University of Southern Queensland

Faculty of Health, Engineering and Sciences

Using UAV's and machine vision in the early detection of combine harvester fires.

A dissertation submitted by

Steven Mosetter

In fulfilment of the requirements of

ENG4111 and 4112 Research Project

towards the degree of

Bachelor of Engineering (Honours) (Mechanical)

Submitted October, 2016

Abstract

Fire during the harvest of crops is an ever present hazard. The combination of hot and dry conditions with a highly flammable crop material creates perfect conditions for fire to start and propagate, the result of which can be loss of production, time, equipment and the crop itself.

The aim of this project is to create a system that can actively detect fire activity so that the harvester operator has a better chance of containing the fire before it spreads out of control. By using the ability of CCD cameras to detect Near Infrared (NIR) and sophisticated machine vision, a cheap and effective fire detection system can be created that can alert the operator to any developing fire before the grows out of control.

An extensive review of available literature regarding combine harvester fires, the use of Near Infrared (NIR) and visual light cameras in fire detection and the use of machine vision to detect fire was conducted. An experimental prototype NIR camera system was constructed with off the shelf components selected on the basis of suitability and cost and a computer program was developed with the purpose of detecting fire in the video feeds.

Testing was done in two phases. The first phase was to test the hardware of the system to determine if the cameras was even able to see fire or related phenomena. The second phase of testing was to determine if the machine vision software was able to quickly and accurately identify fire under different circumstances, and its ability to filter out other phenomena that may cause false positives.

The hardware of the system was able to detect fire in most circumstances. Inexpensive cameras operating in the NIR and Visual spectrums are more than capable of seeing the light, heat and smoke emissions of the fire under all of the conditions that such a system would likely encounter during normal operations. The emissions that the camera detects is highly dependent on the proximity of the camera to the fire, which has significant implications on the software processing algorithm and its ability to accurately detect a fire.

The software algorithm was able to correctly identify a fire during all the software tests. The NIR camera was able to correctly identify fire in all of the testing but was subject to false positives from reflections. The colour program had greater success in bright conditions as the reduced contrast between the fire and its surroundings enabled the colour of the fire to be more easily seen.

Detecting fire with machine vision is still a field that is in its infancy, but the results gained from this project are very promising and with further development could yield a system able to reliably detect fire in harvest conditions.

University of Southern Queensland

Faculty of Health, Engineering and Sciences

ENG4111 & ENG4112 Research Project

Limitations of Use

Certification

I certify that the ideas, designs and experimental work, results, analyses and conclusions set out within this dissertation are entirely my own effort, except where otherwise indicated and acknowledged.

I further certify that the work is original and has not been previously submitted for assessment in any other course or institution, except when specifically stated.

Steven Mosetter

Student Number: 0061042297

Acknowledgements

I would like to thank everybody that has helped me to complete this project. In particular I would like to thank my father, Russ Mosetter, for giving me the idea in the first place and for his practical insight on many matters relating to harvesters.

My supervisor, Tobias Low, for guiding me through the dissertation process and for his vital feedback.

My good friend, Trent Pascoe, for assisting me during the testing phase.

My girlfriend, Shenae Browne, for putting up with me

Finally, to all my friends and family for their support.

Contents

Table of Figures

Tables

1. Introduction

Fire has always been a major hazard for farming operations and has far reaching consequences. Fire can often result in the loss of production, time, equipment and the crop itself; it also has the chance to cause serious injury or death of workers.

Harvest is a particularly hazardous time as it combines hot and dry conditions with a very flammable crop. The risk of fire is only set to worsen as climate change will create longer periods of hot and dry conditions, changes in harvester design that create a more efficient, clean and higher producing machine also create additional fire hazards on the machine.

Currently no new combine harvester has any type of fire detector or fire suppression system. Only this year has a third-party company started to supply its own fire suppressions system. For everybody else the only fire detection they have is their own senses and the method of fire suppression is a handy fire extinguisher.

1.1 An Introduction to the combine and harvest

The modern combine harvester is a versatile machine designed to efficiently harvest a variety of grain crops from the field to deliver clean grains for further processing into food and material for human beings and livestock. (Miu, 2016)

It combines the following three processes

- Reaping, the cutting and gathering of the crop.
- Threshing, loosening the edible part of the grain from the scaly, inedible chaff that surrounds it.
- Winnowing, separating the edible part of the grain from the chaff.

Figure 1: A typical modern combine harvester hard at work. (Lawford, 2010)

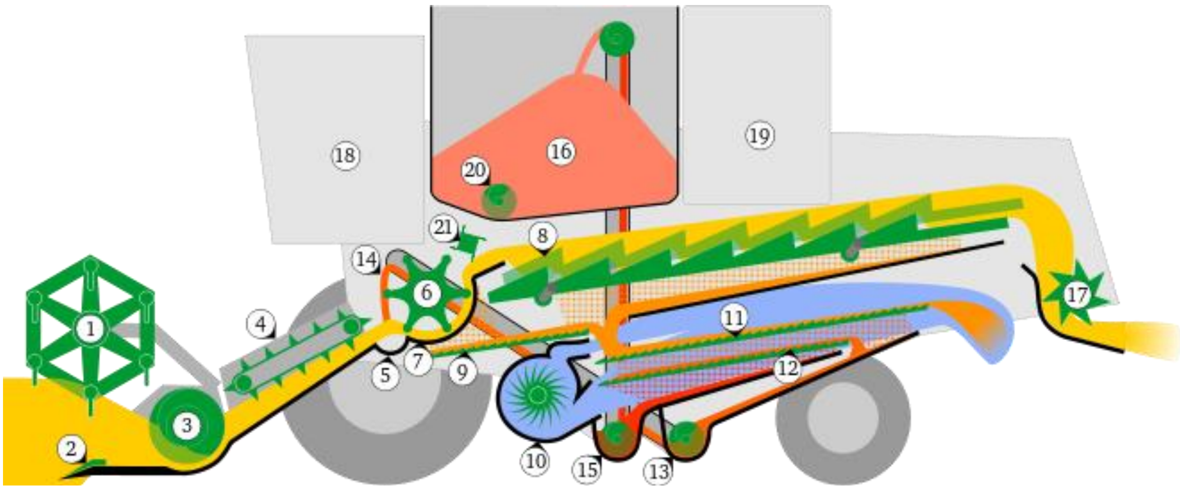

Figure 2: A basic cross section of a typical 'walker-type' harvester. Parts numbers 1 through 4 perform the reaping part of the process. The threshing drum (6) and concave (7), thresh the grain. The straw is carried out by the straw walkers (8) that also vibrate to shake out any additional grain. Parts 10 through 15 perform the winnowing process. Part 16 is grain storage, part 18 is the operator cab and part 19 is the engine. (Hans Wasthuber, 2009)

The demands placed on the modern combine harvester many and varied, the machine must be able to harvester the crop at just the right time when the crop reaches its peak and often only has a very short window to do this once this window is reached, sometimes less than one week. It must be able to harvest a variety of different crops and has a modular construction enabling different parts and sub-assemblies to be meet this demand, and even these parts must have the ability to be adjusted on fly to compensate for different crop conditions.

The increasing size of the average farm and the demand for more and more product has meant that the capacity of the machines has grown over time, but since the size of combine harvesters has reached the limiting width of most roads, cost-effective improvements in capacity must come from increasing the overall efficiency of the machine. (Miu, 2016)

The modern combine harvester is a very large, complex and highly capable machine that processes a very large amount of grain in a very short amount of time. Because of this, modern combine harvesters have increasingly become more and more costly, in the range of hundreds of thousands of dollars for the average machine and nearing a half a million dollars or even more for the very largest, highest capacity machines. (Quick, 2010)

In Australia the main winter harvest occurs between the months of September to February with harvest starting and ending earlier in the more northern states. This coincides with the hotter and drier weather of the spring and summer months in the grain growing regions. Once harvest has begun, farmers, workers and contractors will work from sunup to sunset harvesting, moving, storing and processing crops until the job is completed. 12 hour days are standard and 18 hour days are not uncommon. This leaves very little time for even essential maintenance to be completed.

An unfortunate and very much unwanted part of harvest is fire, as explained below

Australian broadacre harvest conditions are arguably the most hazardous in the world for fires. Each year there are hundreds of harvester fire incidents and approximately a dozen half-million dollar-plus machines burnt to the ground at harvest. In some instances, there are associated crop losses as well. (Quick, 2010)

Figure 3: A harvester that has caught fire during operation. (Law, 2012)

The most effective strategy to combat combine harvester fires is regular clean down and inspection of the machine, preventing the build-up of flammable materials and preventing the creation of possible sources of ignition respectively. In South and Western Australia, if the prevailing environmental conditions become too extreme, as observed by the Grassland Fire Danger index (GFDI), by law the harvester must stop. (Grains Research and Development Corperation, 2013)

In a perfect world, there preventative measures would be enough to stop harvester fires from occurring. Unfortunately, sometimes fire incidents do happen despite the use of best practice. The aim of this system is to provide an extra level of protection to the harvester under these extreme conditions.

