Team Fire Scout

Technological Feasibility

10/23/2020

Sponsors:

Dr. Fatemeh Afghah

Alireza Shamsoshoara

Mentor:

Sambashiva Kethireddy

Team Members:

Nicholas Bollone

Matthew Briody

Jacob Hagan

Kenneth Klawitter

Drew Sansom

The purpose of this document is to discuss the feasibility of the Fire Scout drone project as a whole by examining the technologies it uses, and how they interact with one another.

Table Of Contents

Table Of Contents 1

1.0 Introduction 2

2.0 Technological Challenges 3

3.0 Technological Analysis 4

- **3.1 Drones** 5
  - **3.1.1 Drone Body** 5
  - **3.1.2 Drone Cameras** 8
- **3.2 On-board Hardware** 11
  - **3.2.1 Processor** 12
  - **3.2.2 Cameras** 14
  - **3.2.3 Communication** 16
  - **3.2.4 Battery** 16
- **3.3 AI Technologies** 17
  - **3.3.1 Frameworks** 17
  - **3.3.2 Algorithms/Methodologies** 21
- **3.4 Graphical User Interface** 25

4.0 Technological Integration 28

5.0 Conclusion 29

6.0 Sources 31

1.0 Introduction

Each year wildfires across the globe take lives and cost the economy billions in damages. They burn homes, businesses, national lands, and cannot be fought without risking the lives of first responders. As of October 2020, the California Wildfires have burned down an estimated 2.4 million acres of land across 6,750 fires. The indirect death alone could total 1,200 individuals from reduced air quality and has cost dozens of lives directly.

When fighting wildfires, collecting and analyzing information is crucial to prevent the spread of fire and fully contain it. Fuel sources need to be analyzed, the location and direction of the fire need to be known, and ideally, this information should be coupled with the location of individuals relative to fires to prevent loss of life. Currently, satellite imagery and manned aircraft are used to collect this kind of information, however, these methods are expensive and not real-time. This leaves a critical information gap between what is happening amid a wildfire, and what the stakeholders and decision-makers think is happening.

As technology advances, the methods and tools used to fight fires need to as well. One way to revolutionize the way wildfires are fought is with semi-autonomous Unmanned Aerial Vehicles (UAVs), commonly called drones. UAVs are relatively inexpensive and can be outfitted with sensors and cameras to collect thermal and visual imagery, which allows artificial intelligence models to analyze the fire. These analyses can be performed on the drone or from where the drone was deployed to provide comprehensive and near-real-time information on where fires are located, how they are clustered together, where they are spreading, and if they are threatening bystanders. Most importantly these drones remove humans from the heart of the fire and allow them to collect data in a way that is not only more precise than ever before, but safely and remotely.

The Fires Scout Project aims to harness these advantages that were identified by Dr. Fatemeh Afghah: A member of the School of Informatics, Computing, and Cyber Systems at NAU. Dr. Afghah recognized the need to revolutionize how fires were fought, and with the help of her Ph.D. candidate Alireza Shamsoshoara, the Fire Scout Project was born.

2.0 Technological Challenges

As with all novel and revolutionary technologies, some challenges need to be understood and addressed before the product can come to life. The current scope of the Fire Scout project is to outfit Northern Arizona University's drones with the hardware and software needed to process images on-board the drone, relay that information back to a user, and display it for them. Each step of this pipeline needs to be broken down into an individual and modular component whose technological issues will be analyzed and overcome. These components involve:

1. The drone body. Specifically, it needs to carry payloads that can analyze fires, have an extended flight time, and withstand a variety of weather conditions.
2. The drone's proprietary hardware. This hardware includes the cameras used to collect and train AI models, as well as navigate the drone body.
3. The hardware the Fire Scout team will outfit to the drone. This includes a small processor and additional cameras. By developing the system independently of the drone, it can be attached to any UAV.
4. The AI frameworks used for image analysis. These will be used to develop an AI model capable of classifying fires and will be deployed on the on-board processor.
5. The AI algorithms and techniques that the frameworks use. Unless properly chosen, the frameworks may not run efficiently or reach the desired accuracy.
6. The desktop GUI. This allows users to interact with the Fire Scout drone and displays the information it collects. Fire Scout needs to appeal to a variety of stakeholders and display what it analyzes in a meaningful way.

3.0 Technological Analysis

What follows is how each of the above problems was analyzed, and how a solution to each was chosen. Many factors were considered, including but not limited to budget, provided equipment, current skills and abilities of the Fire Scout Team, and how current choices would affect the future expansion of the project. By careful analysis of the available options, the Team designed a system that allows for seamless integration of the chosen components.

Along with thorough explanations on how hardware and software choices were made, there are graphs that summarize the Team's findings on the suitability of a component. The legend for the tables is as follows:

- = Meets and exceeds needs (ideal and preferred solution)
- = Meets needs (ideal solution)
- = Meets some needs (minimal solution)
- = Barely meets needs (sub-optimal solution)

3.1 Drones

Drones are widely used for their limitless capabilities in high-risk situations. The use of drones enables first responders and other emergency services to do reconnaissance without risking human life. Due to their ability to carry a variety of payloads from light equipment to cameras, move in 3 dimensions, hover in place, and fly at speeds above 40 mph, drones provide first-responders with many new possibilities. These abilities make them perfect for dynamic situations like fire surveillance.

3.1.1 Drone Body

The drone body is the core of the system since it carries the cameras, on-board processor, power banks, and anything else Team Fire Scout may need to attach. As such, the body must withstand the extra weight. The "body" of the drone is the drone itself and does not refer to any additional hardware or cameras attached to it.

Desired Characteristics

To be considered for use, the drone body must meet some basic requirements. First, it must be able to fly for an extended period and withstand a broad range of weather conditions. The drone will likely be subjected to a range of temperatures and fly in decreased visibility situations, so it must keep operating in these situations.

