
- [8] Cork, L. R., Walker, R., and Dunn, S., “Fault Detection, Identification and Accommodation Techniques for Unmanned Airborne Vehicle,” Report, Australian International Aerospace Congress, 2005.
- [9] Donato, P. F. A. D., “Toward Autonomous Aircraft Emergency Landing Planning,” Phd thesis, University of Michigan, 2017.
- [10] Ancel, E., Capristan, F. M., Foster, J. V., and Condotta, R. C., “Real-time Risk Assessment Framework for Unmanned Aircraft System (UAS) Traffic Management (UTM),” *17th AIAA Aviation Technology, Integration, and Operations Conference*, 2017. <https://doi.org/10.2514/6.2017-3273>.
- [11] Ramasamy, S., Sabatini, R., and Gardi, A., “A Unified Approach to Separation Assurance and Collision Avoidance for UAS Operations and Traffic Management,” *IEEE International Conference on Unmanned Aircraft Systems (ICUAS)*, 2017.
- [12] Dempster, A., “Upper and Lower Probability Induced by a Multivalued Mapping,” *The Annals of Probability*, Vol. 38, 1967.

- [13] Shafer, G., *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [14] Federal Aviation Administration, U. S. D. o. T., “SUMMARY OF SMALL UNMANNED AIRCRAFT RULE (PART 107),” Report, Federal Aviation Administration, 06 2016.
- [15] “UAS Remote Identification,” https://www.faa.gov/uas/research_development/remote_id/, 2020.
- [16] “Presidential Memorandum for the Secretary of Transportation,” <https://www.whitehouse.gov/presidential-actions/presidential-memorandum-secretary-transportation/>, 2020.
- [17] Federal Aviation Administration, U. S. D. o. T., “Federal Aviation Administration Program Partnerships,” https://www.faa.gov/uas/programs_partnerships/, 2018.
- [18] “AIM Section 2: Controlled Airspace,” https://www.faa.gov/air_traffic/publications/atpubs/aim_html/chap3_section_2.html, 2020.
- [19] Aweiss, A. S., Owens, B. D., Rios, J. L., Homola, J. R., and Mohlenbrink, C. P., “Unmanned Aircraft Systems (UAS) Traffic Management (UTM) National Campaign II,” *AIAA SciTech Forum*, 2018.
- [20] Ferguson, A., and McCarthy, J., “SHARING THE SKIES (SAFELY): NEAR TERM PERSPECTIVE ON SUAS INTEGRATION IN THE NAS,” *Integrated Communications, Navigation and Surveillance Conference (ICNS)*, 2017, pp. 3B2–1.
- [21] Federal Aviation Administration, U. S. D. o. T., “Section 333,” https://www.faa.gov/uas/beyond_the_basics/section_333/, 2018.
- [22] Garner, B. A., *Black’s Law Dictionary (10th ed)*, Thompson Reuters, 2014.
- [23] Knight, J., Smith, B., and Chen, Y., “An essay on unmanned aerial systems insurance and risk assessment,” *2014 IEEE/ASME 10th International Conference on Mechatronic and Embedded Systems and Applications (MESA)*, 2014.
- [24] Koh, C. H., Deng, C., Li, L., Zhao, Y., Tan, S. K., Chen, Y., Yeap, B. C., and Li, X., “Experimental and Simulation Weight Threshold Study for Safe Drone Operations,” *AIAA SciTech Forum*, 2018.
- [25] Lazatin, J., “A Method for Risk Estimation Analysis for Unmanned Aerial System Operation over Populated Areas,” *14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014. <https://doi.org/10.2514/6.2014-2284>.
- [26] Li, J., Monroe, W., and Jurafsky, D., “Understanding neural networks through representation erasure,” *arXivpreprint arXiv:1612.08220*, 2016.
- [27] Kolodner, J., “An introduction to case-based reasoning,” *Artificial Intelligence Review*, Vol. 6, 1992, pp. 3–34. <https://doi.org/10.1007/BF00155578>.
- [28] Volocopter, “Volocopter,” <https://www.volocopter.com/en/>, 2018.
- [29] Meyer, A., “Xavion,” xavion.com, 2018.

- [30] LAKSHMINARAYAN, I., "Model-based Fault Detection for Low-cost UAV Actuators," Thesis, UNIVERSITY OF MINNESOTA, 2016.
- [31] Weibel, R. E., and R. John Hansman, J., "An Integrated Approach to Evaluating Risk Mitigation Measures for UAV Operational Concepts in the NAS," *AIAA 4th Infotech@Aerospace Conference*, 2005, p. 11.
- [32] Weibel, R. E., and Hansman, R. J., "SAFETY CONSIDERATIONS FOR OPERATION OF UNMANNED AERIAL VEHICLES IN THE NATIONAL AIRSPACE SYSTEM," Report ICAT-2005-1, Massachusetts Institute of Technology, 03 2005.