1.1 Prevalence of combine harvester fires

Information regarding the prevalence of harvester fires is somewhat scarce. In Australia there has been no major study done on harvester fires, that is available in the public domain. The major combine harvester manufactures may have done their own independent studies but they have not released this information. Looking further afield there have been some studies conducted in the United States but the latest of these studies was finalised in 2002 using information dating from 1984-1997 (J.M. Shutske, 2002).

The combine harvester prevention and control summit investigated 8307 combine fires between the years of 1984 and 1995 and the year 1997, unfortunately there was no information for the year 1996. This information was drawn from the National Fire Incident Reporting System (NFIRS) across 38 states. In addition, an additional 620 combine fires that occurred between 1998-2000 from the top 5 states of the previous subset were also evaluated. From this information the study drew the following conclusions:

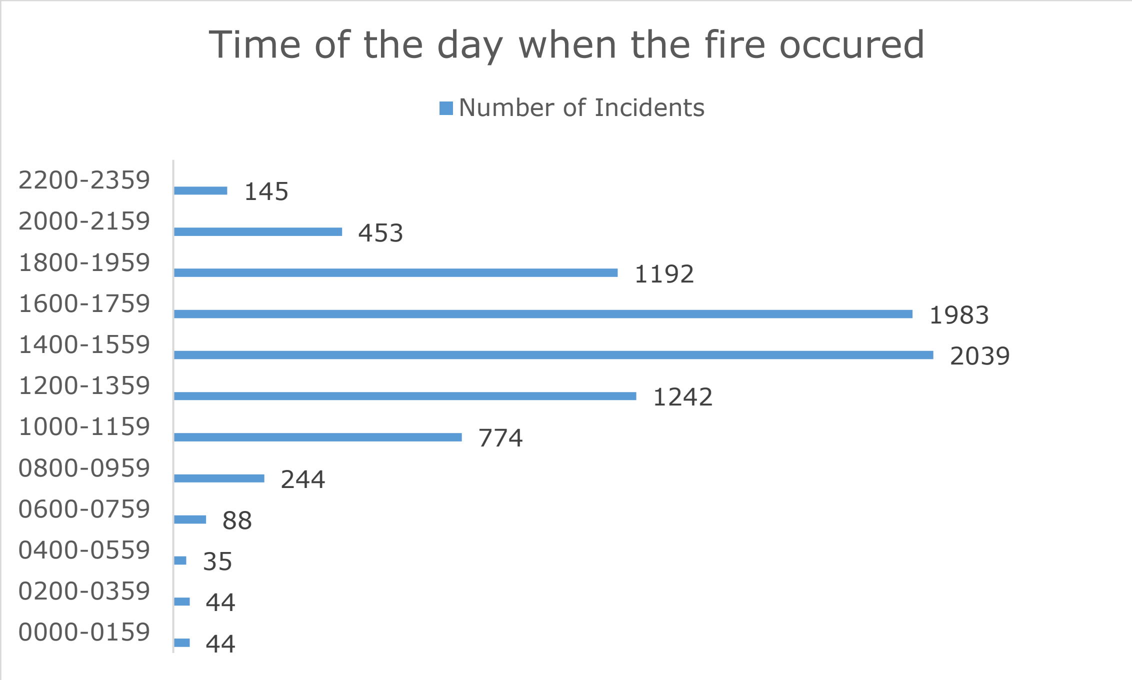

Figure 4: Time of the day when the fire occurred. (J.M. Shutske, 2002)

- 78.2% of combine fires occur between noon and 8:00PM, 48.5% occur between 2:00PM and 6:00PM.
- The majority of fires occur during the week, with the fewest on Sunday. 1984 – 1997 data suggests a higher rate of fires in the middle of the week with 1998 – 2000 data suggesting a shift to the end of the week.
- 67.9%, of fires occurred during the fall harvesting period (late September-November) with a decrease in the frequency of fires during wheat harvest and an increase in fires during the fall harvest from 1984 to 2000.
- 639 reported combine fires occur, on the average, each year in those states that report to the NFIRS.
- 47.2% of combine fires reported mechanical or electrical failure as the ignition factor starting the fire.

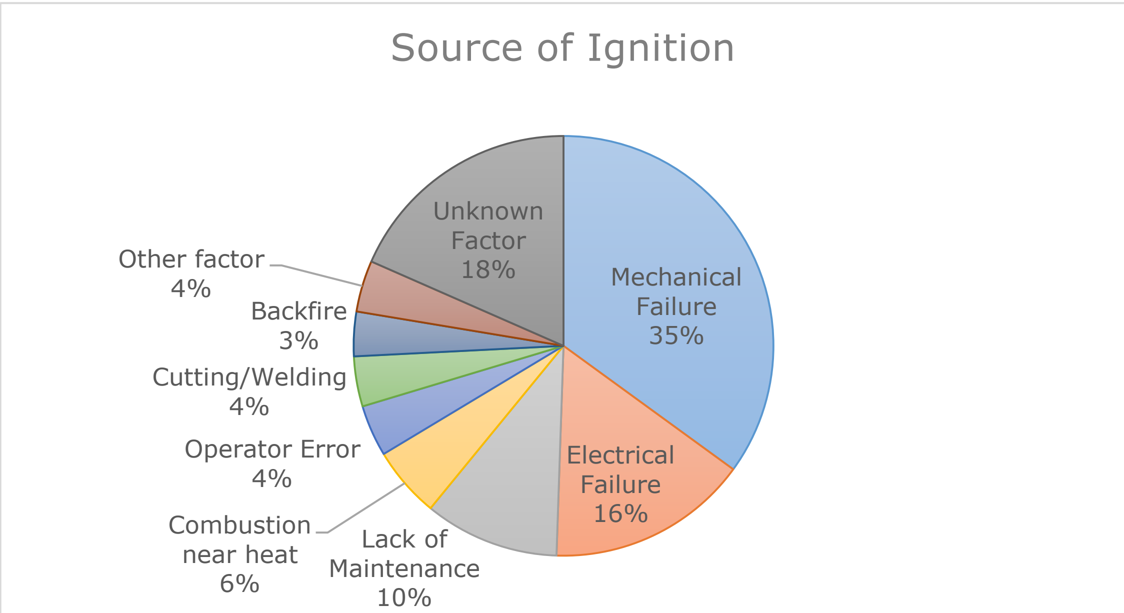

Figure 5: Source of Ignition. (J.M. Shutske, 2002)

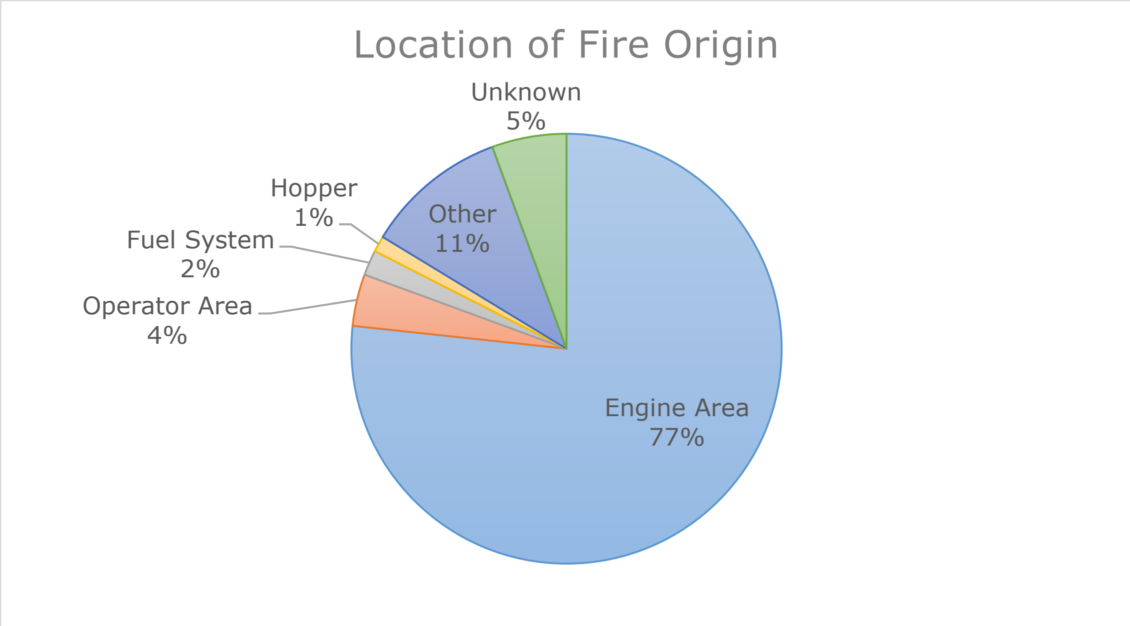

Figure 6: Location of fire origin. (J.M. Shutske, 2002)

- 76.7% of combine fires originate in the engine area
- 41.3% of combine fires have organic material as the type of flammable material first involved in the fire.
- From 1984 – 1997, $94,748,050 in estimated losses from combine fires were reported, averaging $15,182 per fire.
  (J.M. Shutske, 2002)

A second older study collected data relating to 4092 combine and tractor fires between 1984-1988 with the vast majority of these (3655) coming from the NFIRS. The rest came from onsite investigations conducted by the researchers (265), the Indiana State Fire Marshal's office (122) and from surveys (50). From this the researchers came to similar conclusions as the newer paper.