Second, the drone must be capable of carrying a small amount of weight even if the tradeoff is a decrease in flight time. This allows the drone to be outfitted with multiple cameras, sensors, and processors to help analyze fires.

Alternatives

The drones provided to the Fire Scout Team by Northern Arizona University are the DJI Matrice 200, and the DJI Phantom 3 Pro. The DJI Matrice 200 is the newer of the two and is suited for a wider range of applications. The Matrice has an ecosystem of accessories that can be purchased and installed, allowing the user to customize it to their needs.

The DJI Phantom 3 Pro differs in some ways. Unlike the Matrice, it has a built-in camera, built-in gimbal, and a more compact size. The drone can plan its path ahead of flight time and has a convenient “return home” button that will bring it back to the pilot. This could be very useful for quick drone battery swaps, as well as the external power bank that will be utilized by the on-board hardware (see Section 3.2.4).

Analysis

Without being able to field test the drones, the best way to analyze them is to look at the manufacturer’s claimed specs. To start, both have similar flight capabilities. They have the same ascent speed of 5m/s and descent speed of 3m/s. The hover accuracy of the two drones is matched: both have a vertical hover accuracy of ±0.1m, and a horizontal hover accuracy of ±0.3m. Additionally, the vertical accuracy for the GPS positioning is ±0.5m, and ±1.5m horizontally so the team can take confidence in the drone’s claimed location.

However, there are some areas where the DJI Matrice 200 outperforms the DJI Phantom 3 Pro. The Matrice has a wider operating temperature from -4℉ to 113℉, while the Phantom Pro 3, can operate from 32℉ to 104℉. Given the colder temperatures of Northern Arizona and the hot atmospheres above fires, the Matrice seems advantageous.

Finally, the standard battery for the Matrice gives it a maximum flight time of 27 minutes, which is four minutes longer than the maximum flight time of the Phantom 3 Pro. With the upgraded TB55 battery for the Matrice, the maximum flight time is increased to 38 minutes, significantly higher than the Phantom.

CHOSEN APPROACH

While the drones are similar, the preferred drone for Fire Scout is the Matrice 200. Compared to the Phantom 3 Pro, the Matrice outperforms in some critical categories including battery capacity and flight time. The TB55 battery almost doubles the mAh offered by both the Phantom 3 Pro and the Matrice 200's original batteries. This is extremely advantageous since the ability to fly for longer distances and times is immensely important. Even though the Phantom 3 has multiple backup batteries that can be swapped during operation, the Phantom 3 would still have to land more frequently than the Matrice.

Finally, while the Matrice 200 is preferred, it may not be possible to use it as the primary drone since other NAU labs use the Matrice as well. Therefore, both drones will be used depending on their availability.

	DJI Phantom 3 Pro	DJI Matrice 200
Battery	●	●
Max Speed	●	●
Max Service Ceiling Above	●	●
Flight Time	●	●
Temperature Range	●	●

Overview of Drone Functionality

FEASIBILITY

There are a variety of industries that use the DJI Matrice 200 drone. For example, when paired with thermal camera imaging, the Matrice 200 has been used to test temperatures of building structures and power lines. The DJI Matrice 200 offers organizations money-saving, time-effective, and safer alternatives to the traditional human intelligence-gathering procedures.

The DJI Phantom Pro 3 is also utilized in a variety of real-world applications. With its imaging abilities and relatively low cost, it is widely used in the real estate industry to capture videos and images of properties. The Phantom is also used to inspect solar panels and power-lines. Finally, it is used for public safety inspections including highly trafficked areas/buildings where the data it collects helps create evacuation routes.

3.1.2 Drone Cameras

The cameras for drones provide them with their capabilities. Commonly called payloads, the proprietary cameras offered by DJI for their drones include thermal and high-definition cameras that Fire Scout will use to fly the drones (they can also record video feed for later image analysis, but cannot be used for real-time analysis). There are different options for both the Phantom and Matrice, and depending on what drone is available, will determine what cameras Fire Scout uses. Therefore, it is necessary to have an overview of all the hardware options for both drones.

Desired Characteristics

The cameras must be able to autofocus so when the drone is flying at different rates and altitudes, the image quality remains clear. The cameras must be DJI compatible as well so that they can send live feedback to the operator and use the gimbal to stay steady during flight.

The cameras need to be a minimum of full-HD but do not need to be 4K (although it is certainly welcome). Additionally, they must be compact and efficient so that they do not drastically disrupt flight time. Finally, while DJI makes thermal cameras, the thermal camera used by Fire Scout is for the on-board hardware (Section 3.2.2) so the Team will not need to purchase a DJI compatible one.

Alternatives

The HD camera provided for the Matrice is the Zenmuse X4S Pro. It is built into the three-axis Zenmuse X4S gimbal that allows for quick adjustments of viewing angles, and a stable image during flight. It is a powerful camera equipped with a large sensor suitable for object-detection and image-processing.

The DJI Phantom 3 Pro has a camera/gimbal setup that comes pre-installed on the drone. The camera has a 20mm lens that can capture high-quality images and videos. It offers a variety of still photography and video recording modes for a range of conditions.

Analysis

The Zenmuse X4S offers a one-inch sensor, which by comparison, is two-thirds larger than the iPhone 7's sensor. Having a larger sensor size is vital to image quality. This, coupled with its 20-megapixel resolution and 11.6 stops of dynamic range, makes it an ideal choice. The camera also has 4K/60 and 4K/30 capabilities both with a 100Mbps bitrate, well above the Team's minimum requirements. The lens is an 8.8mm/F2.8-11 lens with low distortion and a radical dispersion of 3um. It has a high MTF rating performance and low pass filters. Additionally, this camera can shoot in both JPEG and DNG photos at up to 14 frames per second in "Burst Mode." Finally, the Zenmuse X4S Pro can take photos during video recording allowing it to collect both kinds of data Fire Scout needs at once.