- [33] Weibel, R. E., "Safety Considerations for Operation of Different Classes of Unmanned Aerial Vehicles in the National Airspace System," Thesis, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 2005.
- [34] Lum, C. W., and Waggoner, B., "A Risk Based Paradigm and Model for Unmanned Aerial Systems in the National Airspace," *Infotech@Aerospace*, 2011.
- [35] Rattana-graikanakorn, B., Sharpanskykh, A., Schuurman, M., Gransden, D., Blom, H. A., and De Wagter, C., "Characterizing UAS Collision Consequences in Future UTM," *2018 Aviation Technology, Integration, and Operations Conference*, 2018.
- [36] Breunig, J., Forman, J., Sayed, S., Audenaerd, L., Branch, A., , and Hadjimichael, M., "Modeling Risk-Based Approach for Small Unmanned Aircraft Systems," Report, The Mitre Corporation, 07 2018.
- [37] Neogi, N. A., Quach, C. C., and Dill, E., "A risk based assessment of a small UAS cargo delivery operation in proximity to urban areas," *Proceedings of the 37th Digital Avionics Systems Conference (DASC)*, 2018.
- [38] Luxhoj, J. T., Joyce, W., and Luxhoj, C., "A ConOps derived UAS safety risk model," *Journal of Risk Research*, 2017.
- [39] Kumar, V., Wieland, F., Toussaint, S., and Luxhøj, J. T., "UAS (Unmanned Aerial System) Safety Analysis Model (USAM)," *14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014.
- [40] Ferreira, R. B., Baum, D. M., Neto, E. C. P., Martins, M. R., Almeida Jr., J. R., Cugnasca, P. S., and Camargo Jr., J. B., "A Risk Analysis of Unmanned Aircraft Systems (UAS) Integration into non-Segregate Airspace," *International Conference on Unmanned Aircraft Systems (ICUAS)*, 2018.
- [41] Denney, E., Pai, G., and Johnson, M., "Towards a Rigorous Basis for Specific Operations Risk Assessment of UAS," *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*, 2018.
- [42] Weibel, R. E., Edwards, M. W. M., and Fernandes, C. S., "Establishing a Risk-Based Separation Standard for Unmanned Aircraft Self Separation," *Ninth USA/Europe Air Traffic Management Research & Development Seminar*, 2011.
- [43] Weinert, A., Campbell, S., Vela, A., Schuldt, D., and Kurucar, J., "Well-Clear Recommendation for Small Unmanned Aircraft Systems Based on Unmitigated Collision Risk," *JOURNAL OF AIR TRANSPORTATION*, Vol. 26, No. 3, 2018.
- [44] Washington, A., Silva, J. M., and Clothier, R., "A Review of Unmanned Aircraft System Ground Risk Models," *Progress in Aerospace Sciences*, 2017.

- [45] Atkins, E. M., Khalsa, A. D., and Groden, M. D., "Commercial Low-Altitude UAS Operations in Population Centers," *9th AIAA Aviation Technology, Integration, and Operations Conference*, 2009.
- [46] Melnyk, R., Schrage, D., Volovoi, V., and Jimenez, H., "A third-party casualty risk model for unmanned aircraft system operations," *Reliability Engineering and System Safety*, Vol. 124, 2014, pp. 105–116.
- [47] de Castro Fortes, J. L., Fraga, R., and Martin, K., "AN APPROACH FOR SAFETY ASSESSMENT IN UAS OPERATIONS APPLYING STOCHASTIC FAST-TIME SIMULATION WITH PARAMETER VARIATION," *Winter Simulation Conference*, 2016, pp. 1860–1871.
- [48] Ford, A. T., and McEntee, K. J., "Assessment of the Risk to Ground Population Due to an Unmanned Aircraft In-Flight Failure," *10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, 2010.
- [49] Lum, C. W., Gauksheim, K., Kosel, T., and McGeer, T., "Assessing and Estimating Risk of Operating Unmanned Aerial Systems in Populated Areas," *11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, including the AIA*, 2011.
- [50] Ijaz, K., Sohail, S., and Hashish, S., "A Survey of Latest Approaches for Crowd Simulation and Modeling using Hybrid Techniques," *17th UKSIM-AMSS International Conference on Modelling and Simulation*, 2015, pp. 111–116.
- [51] Kang, W., and Han, Y., "A Simple and Realistic Pedestrian Model for Crowd Simulation and Application," *ArXiv*, 2017.
- [52] Deville, P., Linard, C., Martin, S., Gilbert, M., Stevens, F., Gaughan, A., Blondel, V., and Tatem, A., "Dynamic population mapping using mobile phone data," *Proceedings of the National Academy of Sciences*, Vol. PNAS Early Edition, 2014. <https://doi.org/10.1073/pnas.1408439111>.
- [53] Di Donato, P. F. A., and Atkins, E. M., "Evaluating Risk to People and Property for Aircraft Emergency Landing Planning," *JOURNAL OF AEROSPACE INFORMATION SYSTEMS*, Vol. 14, No. 5, 2017.