- 67% of the NFIRS fires occurred between 10:00am and 6:00pm, with the largest number of fires for all data sets occurring between 2-4:00PM.
- This study did not look at what days the fires occurred on nor did it list what months of the year the fires occurred in.
- 40% of the fires originated near hot components (24% exhaust and 16% hot engine surface), 34% originated from engine electrics for the 50 Indiana combine fires.
- 62.4% of the fires originated in the engine area for the 3655 NFIRS combine and tractor fires.
- 40.0% of the Indiana combine fires involved crop residue as the primary flammable material for the fire.

(J.M. Shutske, 1990)

The only Australian report I was able to find was by Dr Graeme Quick, unfortunately I found the report to be of limited utility as it lacked definite figures. The report did however agree with the results found by the two US studies. That around three-quarters of combine harvester fires start in the engine bay, that a large proportion of these fires are started by crop residue collecting in and around the engine bay which is then set alight by hot exhaust components.

The report noted that new model harvesters have larger, more powerful engines and use new emission control systems that result in an increase amount of heat rejection. Add into this newer farming practices such as desiccating the crop, that is to spray the crop with herbicide to remove green crop material and green weeds that may have otherwise damped fire risks and a change in the type of crops being harvested have probably also contributed to an increase fire risk for Australian farmers.
(Quick, 2010)

During the course of gathering information I contacted the Queensland Fire and Rescue Service (QFRS), the Country Fire Authority (CFA) in Victoria and the Country Fire Service (CFS) requesting any information regarding combine harvester fires. Of these only the QFRS responded. The QFRS provided statistics regarding fires that were attended by the QFRS where 'agricultural equipment' was identified as the source of ignition. Unfortunately, the QFRS only began using a separate 'agricultural equipment' ignition code in their record keeping in July of 2015, limiting the amount of useful information to be gained. In addition, there was no information regarding what the type of agricultural equipment was nor was there information regarding the outcome of the fire. However, the statistics provided tell us that between July 2015 and March 2016

- There was a total of 45 fires where agricultural equipment was identified as the ignition source in Queensland.
- October 2015 was the peak month for fires with 11 incidents, this concedes with the start of the winter harvest in Queensland. August 2015 had the second most incidents with 8 total. September 2015, January and February 2016 were equal third with 5 incidents per month.
- The most common type of fire was identified as 'Mobile Property fire' with 18 incidents overall. The second most common type of fire was identified as 'Scrub/bush/grass fire' with 17 incidents.
  (Queensland Fire and Rescue Service, 2016)

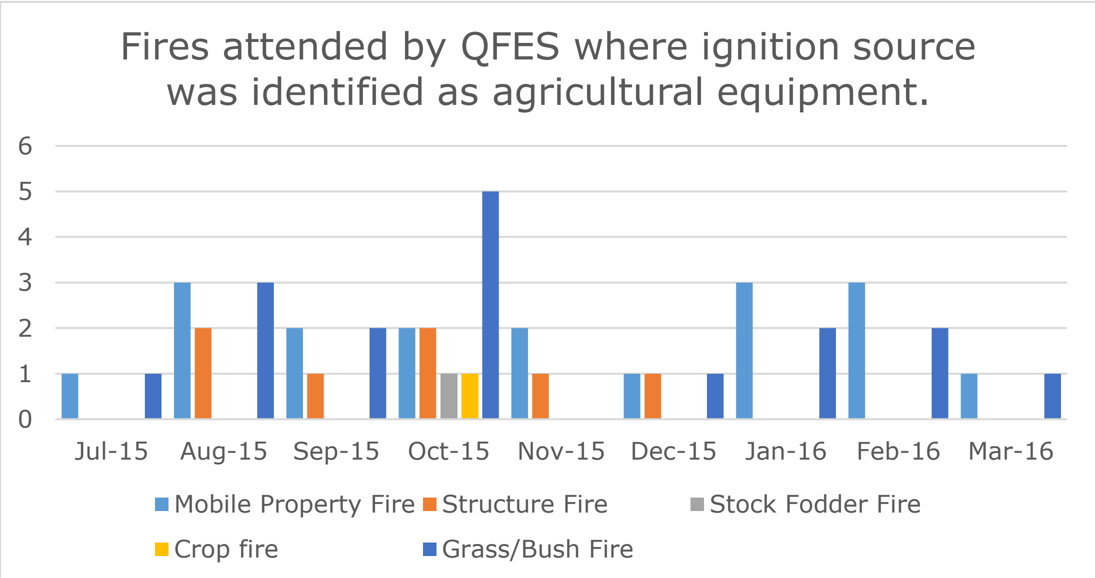

Figure 7: Fires attended by QFRS where the ignition source was identified as agricultural equipment. (Queensland Fire and Rescue Service, 2016)

The data from these studies show that the majority of combine harvester fires start in the engine bay and this is where the primary focus of the fire detection system should lie. The two major fuels for these fires are crop residue and hydrocarbon fuels and the detector systems should be tuned as such to look for specific markers from these types of fires.

1.1 Detecting and extinguishing the fire

Early detection is the critical factor in containing combine harvester fires. Often if the response is not immediate the fire will quickly become too big for one persons to control and even if the fire is controlled, the combination of hot fuel vapours and hot metal surfaces will often cause the fire to reignite.

Table 1: Elapsed time- fire ignition to extinguisher use. (J.M. Shutske, 1994)

Time	Percent
0-1 Minutes	18.8%
1-2 Minutes	25.0%
2-3 Minutes	46.9%
> 3 Minutes	9.4 %

Table 2:Reported effectiveness of extinguishers. (J.M. Shutske, 1994)

Rating	Percent
Very effective	15 %
Extinguished fire after delay	15%
Knocked down fire, but didn't extinguish.	67%
No effect	3%

These two tables show the relationship between detection time and the ability of the combine harvester operator to extinguish a fire by themselves. We can see that in most cases by the time the harvester operator was able to respond to the fire they had significant trouble extinguishing the fire. However, the data from this study is somewhat biased as the responders were initially identified by fire department reports. In most cases if the fire was successfully extinguished the operator may see no need to contact the fire department.

Of the 50 respondents, 72% had a fire extinguisher available to them and the overwhelming majority of these were dry chemical type extinguishers. One major issue with these extinguishers was capacity, in some cases the extinguisher was retarding the fire until it ran out of chemical upon which the fire reignited. The study recommended that every harvester should be fitted with at minimum one 4.55 kg (10lb) ABC dry chemical fire extinguisher and suggested that had second extinguisher should also be carried on board. (J.M. Shutske, 1994)

While a fully automated system detection and extinguishing system may appear to be the ultimate solution to the problem of combine fires, if only from a loss reduction perspective. But such a system is subject to certain limitations.

The system must include an automatic engine shut-off, testing by J.M Shutske indicates that extinguishing a fire may be very difficult while the engine is still running. The same study notes that a fully-automatic system would be a complex and expensive addition to the harvester, needing detectors, wiring, controllers and plumbing, storage tanks and pumps for the dispersal of the fire retardant. It states that a detector only system combined with the use of hand-held fire extinguishers would be just as effective while costing less than one quarter of the amount for the fully automated system. (J.M. Shutske, 1994)

A fully automated system enables the operator to be removed from the fire area, allowing them to extinguish the fire from the safety of the operator cab. A hand held fire extinguisher requires the operator to place themselves in close proximity to the fire, risking injury.

1.2 Project Aim

The broad aim of this project is *to attempt to develop a system that will detect fire on a combine harvester and will alert the operator so that they may take appropriate action to contain and extinguish the fire before it becomes too out of control for the operator to contain by themselves.*

1.3 Specific objectives

These objectives are ranked at two levels of importance. Objectives that use the phrase 'The system must…' are critically important to the success of the overall system. If these objectives are not met, then the ability of the system to meet the project aim is in doubt. Objectives that use the phrase 'The system should…' are performance parameters that would optimise the performance or usability of the system but would result in critical failure of the system if they are not met.