The DJI Phantom 3 Pro camera has a 1/2.3" CMOS sensor with a 20mm lens that has a 94-degree field of view. This camera has an electronic shutter speed of 8-1/8000 s for a wide range of drone speeds and conditions. The Phantom's camera may not be as technically advanced as the Zenmuse but is still a great camera.

Chosen Approach

The Zenmuse X4S Pro is the camera of choice. It has a more advanced gimbal that allows for a wider range of views, and alongside the drone's sensors, it helps the drone remain safe when flying in harsh conditions. Coupled with its higher image quality, it is the ideal camera for Fire Scout although the DJI Phantom 3 Pro's camera is sufficient as well. Both cameras have high-quality gimbals that can stabilize the image during flight.

	Zenmuse X4S Pro	Phantom Pro 3 Camera
Weight	●	●

Operating Temperatures	●	●
Gimbal	●	●
Sensor	●	●
Lens	●	●

Overview of HD Camera Functionality

Feasibility

The Zenmuse X4S Pro and the Phantom 3 Pro camera have both been utilized in real-world applications such as power grid inspections, infrastructure analysis, and precision agriculture. Finally, they both have been used to collect data for Fire Scout's AI model. These cameras open the possibilities to observe dangerous situations, such as a growing fire, with a level of precision and accuracy unobtainable in the past.

3.2 On-board Hardware

The on-board hardware will be attached to the drone manually by the Fire Scout team. This hardware is not a part of the drone, but rather, the equipment used to run the AI model, extra cameras, and Software Defined Radio (SDR). While newer DJI drones are compatible with DJI's Software Developer Kits including their On-board Software Development Kit which allows the drone and its camera to be programmed manually, the older Phantom 3 Pro and Matrice 200 are not. Therefore, to implement image recognition, AI models running on custom hardware and third-party cameras need to be attached.

3.2.1 Processor

The processor is responsible for running the AI model against the camera input and communicating with the GUI through the SDR. By using a third-party processor such as an NVIDIA Jetson Nano or a Raspberry Pi, the Fire Scout system can be installed on any drone capable of carrying the necessary hardware.

Desired Characteristics

There are three primary requirements the on-board hardware needs to meet. One, the hardware must be light enough for the drone to carry without drastically impacting flight time. Second, it must be powerful enough to run image processing algorithms, collect images from the camera, and transmit signals to the SDR in real-time. Finally, the device must be compatible with HD and thermal cameras for image collection and processing.

Alternatives

Two micro-computers were compared and contrasted: the NVIDIA Jetson Nano and the Raspberry Pi 4. These were initially chosen as candidates by Ali Shamsoshoara because they are commonly used for image analysis, are from two reputable companies, and meet the basic qualifications. Both can implement AI frameworks and are small enough to be attached without drastically affecting drone performance. Each micro-computer is also compatible with multiple third-party cameras.

Analysis

The NVIDIA Jetson Nano has a GPU designed for AI applications and is newer to the market than the Raspberry Pi. The Jetson Nano GPU is based on an NVIDIA Maxwell architecture with 128 NVIDIA CUDA® cores and can process information slightly faster than the Raspberry Pi's Broadcom VideoCore VI CPU. Though the Jeston Nano is geared towards AI applications due to its large on-board GPU, the community around the Jetson Nano is significantly smaller than the community around the Raspberry Pi since it is a newer device.

The Raspberry Pi 4 is lighter than the Jetson Nano. Weight is very important since it leads to longer flight times, and the Jetson Nano weighs 248 grams while the Raspberry Pi 4 only weighs 43 grams. The Raspberry Pi 4 is also more commonly used in industry and startup applications, giving the Raspberry Pi 4 more resources and open-source software to help with development. Finally, although it has a less powerful GPU, the Raspberry Pi 4's CPU is still powerful enough to run an AI model and communicate with the SDR.

The processor on the NVIDIA Jetson Nano is a 64-bit Quad-core ARM A57 1.43GHz CPU while the Raspberry Pi 4 has a Broadcom BCM2711 SoC with a 1.5 GHz processor. Though they are very similar in GHz, the Raspberry Pi 4 is faster at information gathering, sending, and receiving through the SDR which is important for real-time data delivery.

Chosen Approach

Fire Scout will use the Raspberry Pi 4. Its lighter weight will lead to extended flight times which benefits the project as a whole. Additionally, its larger user base, ample tutorials, and Team Members' prior experiences with Raspberry Pis will streamline deploying AI models. While not as powerful as the NVIDIA Jetson Nano, it is still sufficient to run a fully trained AI model and has gained the client's approval.

	Raspberry Pi 4	NVIDIA Jetson Nano
Community Resources		
CPU		
GPU		
Weight		

Overview of Micro-Computer Functionality

Feasibility

The Raspberry Pi 4 has been used in many video and image processing projects such as security cameras, facial recognition systems, and object detection systems proving itself capable for Fire Scout. Forums created for Raspberry Pi devices are filled with open-source projects and examples of its uses. These will help Fire Scout learn from others and gain a better understanding of how neural networks can be run on the Raspberry Pi.

3.2.2 Cameras

The cameras are responsible for capturing the images for analysis. The cameras Fire Scout will attach to the Raspberry Pi will be an HD camera for use in high visibility situations, along with a thermal camera for low visibility situations. However, unlike the drone's cameras which are used for navigation and initial data collection, the Raspberry Pi's cameras will be used for the AI model.