- [54] Kaiser, M. S., Lwin, K. T., Mahmud, M., Hajializadeh, D., Chaipimonplin, T., Sarhan, A., and Hossain, M. A., "Advances in Crowd Analysis for Urban Applications Through Urban Event Detection," *IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS*, Vol. 19, 2018, pp. 3092–3112.
- [55] Tzelepi, M., and Tefas, A., "Human Crowd Detection for Drone Flight Safety Using Convolutional Neural Networks," *25th European Signal Processing Conference*, 2017, pp. 743–747.
- [56] Du, X., Dori, A., Divo, E., Huayamave, V., and Zhu, F., "Modeling the motion of small unmanned aerial system (sUAS) due to ground collision," *Journal of Aerospace Engineering*, Vol. 232, 2018, pp. 1961–1970.
- [57] Vian, J., and Moore, J., "Trajectory optimization with risk minimization for military aircraft," *Journal of Guidance, Control, and Dynamics*, Vol. 12, No. 3, 1989, pp. 311–317.
- [58] Johnson, E. N., "Georgia Tech UAV Research Facility," <http://www.uavrf.gatech.edu/platforms/gust/>, 2018.

- [59] Chowdhary, G., and Johnson, E. N., "Adaptive Neural Network Flight Control Using both Current and Recorded Data," *AIAA Guidance, Navigation, and Control Conference*, Vol. AIAA 2007-6505, 2007.
- [60] Kannan, S. K., and Johnson, E. N., "Adaptive Control with a Nested Saturation Reference Model," *American Institute of Aeronautics and Astronautics Conference*, 2003.
- [61] Twigg, S., Calise, A., and Johnson, E., "On-Line Trajectory Optimization Including Moving Threats and Targets," *AIAA Guidance, Navigation, and Control Conference*, Vol. AIAA 2004-5139, 2004.
- [62] Logan, M. J., Gundlach, J., and Vrana, T. L., "Design Considerations for Safer Small Unmanned Aerial Systems," *AIAA SciTech Forum*, 2018.
- [63] Yomchinda, T., Horn, J., and Langelaan, J., "Flight Path Planning for Descent-phase Helicopter Autorotation," *AIAA Guidance, Navigation, and Control Conference*, 2011. <https://doi.org/10.2514/6.2011-6601>.
- [64] Hejase, M., Kurt, A., Aldemir, T., Ozguner, U., Guarro, S. B., Yau, M. K., and Knudson, M. D., "Quantitative and Risk-Based Framework for Unmanned Aircraft Control System Assurance," *JOURNAL OF AEROSPACE INFORMATION SYSTEMS*, Vol. 15, No. 2, 2018.
- [65] Fang, S. X., O'Young, S., and Rolland, L., "Online Risk-Based Supervisory Maneuvering Guidance for Small Unmanned Aircraft Systems," *JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS*, Vol. 41, No. 12, 2018.
- [66] Lu, Y.-J., and He, J., "Dempster-Shafer Evidence Theory and Study of Some Key Problems," *JOURNAL OF ELECTRONIC SCIENCE AND TECHNOLOGY*, Vol. 15, No. 1, 2017. <https://doi.org/10.11989/JEST.1674-862X.5030211>.
- [67] Dunham, J., Johnson, E., Feron, E., and German, B., "Automatic Updates of Transition Potential Matrices in Dempster-Shafer Networks Based on Evidence Inputs," *Sensors*, preprint.
- [68] Dunham, J., and Johnson, E., "Unmanned Aerial Systems Health Monitoring Architecture," *IEEE Aerospace Conference*, 2019.
- [69] Guttorp, P., and Thorarinsdottir, T. L., "What Happened to Discrete Chaos, the Quenouille Process, and the Sharp Markov Property? Some History of Stochastic Point Processes," *International Statistical Review*, Vol. 80, 2012, pp. 253–268.
- [70] VELOS ROTORS, L., "VELOS ROTORS," <https://velosuav.com/>, 2020.
- [71] Murphy, C. K., "Combining belief functions when evidence conflicts," *Decision Support Systems*, Vol. 29, 1999.
- [72] Zhang, Z., Liu, T., Chen, D., and Zhang, W., "Novel Algorithms for Identifying and Fusing Conflicting Data," *Sensors*, Vol. 14, 2014.
- [73] EMLID, "EMLID," <https://emlid.com/>, 2020.
- [74] BoYang, J., and LingXu, D., "Evidential Reasoning Rule For Evidence Combination," *Artificial Intelligence*, Vol. 205, 2013.
- [75] Shenoy, P. P., "Valuation-based systems: A framework for managing uncertainty in expert systems," *Fuzzy Logic for the Management of Uncertainty (L. A. Zadeh and J. Kacprzyk, Eds.)*, 1992, pp. 83–104.