- The system must be able to perform its primary goal of identifying fire and the beginnings of fire in the engine compartment of the harvester. This includes abnormal hotspots, smoke and flame from the fire.
- The system must be able to differentiate between hotspots that occur during normal engine operation e.g. (A hot exhaust system) and hotspots that occur from fire so that the operator is not overwhelmed and distracted by false positives that may cause the operator to ignore future warnings.
- The system should be able to provide the operator with feedback regarding the type of fire situation that is developing in the engine compartment. This may let the operator determine what type of fire is developing, e.g. bearing failure, chaff fire diesel fire etc. and allow them to respond more appropriately to the situation.
- The system must be able to perform its primary goal under the harsh operating conditions that the harvester operates in during harvest. These conditions include high environmental background temperature, minimal visibility due to dust and particulate, high sunlight load and high air flow; all of these are environmental factors that occur under harvest conditions.
- The system must be able to alert the operator of the developing fire in time for said operator to have a 'reasonable chance' of effectively combating the fire. In this case the response time should be less than 30 seconds for an open flame and less than a minute for an abnormal hotspot. Any longer than this may well result in the fire developing beyond being controllable with what fire extinguishing equipment the operator has on hand.
- The system should require no operator input (No calibration or monitoring required) and must be able to operate with minimal operator input (Occasional calibration and/or monitoring). If the system requires too much attention it will distract the operator from driving the harvester, which may result in poorer harvest performance and at worst may cause an accident.
- The system should be able to be easily retrofitted to both new and old harvesters and should require a minimum of modification to the harvester during fitment.
- The system should be able to operate with minimal maintenance during the harvest period.
- The system must be able to operate for extremely long continuous periods. In an extreme case this could be up to 30+ hours.

2.Literature Review

2.1 Current Standards

There are currently no Australian Standards regarding fire safety and combine harvesters apart from the mounting of a fire extinguisher on the vehicle (AS/NZS 2153.7/1997) (Australian/New Zealand Standards, 1997), there is a small section regarding the safety of the machine operator and hot parts in the general standard regarding tractors(AS/NZS 2153.1/1997) (Australian/New Zealand Standards, 1997) but there is nothing regarding the prevention of hot parts contacting flammable materials.

There is a standard regarding automated fire protection systems for mobile and transportable equipment (AS 5062-2016) (Australian Standard, 2016). The standard specifics the requirements for the design, installation, commissioning and maintenance of fire protection systems for mobile and transportable equipment and it will be used as a guide for the development of any fire detection system produced by this project. Sections of note include

Section 2 and 3 of the standard deal with fire risk management and fire risk reduction respectively. Section 5 of the standard defines fire protection systems into four basic types. A fire alarm only type system that is capable of rapidly detecting and warning of an outbreak of fire but lacks the capacity to take any action to combat the fire. A manually operated fire suppression system, able to combat the fire but has no fire detection ability. It must also first be activated by the operator. A fire alarm system combined with a manual fire suppression, able to detect and combat an outbreak of fire but still requiring the operator to initiate the discharge of a fire suppression agent. Finally, a fully automated system that is able to combat any outbreak of fire without any human input. (Australian Standard, 2016)

The project will focus on creating a fire alarm only system that will rely on the harvester operator to take action to prevent any fires. The basic requirements for a fire alarm only system are:

a) Rapidly detect the outbreak of fire.
b) Initiate an alarm signal to allow manual safety functions.

(Australian Standard, 2016)

The rest of the standard deals with the implementation of a fire suppression system and is mostly not applicable to a fire alarm only type system.

2.2 Fire detection methods

Whatever type of fire detection sensor used it will have to contend with a difficult environment.

> *Engine compartments of heavy duty vehicles are, in general, spaces where detecting fires with inexpensive and simple detection systems is arduous. High air flows and large amounts of suspended pollutants in the compartment, together with the complicated geometry and the wide range of surface temperatures typically occurring during the normal operation of the vehicle, complicate the operation of all types of detectors. The deposition of pollutants on the components of optical detectors can impair their operation as well as obstruct the channels of aspirating systems, thus hindering their operation or shortening their service interval. In addition, thermal point detectors can have an extremely limited effectiveness under high air flow conditions unless these are located in the vicinity of an eventual fire where these can be effectively heated by the ensuing smoke and fire plumes.* (Brandt J., 2013)

The conditions for any system operating on a combine harvester will be working under even more difficult conditions than this. In addition to the challenges already mentioned there will large amounts of dust and crop residue swirling in and around the engine bay.

When a fire starts, regardless of the situation, it is characterized by several distinct physical and chemical manifestations. Theses phenomena include the flames, its size, colour, its movement, the smoke particles and the gases created during combustion. Finally, the fire creates a wide range of electro-magnetic radiation including visual light, infrared which manifests itself as the heat we feel on our skin and even more exotic types such as ultra-violet. (Brogue, 2013)

We can use all of these phenomena as mediums for detecting fire, however no method by itself can detect a fire with complete and dependable certainty. Because of this, most fire detection system use a combination of different sensors in order to increase the accuracy and reliability of the system. (Brogue, 2013)

The main types of sensors are particulate (smoke), light, heat and combustion gas detectors. Particulate detectors use the particulate matter, the smoke and associated products as a medium for detecting a fire. One of the cheapest and easiest methods uses a light source and a photodetector, when the particulate matter in the air reaches a point of saturation where the light can no longer be seen by the photodetector the alarm is tripped. An alternative method uses a very small amount of a radioisotope to emit alpha particle radiation into air flowing into the detector, ionizing the air and allowing a small current to flow through it. The presence of smoke in the air cause a disruption of this ionizing process causes the current flow to reduce, tripping the alarm. Both of these methods are commonly used in household fire detectors. (Brogue, 2013)

A light based system can not only just be used to detect the flame from an established fire but can also be used to detect the smoke in the air and the heat created depending on the exact configuration of the sensor. This type of sensor can recognize a flame by its colour and its movement using pattern recognition techniques. It can also look for smoke moving through the air and even the infrared (IR) and occasionally ultra-violet (UV) radiation produced. However, these systems are considerably more complex and sophisticated than other types and require a direct field of vision of the fire. (Brogue, 2013)

Heat type detectors can be broadly places into two different categories, point sensors and linear or distributed sensors. Point sensors use thermistors, as the temperature increases or decreases the resistance of the thermistor changes proportionally. They are very inexpensive but are quite inaccurate. Linear systems use a length of sensing cable to detect changes in temperature along its length using several different operating principles. Finally, gas detectors detect the gases from combustion, specifically the carbon monoxide created, as it is considered to be the only reliable gaseous indicator. (Brogue, 2013)

As previously described the engine bay of any heavy vehicle and particularly the engine bay of a combine harvester is an extremely difficult environment to detect fire in. The dust and particulate matter in the air will dramatically reduce the effectiveness of the smoke detection based systems as nearly all these systems detect fire based on how much particulate is in the air sample, it is not able to distinguish if the particulate is from a fire or if it from another source.

A visually based system will also have its effectiveness reduced but not to the extent that a particulate system. Systems that visually based can also utilise other methods of detection like IR or UV light that have the ability to 'see' through the dust where a visual light system would fail.

Figure 8: A CASE combine harvester harvesting lentils. Notice the extreme amount of dust created during the harvesting process. The object behind the harvester is a chaff cart and almost completely obscured by dust. (Quick, 2010)

Further complicating the particulate problem is the large air flow involved. Modern high performance diesel engine need large heavy duty cooling systems to operate optimally, inside the harvester itself there are large blowers used to lift lighter chaff and dust away from the heavier seeds. The complicated geometry inside the harvester makes it difficult to predict the airflow patterns and place the sensors in an optimal position for detection. Gas detectors are similarly handicap by the particulate and high air flows.

These two systems types, even if the smoke or gasses do make their way into the detectors will still have a significant delay as the smoke will need to build to a level sufficient to trip the alarm. As previously mention the time between detection and extinguishment is the critical factor in preventing the fire from becoming uncontrollable, even a delay of 30 seconds could be the different between the fire being a minor incident and the loss of the harvester.

A heat based sensor is unaffected by these environmental factors but still has its own issues that make detection of a fire unreliable in this situation. An engine bay is a hot place, the engine itself has an operating temperature of 100°C but that is quite cool when compared to the exhaust and especially the turbocharger, which can reach temperatures of 500°C or more. These parts present the most obvious fire risk but could confuse a simple heat sensor as it has no way of knowing if the heat is from a fire or just from a hot component. We could move the sensors away from these hot components but this would increase the delay in detection as the fire would have to move from its starting point towards the sensor in order to be detected.

Until recently, a system based on visual detection of fire using computer vision and video cameras was pure fantasy, the cameras were too fragile and expensive and the computational requirements were impractical for an embedded system. However, such a system would have numerous benefits and advantages over other types of systems. A camera is a volume type sensor, it is able to monitor a large area and it is able to identify exactly where a fire is occurring within its field of view as opposed to the other types of sensors which must wait until the smoke or heat diffuses enough to reach the sensor. (Byoung Chul Ko, 2010) Since the processing and detection occurs away from the camera, a single video feed can be analysed by numerous different methods, increasing the speed and accuracy of detection.