Desired Characteristics

The cameras have several requirements. One, the cameras must be able to autofocus as the drone moves at different speeds and elevations. Second, the cameras must be light enough to not hinder the flight capabilities of the drone. Finally, like the DJI's cameras, full HD recording is necessary for the HD camera, but not 4K, and the thermal camera does not need to be HD, just clear enough for a human to identify the shape of a fire.

Alternatives

The first HD camera is the Arducam Auto Focus Camera. The Arducam has manual zoom autofocus with a software-controlled moveable lens that can be programmed through the OpenCV framework (discussed in Section 3.3.1). The camera has a range of 8, 10, or 12 megapixels allowing it to record 1080p or 4K resolution. Finally, it has a detachable thermal/night vision lens available for purchase. While the autofocus will require extra programming, this is a cost-efficient $35 camera.

The second option is the SEE3CAM. This camera has an autofocus feature that does not need to be programmed. The 4K autofocus SEE3CAM 130 offers 13-megapixel imaging. While this can autofocus manually and has a higher resolution, the Arducam's ability to communicate with OpenCV may prove advantageous at ensuring the images captured work with the chosen frameworks. The SEE3CAM is also less cost-effective at $99.

The thermal camera given to Fire Scout is the FLIR Vue Pro R. There are 12 different variations of the FLIR, and the product variation that was provided to team Fire Scout is the Vue Pro R 336, 45° FOV, 6.8mm, 30Hz. With a forty-five degree angle for the field of view, a 6.8mm lens size, and a framerate of 30Hz, the camera exceeds Fire Scout's needs for thermal imaging. With an adjustable zoom feature, precision mounting holes, and compact lightweight design, the FLIR Vue Pro R is perfect for the Matrice 200. The weight of this camera is just under 120g, and the size including the lens is 2.26” x 1.75” making it optimal for drone use. Finally, this camera has an operating temperature from -4℉ to 122℉ with an operational altitude of plus 40,000 feet.

Analysis

The Arducam can photograph objects at both the millimeter level and long-distance. The Arducam allows for precise autofocus through OpenCV allowing the camera to capture clear images. The camera is small enough to not weigh down the drone and has many different lens attachments. There are numerous resources for the Arducam as well which will help if any difficulties are encountered.

The SEE3CAM 130 comes equipped with built-in autofocus and an electronic rolling shutter. It is plug-and-play with no drivers required. This camera offers a full high definition resolution in 1080p at 60fps or a 4K resolution. Additionally, the SEE3CAM has an operating range of -22℉ to 158℉ so flying in extreme weather will not damage it.

FLIR Vue Pro R is a high-functioning thermal-imaging camera that captures calibrated radiometric temperature measurements pixel by pixel. The thermal camera is extremely important for this project since in very smoky conditions it may be the only way to view the fire, and provides information on the fire’s temperature. Finally, it is compatible with both drones.

Chosen Approach

The Arducam will be the camera of choice. The programmable autofocus coupled with the ability to attach thermal and HD camera components makes it ideal. With a resolution of 1080p, low cost, and a light weight, it meets the requirements. The SEE3CAM would also work but its higher price and lack of accessories may be problematic. This is subject to change based on the Electrical Engineers’ input who are also members of Fire Scout.

The FLIR Vue Pro R is of course the chosen approach since there was no thermal camera alternative; however, it will still be invaluable to Fire Scout. It offers exceptional thermal imaging and will be useful when analyzing fire temperatures and for image segmentation.

	Arducam	SEE3CAM
Auto Focus	●	●
Megapixels	●	●
Resolution	●	●
Weight	●	●
Attachments	●	●

Overview of Camera Functionality

	FLIR Vue Pro R
Weight	●
Operating Temperatures	●
Gimbal	●
Thermal	●
Lens	●

Overview of Thermal Camera Functionality

Feasibility

There have been many different applications of the Arducam using the different lens attachments. It has been used on drones before the Fire Scout project, and for hobbyist activities such as time-lapse photography. Finally, Arducam's large community provides many tangible uses and implementations that will help with the camera. Though the Arducam appears to work for the CS team, it needs to work with the EE team as well, something that cannot be decided upon at this time. Finally, the provided thermal camera has already been used to collect data for the AI model.

3.2.3 Communication

The Pi will transmit data to the desktop GUI through a software-defined radio (SDR). This component will be implemented entirely by the Electrical Engineering team working on

Fire Scout and will interface with data signals created by the Computer Science team, allowing the drone's user to select what information they want to be displayed on the GUI. This information includes, but is not limited to, still-frame images of fires, video feed, real-time AI classification, segmentation, and GPS coordinates.

3.2.4 Battery

The battery that will power the above mentioned on-board hardware, is the INIU Power Bank. This was chosen by the client who provided Team Fire Scout with three of these batteries, specifically the INIU Power Bank 10,000 mAh model: B1-B1, with 2.4 A max output. They can be quickly swapped out as needed to supply constant power to the Pi and its hardware. They deliver their power through a light 7-gram charger. Finally, their dimensions are 5.2" L x 2.7" W x 0.5" H making them very easy to attach to either of the drones.

3.3 AI Technologies

The artificial intelligence model is the heart of Fire Scout. It will be trained on local machines before being deployed on the on-board processor. Much of the real-world value depends on how accurately the drone can analyze what it sees. Since AI is a large and growing field, there are multiple options for the frameworks and types of algorithms that can be implemented on the drone. These include neural networks, such as Deep Neural Networks and Convolutional Neural Networks, as well as classic image processing techniques. What is important is that the chosen frameworks and algorithms provide highly accurate data that can inform decision-makers.

3.3.1 Frameworks

The artificial intelligence and image analysis frameworks are the software that the Fire Scout team will use to create and train AI models. They are the toolkits that allow the development team to focus on the implementation of artificial intelligence, rather than the creation of algorithms from the ground up. Choosing the correct framework will speed development.

Desired Characteristics

Ideally, these AI frameworks need to support a multitude of machine learning algorithms and techniques. Since Fire Scout's end goal is to implement Artificial Neural Networks (ANNs), the frameworks should support two main kinds of ANNs: Deep Neural Networks (DNNs) which process structured data, and Convolutional Neural Networks (CNNs) which process unstructured data.

The frameworks need to be proven and have a large user base. The first advantage this offers is quicker software development and faster problem solving for challenges that might be novel to the development team. Additionally, since Fire Scout will be used in high stake situations, the proven track record of a well-established framework inspires confidence in the developers and shareholders that the underlying code is reliable. Finally, choosing a popular framework means that after the current Fire Scout team graduates, it is more likely that the next team will be familiar with the chosen frameworks and will be able to benefit from the same advantages as the current team.

Since AI frameworks can be large and complex, a large user base also means there will be numerous tutorials that are advantageous to development. Rather than trying to implement an entirely new solution, any open-source contribution or tutorial can be used as long as it is cited properly.

The two final necessities for AI frameworks are compatibility: the frameworks need to be compatible with the chosen hardware, and compatible with each other. This again illustrates why having a large user base is advantageous; there is a greater likelihood that someone has combined one framework with another, or ran the software on a Raspberry Pi already.

Alternatives

Currently, one of the most popular AI frameworks is TensorFlow. Developed and maintained by Google for ANNs, it is used in-house for applications such as Google Translate and supports languages including Python, C++, and R.

The main rival to TensorFlow is PyTorch. Developed by Facebook, it is primarily used in research applications and is also designed for ANNs. It is written in Python and has a C++ interface as well.

Both TensorFlow and PyTorch have large user bases and implement many of the same features. Since they are the predominant ANN frameworks, when one framework implements a feature users find desirable, the other is quick to copy it and implement their own version of said feature. The results of this competition are very capable and functional frameworks.

Similar to these but still distinct is Keras, which is another framework for ANNs but uses a much higher level, and less configurable API. It emphasizes ease of use and abstracts the more complex parts of ANN models away from the developer. This means that Keras is easier to learn than PyTorch or TensorFlow, and models can be created faster with Keras.

Finally, OpenCV is an open-source computer vision library used by companies such as Google, Intel, and Toyota. It supports machine learning and can implement facial recognition, object detection, and real-time image processing. Like the other frameworks, it has a Python and C++ interface, a large user base, and is free to use. It differs from the above frameworks in that while it does have the ability to implement DNNs, its primary purpose is for computer vision programming. It excels at traditional image processing, rather than image processing using ANNs, which is useful for prototyping and various object detection applications.

Analysis

Many of the frameworks are new to the members of the Fire Scout Team, so testing each one would be a very time-consuming way to analyze them. Instead, careful research was conducted to evaluate which frameworks met or exceeded the desired characteristics.

The best way to conduct this research was to find examples where each framework has been used, how it was used, and how many tutorials were available for each. When possible, facts and figures about the sizes of the user base were weighed out as well. This gave a strong indication of what framework had the most lively user base, and therefore, the most useful resources.

To evaluate the compatibility of frameworks with one another, examples were found of projects and tutorials that featured combinations of the potential frameworks working together. Developer testimonies were also considered when evaluating how well each framework fit together. Finally, if one framework officially supported another in its API, it was given priority.

Chosen Approach

After careful analysis, Fire Scout decided to use TensorFlow, Keras, and OpenCV. These frameworks are the most compatible with each other and have the most lively user bases with some tutorials even existing on how to use them to detect fires.

Keras was chosen first because it allows for the quick configuration and deployment of ANNs. This is great for prototyping networks and later refining them to detect fires. Additionally, it is the most beginner-friendly of the frameworks and approachable by all members of the current Fire Scout team.

This led to the decision to choose Tensorflow over PyTorch. While both excel in creating ANNs, as of 2019, Google included Keras into the TensorFlow API due to its high use with TensorFlow in the past. This means that the two are essentially part of the same framework and can be used within the same program. Higher-level aspects of the neural networks can be set up in Keras, and then configured in TensorFlow as needed. Finally, TensorFlow has a larger user base and more tutorials which will help development.

OpenCV was chosen because it fills in the gaps. TensorFlow and Keras are designed for Neural Networks, and while OpenCV does have some ANN capability, it will be the go-to framework for quick object detection and segmentation of fires. It has an expansive user base, a multitude of tutorials, and can be used to solve any image-related problems that may arise.

Below is a table that summarizes these findings. Worth noting is that OpenCV and the AI Frameworks implement image analysis and Object Detection differently. So, OpenCV scores higher in these categories not necessarily because it is better, but because it is designed specifically for image processing.

	Keras	TensorFlow	PyTorch	OpenCV
ANNs	●	●	●	●

User Base	🔵	🔵	🟠	🔵
Available Tutorials	🔵	🔵	🟠	🔵
Image Analysis	🟢	🟢	🟢	🔵
Ease of Use	🔵	🟠	🟠	🟢
Object Detection	🟢	🟢	🟢	🔵
Compatibility	🔵	🔵	🟠	🟢

Overview of Framework Functionality

Feasibility

While no proof of concept has been implemented for any of the frameworks, there are examples and tutorials of each framework being used in a way that is similar or identical to our intended uses. Keras has not only been used to identify fires, but instructions on how to do so are free online. OpenCV comes with object detection methods for cars and people, two of the objects Fire Scout needs to detect, and multiple tutorials exist of TensorFlow being used to create CNNs to detect objects at high accuracy. Finally, there are examples and instructions of these frameworks being installed and run on Raspberry Pi devices. This means that there are concrete examples of the chosen frameworks being used to meet the hardware and software requirements in every main use case of Fire Scout.