This system type has a much faster detection speed than both smoke and thermal based systems. A typical thermal sensor can have a response time of 20-80 seconds while a Near Infrared(NIR) optical sensor can have response times as short as 350ms. (Y. Le Maoult, 2007)

Finally, in the event of an alarm, the operator can simply look at the video feed and rapidly determine if the alarm is genuine rather than having to investigate the area in person.

2.3 Camera based detectors

For a fire detection system using digital video as the means of detection, there are two different types of camera sensor to choose from. The Charged Couple Device (CCD) or Complementary Metal Oxide Semiconductor (CMOS). CCD sensors were considered to be the superior type of digital camera sensor for many years, but CMOS sensors are now reaching a level of development and sophistication where they can no longer be relegated to second place.

A CCD camera sensor consists of closely spaced electrodes that are separated by a thin oxide layer from a semi-conductive substrate. When a voltage is applied to the electrode, a depleted region is formed beneath it in the semiconductor layer; this region is often described as a ‘well’ or ‘bucket’ and corresponds to one pixel in the processed image. When the CCD sensor is exposed to electro-magnetic radiation, say normal visual light, the ‘well’ fills up with the light photons impacting it. These ‘wells’ can then be emptied by pulsing the voltage through the electrode. (Waltham, 2013)

In a CMOS sensor each pixel has its own photodiode and readout transistor along with ancillary electronics that address, array and buffer the analogue video signal. In most cases the processing of the signal from analogue to digital is done on the pixel as well. (Waltham, 2013)

The fundamental different between the two types is that a CCD sensor physically captures the light photons and processing into an electrical signal, a CMOS sensor simply reads the light falling upon the photo diode sensor and creates its own signal. So which one to use? CCD is the older, more mature technology but CMOS is able to exploit the same advancements in silicon chip technology as other integrated circuits, it’s use in mobile phones and digital cameras has meant that CMOS sensors have undergone much more R&D in the past decade that CCD sensors have. Now we have a situation where CMOS sensors are now equalling and even surpassing CCD sensors in metrics where the CCD was considered superior.

The very latest CMOS sensors are able to outperform similar CCD sensors, on almost all fronts. The image quality is better, its sensitivity is comparable, a CMOS sensor is lower voltage, cheaper and less bulky; they are also not subject to blooming, where bright lights cause oversaturation of pixels, in the same way CCD sensors are. (Dempster, 2014)

When we are talking about the sensitivity of a camera there are two important parameters to understand, Quantum Efficiency (QE) and Read Noise (RE). Quantum efficiency is the measure of how efficiently the sensor converts the light photons into charge electrons. The higher the output level of the sensor, the more sensitive the sensor is to that particular wavelength of light. A QE of 1 indicates that every light photon generates one electron. (Adimec, 2015)

Read noise is the equivalent noise level (in electron RMS) at the output of the sensor in the dark. The lower the noise level, the lower the minimum number of signal electrons that can be detected. Combining these two gives the overall sensitivity of the sensor as QE/RN or the minimum amount of light that can been seen by the camera. (Adimec, 2015)

For the purposes of this project, one of the more important factors is sensitivity to spectral bands other than visible light, namely infrared and ultra-violet. CCD cameras can be used to detect different spectral bands: UV (0.25-0.39µm), visible light (0.39-0.75µm) and Near Infrared (NIR) (0.75-1.1µm). (Y. Le Maoult, 2007) CMOS sensors can also be used to detect these bands.

The latest industrial CMOS sensors are able to dramatically outperform industrial CCD in visible light sensitivity. But when we are looking at the NIR spectral band the gap starts to close. (Adimec, 2015)

But the latest and greatest has a large price tag attached and the potential for the sensors to get damaged or destroyed cannot be ignored. To get a more realistic picture of how cheaper, more easily available and more disposable sensors might perform it is worth while looking at how the older generation of CMOS and CCD sensor compare to each other.

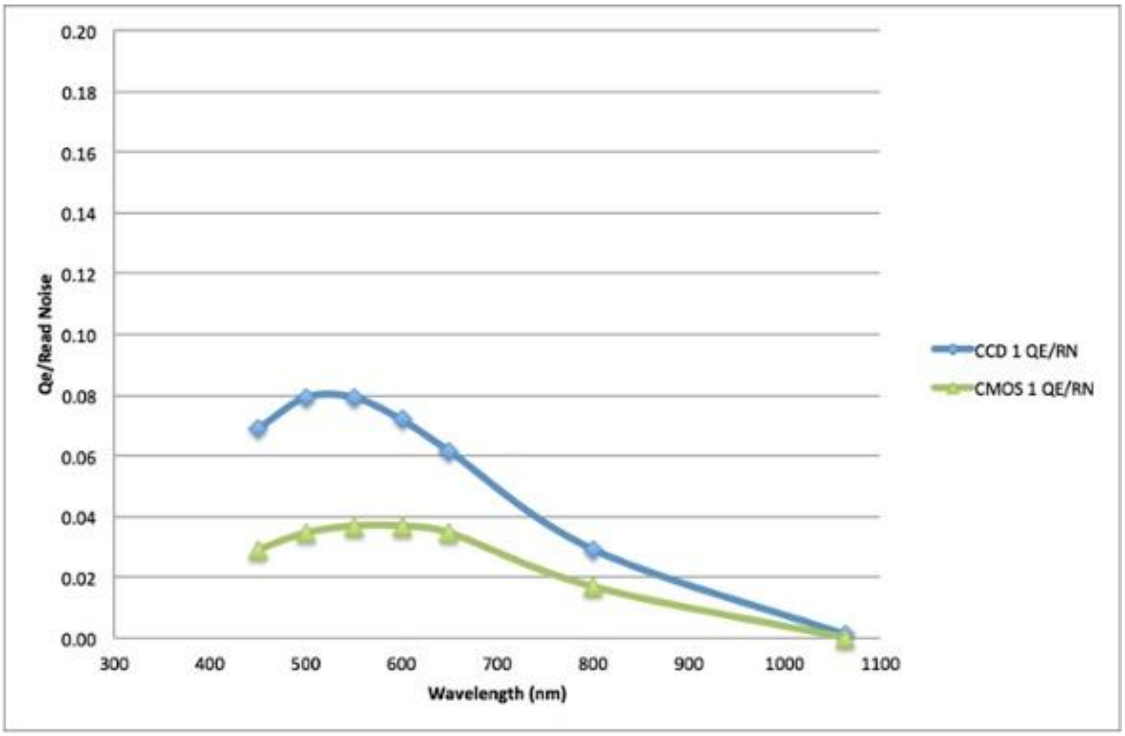

Figure 9: CDD vs CMOS in 2011. (Adimec, 2015)

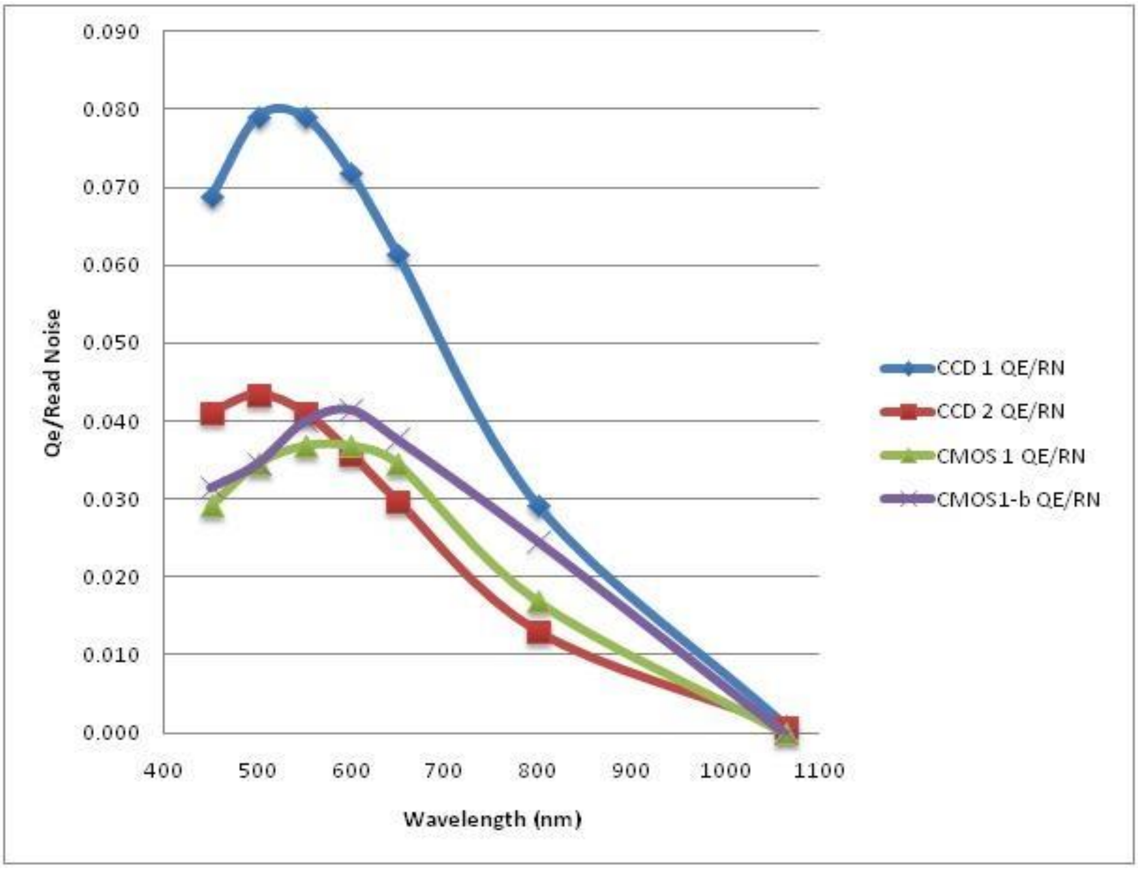

Figure 10: Various CCD and CMOS sensors compared in 2011. The CMOS1-b sensor has been specially made to be more sensitive in the NIR band. (Adimec, 2011)

In cheaper equipment we can expect to see a CCD sensor outperforming a CMOS sensor most of the time. But we can also expect this to change within the next few years.