3.3.2 Algorithms/Methodologies

Although the chosen frameworks are powerful tools, without the knowledge of how to use them and what kind of algorithms to implement, they do not provide much value. Therefore, it is important to have an understanding of what algorithms exist and how they can be used to classify fires, segment fires, as well as to detect objects such as people and vehicles. The AI algorithms and methodologies will use a training set of data to create a fire classification model for the Raspberry Pi. This is the algorithm that will analyze fires in real-time, and relay its information back to the GUI via SDR (see section 3.2.3).

Desired Characteristics

The Fire Scout Team's datasets come from videos and pictures taken by people who previously worked with Dr. Afghah. These videos, mainly .mov and .mp4 files, were recorded by the drones proprietary cameras. When Fire Scout was given these videos, the team parsed out the frames and separated them into fire and no fire datasets (binary classification). So, the algorithms chosen must be able to process this flow of data, specifically the fire and no fire JPEG datasets.

Next, these algorithms should be able to be trained on the frameworks Fire Scout has already chosen. After properly trained, the algorithms need to have low overhead and run efficiently on a Raspberry Pi. Finally, they need to be accurate.

Alternatives

Artificial Neural Networks (ANNs) are inspired by the biological neural networks that are found in the human brain. ANNs can be either shallow or deep. Shallow is when there is one hidden layer (the layer between input and output), and deep is when there is more than one hidden layer leading to Deep Neural Networks. ANNs consist of many different types but the Team's analysis was between Deep Neural Networks (DNNs) and Convolutional Neural Networks (CNNs).

DNNs can be supervised or unsupervised but must have multiple layers. With an increase in layers comes an increase in accuracy. CNNs allow for image classification. Specifically, they can observe the changes in pixels based on RGB or HSI scale and detect edges, which is useful for detecting fires.

In a supervised learning method, the AI algorithm begins to learn with a labeled dataset (fire or no fire). It receives input values with the correct output values and trains itself until it produces results equal to the correct output values. Common supervised learning algorithms include logistic regression, Naive Bayes, and Random Forest.

Unsupervised learning models have unlabeled datasets and can not be directed or controlled. Since datasets are unlabeled, the programmer cannot control how that data is structured. Unsupervised learning instead clusters the data together by patterns it finds automatically, and without the guidance of a programmer.

Analysis

DNNs work with structured data, such as tables and spreadsheets, while CNNs work with unstructured data such as photos which can be difficult to collect, process, and analyze. By adding more layers to a neural network, the network is given more capabilities which leads to more accurate predictions, so the networks used by Fire Scout will need to be deep.

Supervised learning algorithms, like classification, are trained to recognize what the developer wants them to. This is ideal for image detection of unique shapes like fires. With unsupervised learning, there is no control over data classification or data clustering. This means the model may not correctly identify fires. This is problematic.

Chosen Approach

Due to this, Fire Scout will be using supervised learning. It is relatively easy to create fire and no fire datasets, and since supervised learning can train a model to reach the desired output, it can train the model to recognize fires based on the existing datasets. Coupled with CNNs that process unstructured data through a deep network, Fire Scout will be able to predict with high accuracy what regions of land have fires or not. Since CNNs are commonly used for image recognition, they are the right choice of ANN.

While some of these methodologies and algorithms are not directly comparable, the table below shows how each one meets Fire Scout’s requirements. This gives an overview of each one’s strengths and weaknesses.

	DNN	CNN	Unsupervised	Supervised
Overhead	●	●	●	●
Time	●	●	●	●
Ease of Use	●	●	●	●

Accuracy	🔵	🔵	🔴	🟢
Image Analysis	🟠	🔵	🟡	🟢

Overview of Algorithm Functionality

Feasibility

With the fire and no fire datasets, frameworks like OpenCV and Keras can be used to implement supervised learning and CNNs. However, Fire Scout needs to collect more data before this can be done successfully. The technology is promising because there are tutorials and online examples of people using the Team's chosen frameworks and algorithms to detect fires. There are also a variety of projects that implement CNNs with Keras. Finally, this approach is not novel since major companies, like Yelp and Facebook, use CNNs to process images.

3.4 Graphical User Interface

Data is useless if it is not presented in a convenient and readable fashion. There must be a visual way for the end-user to interact with the information being processed. This introduces the need for a Graphical User Interface (GUI).

Fire Scout's GUI is going to be delivered on a ground-based unit such as a laptop. The data presented will be received via the software-defined radio (SDR) implemented by the engineering team. The results of their implementation will determine what information will be available to present on the GUI.

Desired Characteristics

Something very important to the Fire Scout team is accessibility. The GUI needs to be easy to use and present intuitive ways to control the drone for all parties such as firefighters and search and rescue teams. The information the GUI presents includes the location of the drone, as well as the feed from the Raspberry Pi camera.

The GUI should also be customizable. For example, if the user is focused on one camera feed over the other (thermal over HD), the GUI should allow the user to view only the desired camera feed. This reinforces usability and personalization making our end-user experience a positive one.

Alternatives

There are two contenders for GUI frameworks: TkInter and Kivy. TkInter is a Python binding to the Tk GUI toolkit and is the standard Python interface for building GUI applications. TkInter is characterized by its large community, simple design, and ease of use.

Similar to TkInter, Kivy is an open-source Python library that helps create innovative user interfaces. Originally released in 2011, its most stable release was in July 2019 so it is still being maintained and used. Kivy is compatible with Linux, Windows, Mac OS X, Android, and Raspberry Pi making it flexible for all kinds of OSs. It is typically used for creating touchscreen interfaces and games.