Cameras that can detect the entire IR spectrum are known as wideband sensors, but these sensors are expensive, fragile and often require external cooling. Any system based on these sensors would need to be well protected increasing weight and cost further. (Y. Le Maoult, 2007). In this case that extra range is probably not necessary.

The detection of hotspots for all three bands is based upon the emission of blackbody radiation from the object in question. Of the three bands, NIR offers the most potential but both UV and visible light based systems could be able to detect hotspots under the right conditions. The minimum temperature that can be detected using this type of system around 350 degrees Celsius (Drysdale.D, 1996). While this means that the system should for the most part be insensitive to normal temperature changes even during an Australian summer but could produce false positives when viewing the hot exhaust and engine.

In the UV, visible and infrared spectrum, fires show distinct and well known set of emission bands. These bands are the primary method that is used by detectors to sense the fire. A non-exhaustive list of the detection criteria could include

- Energy threshold on a single spectral band or on several ones;
- Ratio of energy for two different spectral bands: the typical infrared signatures of a fire. As fire is also a dynamic phenomenon, a temporal criterion can be added:
- The flickering analysis of energy in a spectral band due to the ‘puffing’ frequency of the fire

A combination of these criteria could be used to avoid false positives, increase the accuracy and speed of detection. (Y. Le Maoult, 2007)

In the UV spectral band, a CCD camera is capable of detecting hotspots above 350 degrees' Celsius blackbody temperature but this requires a much larger number of photons (1.65 x10^10) to hit the detector and has a correspondingly longer detection time because of this. Depending on the type of fire it can be able to detect flame but it is insensitive to smoke. There is also the distinct possibility of interference from solar radiation and manmade sources. (Y. Le Maoult, 2007)

In the visible spectral band, a CCD sensor cannot detect hotspots until they reach a much higher temperature of 500 degrees' Celsius blackbody temperature. But unlike both UV and NIR it can detect this with a minute emission (24 photons). Flame Detection is based on the flickering of the flame and emissions in the CH Bands. It can also detect smoke by opacity measurement with a reflecting target and visible light source. However, the ability of a visual light sensor to see is severely restricted by dusts and particulate matter in the air. (Y. Le Maoult, 2007)

The NIR band (0.75-1.1µm) allowed the detection of hotspots down to 350 degrees' Celsius blackbody temperature but required around half the emission required by UV (1.813 x 10^5 photons). When a NIR filter of 950nm ±100nm was used the detection temperature went up to 410 degrees' Celsius blackbody temperature for the same level of emission. (Placeholder1)

Flame detection is based on the same principles as those used for UV and visible light, looking for the distinct emission bands and flickering mode of the flame. Using a 950nm filter produces a weaker signal.

Smoke detection works similarly to the method used for visible light except the light source is replaced by one emitting in the NIR spectrum. (Y. Le Maoult, 2007)

For optimal performance from a system using this type of sensor, it will most likely need to used multiple methods of detection, probably flame and hotspot detection as smoke detection may be extremely difficult to implement considering the amount of dust and particulate matter already in the air during normal operation. The system will also need to use different spectral bands, in this case NIR and visible light. UV light has limited usefulness considering the amount of interference that solar radiation and other sources will cause.

2.4 Signal Processing

In order for a fire detection system to be considered a true 'set and forget' type system, it needs to have the ability to identify a fire within its sensor range. With a system that is based on optical video, we need to use machine vision technology that is able to recognize the flame based one or more different indicators. This system needs to work with a minimum of delay in order to raise the alarm as soon as possible.

The hardware package must minimal in size, rugged enough to withstand the conditions it will be placed in and must still have enough processing power to run the signal processing software. More sophisticated systems use special purpose systems with integrated digital signal processors, video encoders/decoders and communication modules (Xuejun Chen, 2015)

In a study done by (B. Uğur Töreyina, 2006), the process used a video-based fire detection algorithm looking for motion and colour cues, flame and fire flicker, quasi-periodic behaviour in the flame boundaries and irregularity of the boundary of the fire region. Their system used four basic steps to determine if any of these conditions had been met.

First the algorithm looked for moving pixels or regions in the current video frame, this was done using a simple hybrid background estimation method that compared the intensity of the current pixel to that of a background model that was taken on the first frame. If the intensity was more than a threshold value it is considered to be in the foreground and to be moving. They chose this method because of its computational efficiently. (B. Uğur Töreyina, 2006)

Next the colours of the moving pixels are checked to see if the match pre-determined fire-colours. These colours were first determined using a mixture of Gaussian distribution models made from sample images containing fires. If the pixel being processed lies within the standard deviation of the model, it is considered to be fire coloured and marked as so. (B. Uğur Töreyina, 2006)

Finally, the algorithm does a wavelet analysis of the moving regions in temporal and spatial domains. The temporal wavelet analysis is looking for the repeated flicker or oscillation of the pixels as the flame moves around in the video frame while the spatial wavelet analysis is searching for the repeated oscillation in the colour of flame. The idea behind these two steps is to attempt to filter out objects that are fire coloured and moving. For example, a person wearing a red coloured shirt that walked past a fire detection system that only looked for movement and colour would case a false alarm. That person's movement is much more consistent than the flame and the coloured shirt has a much more consistent red colour than a flame would have. By using the two filtering steps it is possible to prevent such a false alarm from occurring. (Y. Le Maoult, 2007)

This team was using video in normal visible light wavelengths to detect the flame. Since the process relies on determining the colour of the moving region before moving onto wavelet analysis, its usefulness in the NIR region may be limited. In addition, they used many heuristic thresholds making them impractical in real-life applications as the results will vary depending on the input, as noted in two studies by Byoung Chul Ko, Kwang-Ho Cheong and Jae-Yeal Nam. (Byoung Chul Ko, 2009) (Byoung Chul Ko, 2010)

This team has proposed two methods of signal processing for fire detection. In the first study the team used a process of detecting the fire by first looking for pixels that were fire coloured, they noted that what is considered to be 'fire coloured' changes depending on the environment and the fire fuel. So instead they generated a RGB probability model using a unimodal Gaussian from sample pictures which was then used to detect fire pixels. After the pixels are determined to be candidate fire pixels, the system removed non-fire pixels by comparing and analysing the difference between two consecutive frames. If the pixel didn't move enough between frames it was classified as a non-fire pixel and combined with the background, otherwise it was declared to be a fire pixel. (Byoung Chul Ko, 2009)

Using just these two steps makes it difficult for the system to distinguish between what is an actual fire and a moving object with a fire colour. In this study they used temporal luminance variation as the third step to remove additional non-fire pixels, the idea being that a fire region will have higher luminance contrast over several frames than a non-fire region. Finally, the system used a support vector machine (SVM) to classify which regions were fire or non-fire within the picture frame. (Byoung Chul Ko, 2009)

When this program was tested on various test videos containing fires and other moving regions the overall detection rate was 86.5% with it being a true positive 86.1% of the time overall. In comparison the Töreyina method had an overall detection rate of 71.3% with it being a true positive only 66.4% of the time. (Byoung Chul Ko, 2009) (B. Uğur Töreyina, 2006)

A second study was conducted in 2010. The process for this program inverted the first two steps, the movement in the video feed was detected first and then it was determined if the pixel was fire coloured or not. The movement detection used a very simple adaptive background subtraction model to separate the foreground from the background in the frame by comparing the intensity value of each pixel with a background model created on the first frame of the video, if it is above an adaptive threshold value it is considered to be the moving foreground. The same RGB probability model was used in this program to determine if the pixel was fire coloured or not. (Byoung Chul Ko, 2010)

The last step of this model used a completely different process, the third step and final step used a hierarchical Bayesian network to verify if a pixel was in fact a real fire pixel. While somewhat more complicated than the previous methods it shows more promise as a real world model. In this study the overall detection rate was 95.7% with the overall true detection rate being 95.3%. (Byoung Chul Ko, 2010). It should be noted that for the Töreyina method and both of the Byoung Chul Ko methods the same footage was used in testing.