Analysis

TkInter is a mature application that has been used for countless projects over its lifetime. It is an ideal tool for quick deployment since it implements Python, a language all Fire Scout team members are familiar with, and has a very large user base. This user base has led to extensive TkInter tutorials, some even in conjunction with OpenCV. TkInter, however, is often criticized for its execution speed. This could be a problem since the GUI must display video feed in real-time.

There is not nearly as much documentation on Kivy when compared to TkInter since the framework is significantly newer. Its main advantages over TkInter are its cross-platform capabilities and its frequent use for video game displays (something that is irrelevant for Fire Scout). Kivy looks to be an impressive tool, but the lack of community is concerning.

Chosen Approach

For this project, TkInter looks to be the most promising framework for developing the GUI. The active and mature community, as well as native support in Python, are key benefits. Though Kivy is integrated with Python as well, it is not included in the standard library that comes with all distributions of Python like TkInter does. Kivy feels excessive for the scope of our project since it's typically used for video game user interfaces; this proves to be far more complex than what Fire Scout's GUI needs to be. Kivy's lack of community support also makes TkInter the more suitable framework.

	TkInter	Kivy
Maturity/Community	●	●
Python Integration	●	●
Speed	●	●

Overview of GUI Functionality

Presented below is a rough sketch of what the GUI may look like, but will likely change as the Team gets feedback from the clients. Some initial features may include:

- HD camera display showing object detection.
- Thermal camera display to assist in the decision of whether the drone is looking at a fire or not.
- Coordinates to assist in locating the fire.
- Measurement of travel speed and direction of movement.

While this is a rough sketch, the GUI is expected to evolve based on the client's demands, as well as what information the Electrical Engineering team's SDR is capable of communicating (refer to Section 3.2.3).

Feasibility

TkInter has been used for many GUIs, including some similar to Fire Scout's rough sketch. One example is a project from PyImageSearch.com. This project used TkInter to create a GUI that would capture frames from a video feed on demand.

A larger project that uses TkInter is IDLE, the standard IDE that comes with all installations of Python. It uses TkInter to create its environment that includes functionality such as opening/editing Python scripts, as well as some debugging features.

4.0 Technological Integration

While all of these technologies are great on their own, they only provide value when pieced together to create the final product. Starting with the drone, its first task is to collect data to train the AI model Fire Scout will create in Tensorflow, Keras, and OpenCV. Its proprietary camera will be used to collect this data and to navigate the drone itself via the DJI controller. Since DJI already has software that can save the feed from its drone's camera, no software needs to be created or installed on the Raspberry Pi to complete the same task, making it ideal for initial and future data collection.

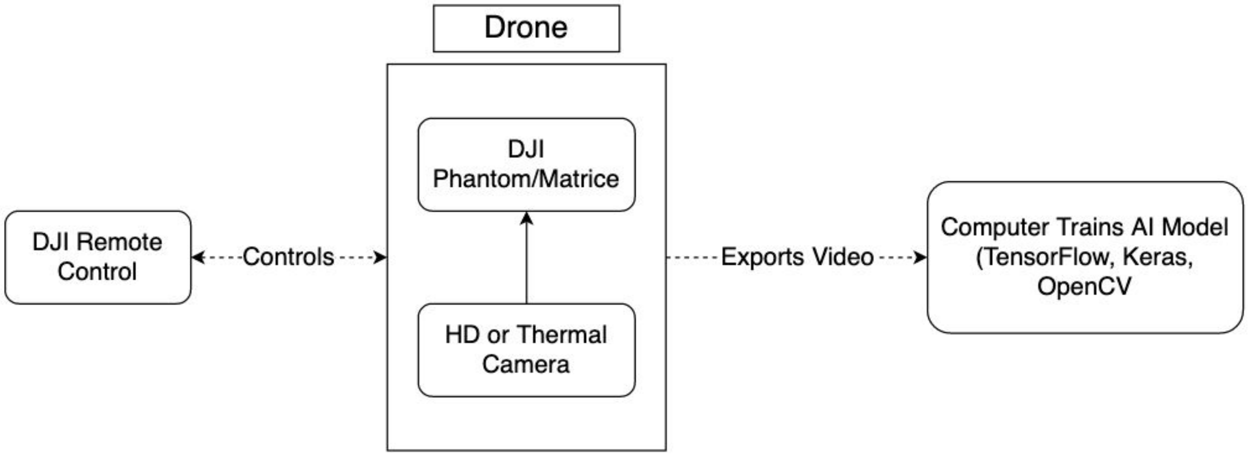

Figure 1

When deploying the drone in a fire, its proprietary camera is invaluable for navigation but is useless for real-time analysis. This is because DJI's video stream is proprietary, making it impossible to tap into on the drone itself. Since the goal of Fire Scout is to have a drone analyze fires in real-time, the Fire Scout System needs to be carried on the drone as well (which is the drone's second task).

Figure 2

After training the AI models with the data collected from the drone, the models will be deployed to the Raspberry Pi 4 (Figure 2). From here, the Fire Scout System is assembled and attached to the drone (Figure 3). This system performs real-time analysis and is made up of non-DJI components. Its platform independence means any drone body that can carry this system can be used to analyze fires.