A different approach was taken by Turgay Celik and his team in 2009. Their proposal was to use YCbCr colour space instead of a RGB colour space. By doing this it allowed them to separate the luminance of the image from the chrominance more effectively.

Figure 11: The original RGB image in a) and it's Y,Cb and Cr channels in b),c) and d) respectively. (Turgay Celik, 2009)

In the RGB colour space, the intensity of the pixel cannot be separated from the chrominance of the pixel, the chrominance being used to model the colour of the pixel, by changing the image into an YCbCr colour space, the image can be more easily analysed. The study claimed to have a detection rate of 99% and false alarm rate of 31.5%. Unfortunately, they used different studies and footage to test and compare results. (Turgay Celik, 2009)

In a similar vein (Y. Le Maoult, 2007) used the variation of the height of the flame, the movement of the gravity centre of the flame and the mean NIR energy on the surface of the flame. The use of a NIR camera meant that the object had to be emitting large amounts of thermal radiation to be seen, reducing the detection of false positives. Since the study had information regarding the accuracy of the system it is impossible to say how effective a NIR only system is.

The system must have the ability to distinguish between normally hot objects e.g. A hot exhaust muffler or the sun. Since these objects stay very static in comparison to a fire then software will need to be 'tuned' to look for rapidly changing dynamic ques. I believe by using multiple cameras that can detect both visible light and NIR the system can cross reference and compare both sources of information, drastically increasing the accuracy of detection.

2.5 Conclusion

Surmising what has been discussed during the literature review

- There is an applicable standard for automated fire systems (AS5062-2016)
- The engine compartment of a combine harvester is an extremely difficult environment to reliably detect fire in.
- When a fire starts it is characterized by several distinct phenomena including light, heat, flame, smoke and combustion gases. These can all be used as mediums for fire detection.
- Most fire detectors require a build-up of smoke or gases in order to trip the alarm. Camera based detectors require a line of sight to the fire.
- Camera based alarm systems can use either CCD or CMOS digital camera sensors, CCD sensors are quickly being outpaced by CMOS sensors.
- Both CCD and CMOS sensors can detect UV, visual and NIR light with varying sensitivity.
- Machine vision programs can detect fire with some success.
- Machine vision looks for colour and movement ques to detect fire.
- The accuracy of detection is increase by using more sophisticated filtering techniques.
- Regardless of how sophisticated the system is, a human operator is still needed to confirm if a fire has been detected.

Based on this I believe that it is quite possible to detect fire in the engine bay of the harvester, this system does not need to be overly expensive or complex. By combining a camera or cameras with the ability to see both NIR and visual light, and a machine vision process to automatically distinguish fire or the precursors to a fire, a harvester fire might be detected much more quickly that it is currently.

By giving the operator of the machine early warning they stand a much better chance of being able to put out the fire before it becomes uncontrollable by a single person armed with a fire extinguisher.

3.Project Methodology

3.1 Fire System Design

3.1.1 A hypothetical retrofit

A farmer has recently purchased a second hand harvester and wishes to retrofit a fire protection system to safeguard his investment. The harvester is a 2012 model with 3500 engine hours and has cost the farmer $250,000 AUD. Under the farmer's current farm insurance, he is able to claim up $10,000 of fire prevention and extinguishment. He is willing to use his harvester as a testbed for our fire prevention system under certain conditions.

- There are no major permanent modifications to the harvester.
- The system does not interrupt the normal running of the harvester.

Figure 12: The farmer's new harvester

3.1.2 The system

The fire protection system will use the following components.

- Multiple cameras placed around the harvester.
- A central signal processing unit.
- A user interface
- Associated wiring, mounting brackets and other miscellaneous components.

3.1.3 Cameras

The cameras will be used to sense two different types of light, normal visual light and NIR. A third type of light, UV, could also be used to detect fire but could have interference from normal sunlight.

The primary constraints of the cameras are size, image quality and cost. The cameras need to be small enough so that they can be fitted into tight spaces in and around the harvester for optimal coverage of fire starting trouble spots. These cameras need to have good enough image quality that the fire detection algorithm or indeed a person can make out a fire. Finally, since the system will need multiple cameras, the cost of each individual unit should be as low as practically possible. Indeed, there is a large possibility that the cameras will be destroyed or damaged in the event of a fire.

As discussed in the literature review CMOS sensor cameras have caught up to the image quality of CCD sensors and in the near future it is probably that they will surpass them. But for now it tough choice to pick between the two. Both types have slight advantages over one another but there is no clear-cut winner. The low cost CCD cameras seem to have an increased sensitivity to NIR, hence their use in low-cost night vision system that use NIR floodlights and an unshielded sensor. So for now they will be used in the NIR part of the system.

As for the colour cameras in the system a cheap CCD provides better colour information than a cheap CMOS camera but it is arguable that this extra colour information would be useful. A system that used colour cues to detect fire may have better accuracy with better colour sensitivity, it’s hard to know without testing.

Figure 13: The type of CCD camera that will be used for this fire detection system. The camera is roughly the same size as a 50 cent coin.

Based on this the system will use CCD sensor cameras for both the NIR and colour cameras. The cost difference is negligible, there are numerous cameras of both types on the market. Cameras with SD TV quality range in price from 20 to 100 USD and should be perfectly adequate for this purpose. The reason for not using a higher resolution camera is due to the increase in processing power needed.

Increasing the resolution of the cameras results in an exponential increase in the amount of data that must be processed by the fire detection algorithm. In a PAL standard system, the horizontal resolution is 720 pixels and the vertical resolution is 576 pixels. Each frame therefore has 414,720 pixels. If the system has 24-bit colour (the system is able to recognise and reproduce approximately 16 million different shades) then the data per frame is 9,953,280 bits. For live video at 25 frames per second the final bit rate is 289.56 megabits per second (Mb/s). This data rate would fill a DVD (4.7GB) in 127 seconds.

If the system was to use High Definition video the frame size is increased to 1920 horizontal pixels and 1080 vertical pixels. The resultant frame has a total of 2,073,600 pixels and with the same colour depth and frame rate as before the resultant video has a bit rate of 1493.03 Mb/s. Using our same DVD this time it would be filled with data in 25 seconds.

There is no real need for this increase in bandwidth as the sensitivity of the camera to light is a much more important factor in detecting the fire. All the extra resolution does is increase the processing requirements placed on the algorithm and the overall cost of the system.

The cameras will be mounted inside a rugged housing designed to protect the camera sensor from the elements. The housing should meet IP67 standard, meaning that the housing should be dust proof and able to be immersed in water up to one meter. Each camera will require a 12v electrical connection, readily available from the harvester's electrical system and a wired RCA for video output to the processor unit. Wireless transmission of the video is a possibility but the position of the cameras inside a complex metal structure may well make transmission difficult.

3.1.4 Processor and user feedback.

The central signal processing unit is the core piece of the fire detection system. It runs the algorithms that will detect a potential fire in the video feed. There are two ways that the processor can be constructed either as device similar to a normal personal computer or as a specialised embedded Digital Signal Processor (DSP).

The difficulty in using a DSP comes into the initial development of the code for running on the embedded system. Personal computers tend to be complex instruction set computing (CISC) where a single instruction can execute several low level operations (loading and storing from memory for example). DSPs and other embedded chips use reduced instruction set computing (RISC), each instruction only performs one action. This simplified instruction set is designed to result in an increase in performance for a properly written program. But a program written for a CISC platform may not run at all on a RISC based platform, let alone run optimally.

While it is perfectly acceptable to use a normal pc for the initial construction and testing of the fire detection algorithm, for the completed product a DSP provides significant advantages in cost, power usage and size. Depending on how well the algorithm is optimised it may well also be much faster than a PC based solution.

Several DSPs running in parallel could be integrated into a signal unit with another processor controlling the overall system, giving a DSP with a potential fire detection priority over the rest of the units. It processor should be mounted in the cabin, protected from the elements.

Once the signal is processed there needs to be some way of alerting the operator to a fire or even to just provide information from the system. Some new harvesters come fitted with cameras in the grain tank, rear panel and unloading auger and utilise a video screen inside the cab to display the footage from the cameras to the operator. (Deere & Company, 2016)

This screen would provide an ideal way to display video feed from the fire system to the operator. In the event of a fire the system could override the existing video feed and automatically switch to the cameras that had detected a potential fire. The system would also incorporate some sort of audio warning such as a buzzer to attract the attention of the operator. He could immediately determine if the fire alert was genuine and take steps to contain it.