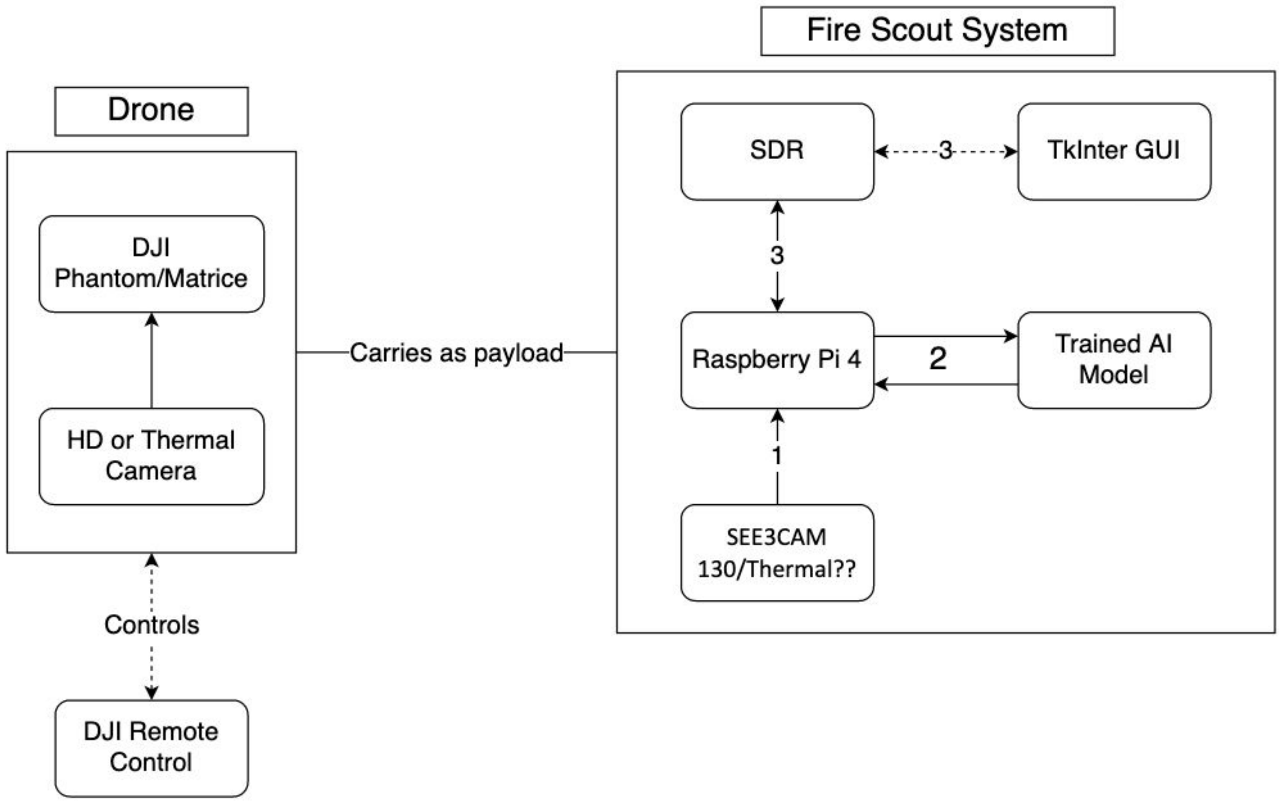

Figure 3

The Fire Scout System's flow 1) gathers information from the Arducam, 2) runs the machine learning algorithms deployed on the Pi, and 3) communicates with the TkInter GUI via the SDR so information can be seen by responders, and so that signals can be passed back to the Raspberry Pi (figure 3).

5.0 Conclusion

With how devastating wildfires are and how quickly technology is evolving, the opportunity to change how they are fought is more feasible than ever. A novel solution that prevents the loss of life, billions in economic damages, and keeps responders safe becomes not just an ideal possibility, but a feasible solution.

Fire Scout is the path forward for fighting wildfires. When using Unmanned Aerial Vehicles responders are kept out of harm's way. By providing real-time information and aerial images of the fire, stakeholders are kept informed of the fire's status in an unprecedented manner. Coupled with deep learning, the information provided is more meaningful than traditional image collection. The AI model can quickly and accurately show where fires are, how they are clustered together, and alert responders of humans in proximity to the fires.

Additionally, by creating a modular system based on third-party components such as the Raspberry Pi and its accessories, Fire Scout can be deployed on any drone that can carry the necessary hardware, and on any processor powerful enough to run the ANN models. This allows the system to be applied to a wide range of situations by many kinds of users and developers. It is the perfect blend of technology and inspiration relying on proven methods like AI and UAVs, yet using them in new ways.

In the future, the system could run on dozens of drones communicating with each other as fire's temperatures, directions, and characteristics are saved and mapped in real-time. The drones could be programmed to fly in sync with one another, to cover large areas with greater efficiency. If more advanced capability is needed, DJI's newer drones feature APIs that let developers view their proprietary video feed and control the attached cameras directly. The models trained for the initial stages of Fire Scout could be deployed on these more advanced systems. Ultimately, Fire Scout presents a new way forward when fighting wildfires, one that is safer, more powerful, saves more lives, and provides more value than ever before.

6.0 Sources

1. https://www.fire.ca.gov/stats-events/
2. https://www.cnbc.com/2020/09/11/wildfires-death-toll-rises-in-california-oregon-and-washington-.html
3. https://marutitech.com/top-8-deep-learning-frameworks/
4. https://blog.udacity.com/2020/05/pytorch-vs-tensorflow-what-you-need-to-know.html
5. https://opencv.org/about/
6. https://www.geeksforgeeks.org/opencv-overview/
7. https://realpython.com/pytorch-vs-tensorflow/
8. https://kivy.org/#home
9. https://wiki.python.org/moin/TkInter
10. https://en.wikipedia.org/wiki/Kivy_(framework)
11. https://www.oreilly.com/library/view/python-programming-on/1565926218/ch20s01s02.html
12. https://www.dji.com/zenmuse-x4s/info#specs
13. https://www.dji.com/phantom-3-pro/info
14. https://www.dji.com/matrice-200-series/info#specs
15. https://www.flir.com/products/vue-pro-r/?model=436-0019-00
16. https://www.pyimagesearch.com/2016/05/30/displaying-a-video-feed-with-opencv-and-tkinter/