Figure 14: An observation screen inside the cab of a harvester. This screen would be an idea way of providing visual feeds from the cameras and warning alarms. (Deere & Company, 2016)

3.1.5 System sensor mounting

In order for the system to be effective, the sensors need to be mounted in positions where they can view fire trouble spots. The most important area to have complete coverage is the engine bay. To provide this complete coverage the engine bay would require multiple sensors. These sensors could be mounted in the engine bay, on an overlooking the engine bay or even on an Unmanned Aerial Vehicle (UAV).

Sensors mounted in the engine bay would be able to directly observe common sources of ignition, such as the turbocharger or exhaust system, and would be able to immediately identify any potential fire. Any sensor in this position would need to be designed to deal with heat, oil, dust and crop matter that collects in the engine bay during operation. The ability of the sensor to detect fire may well be reduced or even nullified if these environmental factors are allowed to build to a critical level. (Brandt J., 2013) Additional sensors mounted in other areas would also have to contend with these environmental factors, but not the same extent.

As such these cameras may need an addition system to keep the sensor clear of debris, one example would be a compressed air nozzle that produced a blast of air periodically to clear the build-up of chaff or dust on the sensor. The system could also incorporate a degreasing fluid to remove oil as well. Of course any such cleaning system would add additional complexity and cost to the fire detection system.

Mounting the camera in a position overlooking the engine bay moves the camera away from the worst of the environmental factors, it also allows one set of cameras to see the engine bay.

This type of mounting would also reduce the field of view of the cameras. Because they cameras would not be in close proximity to potential sources of ignition it may take longer for the fire to be detected. This type of mounting is also highly dependent on the configuration of the harvester. In some cases, the cameras could be mounted on the grain tank making this type of mounting very easy and convenient. But in most cases this would not provide a sufficient view of the engine bay, making the use of a mast necessary. It also introduces the risk of electrocution as the mast may hit overhead powerlines. (Worksafe Queensland, 2015)

3.1.6 Unmanned Aerial Vehicles

Using a UAV moves the sensor package out of the dust and away from the machine completely enabling independent operation in a much cleaner environment. The UAV can move independently of the harvester allowing the sensor package to get a better view of the engine bay. In the event of a sensor failure it provides redundancy, being able to cover the area of multiple sensors. The wide field of view would also allow the sensor package to detect fires that start outside of the engine bay or fires that have been started in the paddock behind the harvester.

Of course using a UAV introduces a number of challenges including, keeping the UAV aloft long enough for the sensor package to be able to monitor the harvester, avoiding obstacles and legal issues relating to flying UAVs.

In this arena the main constraints are payload and endurance. For the platform to be useful it must be able to mount one or more cameras on a stabilised gimbal mount with the ability to transmit live video back to the system processor. It must also have the endurance to be able to fly for long periods without recharging or refuelling, so that it can provide uninterrupted coverage of the harvester. The UAV platform must be able to meet these to requirements regardless of the prevailing weather conditions.

As an aerial vehicle it comes under the regulatory authority of the Civil Aviation Safety Authority (CASA) and their requirements for compliance will greatly depend on how the system will eventually operate. Unless the harvester driver wishes to undertake a remotely piloted aircraft operator's certificate (ReOC) or a remote pilot licence (RePL), the completed drone must weigh **less than 2 kg.** The standard operation of the drone must also

- Be conducted within a visual line of sight.
- Must be kept at least 30 meters away from people.
- Not flown any higher that 120 meters AGL
- Not flown within 5.5km of controlled airspace.
- Not flown in a manner that creates a hazard to people, property or other aircraft. (Civil Aviation Safety Authority, 2016)

UAV Platform types

There are three different types of UAV that could potentially be used for this application

- Fixed wing airplane
- Single rotor helicopter
- Multi-rotor helicopter (Quadcopter/Octocopter)

The airplane type uses has one or more fixed aerofoils to provide lift and an engine or engines to provide forward thrust. Of all the different types of UAV platforms it is the one with the best potential for endurance, range and payload. However, it comes with several disadvantages including needing space and possibly assistance for take-off and landing as the wing has a minimum speed before it starts to generate lift. The UAV has to either get to speed on the ground using a runway or has to be thrown by the human operator. Related to this is the minimum speed that the UAV can fly at. Harvester operation is done at a slow ground speed and the UAV would need to match this speed to provide constant surveillance.

This could be achieved by using a long 'sail-plane' type wing designed to provide lots of lift at these slow speeds but they are easily damaged and more subject to wind. A small cross-wind could easily blow the UAV off course away from the harvester.

Figure 15: A large airplane type drone used in surveying. (Joe, 2010)

A single rotor helicopter solves some of these issues, it can take off and land vertically but the large single rotor is still subject to the wind. It also has reduced range, endurance and payload in comparison to a similar capacity airplane type.

Figure 16: A Yamaha R-max. A helicopter type UAV used in agricultural applications. (Gtuav, 2014)

This leaves the multi-rotor helicopter. Instead of using one large rotor, these use multiple rotors, commonly four or eight mounted at the extremes of the frame. By using multiple small blades, a multicomputer is able to fly in windy conditions that would carry the other two types away. It is also able to take off and land vertically like a helicopter. But this comes with a severe penalty in payload and endurance.

Figure 17: A typical multi-rotor UAV. This UAV is carrying a gimbal mounted camera for aerial photography. (Glinz, 2013)

Despite these disadvantages, if a UAV was to be used with this system it would be of a multi-rotor type as the need for vertical take-off and landing, and the need for flying in windy conditions means that fixed wing or single rotor type would be rendered useless despite their endurance and payload advantages.

Payload

Based on our previous design, the sensor package that would be carried by the UAV would consist of the two cameras, one visual and one NIR. While the cameras are not in themselves heavy, the camera body needed to protect them will add some weight. For these cameras to be any use they will need to be mounted on a stabilised gimbal mount, so that the harvester can be seen by the cameras regardless of how the UAV is manoeuvring.

The gimbal mount is also one of the heaviest if not the heaviest part of the payload. It has multiple parts and motors that add to the complexity and weight of the overall UAV design. The video from these cameras will also need to be transmitted back to the signal processing unit in the harvester.

One of the primary uses of commercial UAVs to date has been aerial photography. Much of the equipment that is used can be quite ready adapted for our purposes. One particular area of interest is First Person View (FPV) racing. Using small but very high powered quadcopter drones, they are raced against the clock and each other through various courses. A small camera mounted on the front of the drone transmits live video back to the operator who uses video googles to provide a literal first person view.

While we certainly have no need for a first person view of the fire, the live video transmission equipment could easily be adapted for our much more mundane use. The most important factor for payload is weight, any weight saving that can be made will allow an increase in endurance as this saved weight can be dedicated to extra battery capacity.

A rough weight budget could be

Cameras: 15g for CCD of TV quality

300g + for HD quality

X2 plus a NIR filter for one of the cameras.

Gimbal: 25g for basic pan and scan mount

120g for 3 axes stabilised gimbal with motors and controller.

Video transmission: 16g for 2.4ghz 16 channel video/audio TX

Total Payload: 56 to 436 grams

There is quite a variance in weight depending the quality and steadiness of the image that is desired.

Endurance

When we start to look at the endurance of the drone is where problems start to emerge. The vast majority of drones use electrical motors and lithium-ion batteries. Li-ion batteries are a marked improvement over older types of rechargeable batteries such as NiCad or NiMh, they have greater power density, low self-discharge and don’t develop a memory. But there is still one glaring issue.

Recharge time.

Charge V/cell	Capacity at cut-off voltage	Charge time	Capacity with full saturation
3.80	60%	120 min	~65%
3.90	70%	135 min	~75%
4.00	75%	150 min	~80%
4.10	80%	165 min	~90%
4.20	85%	180 min	100%

Figure 18: Charge characteristics of typical lithium-ion cells. (Buchmann, 2016)

A typical lithium-ion battery requires around 180 minutes to be fully recharged from a depleted state. In order to prolong the battery life, it is usually only charge to 85% of its maximum capacity. But note that the charge is not linear, the battery needs two hours to reach 60% of its capacity, but to get an extra 25% requires another hour. The trick will be finding the optimal balance between recharge time and charge capacity.

Most quadcopter UAVs have a flight time of 10 to 30 minutes at most. A UAV that spends the majority of its existence recharging instead of monitoring for fire not much use to us. Fully charging (beyond the typical 85%) the batteries all of the time will also reduce the lifespan of the batteries.

So why not charge the battery faster? All batteries have what is known as a C-rate which governs how fast a battery can charge and discharge is stored energy. For a battery rated at 1C, the fully charged battery rated at 1Ah could provide 1A for one hour. If we charge at rate above the batteries rated C-rate, then the lifespan of the battery will be reduced. If we charge the battery at a rate far beyond its specifications the battery may well be damaged or destroyed during recharging. In practice the c-rate for the charge and discharge of the battery differ with peak discharge rates of the battery being much higher than charge rates.

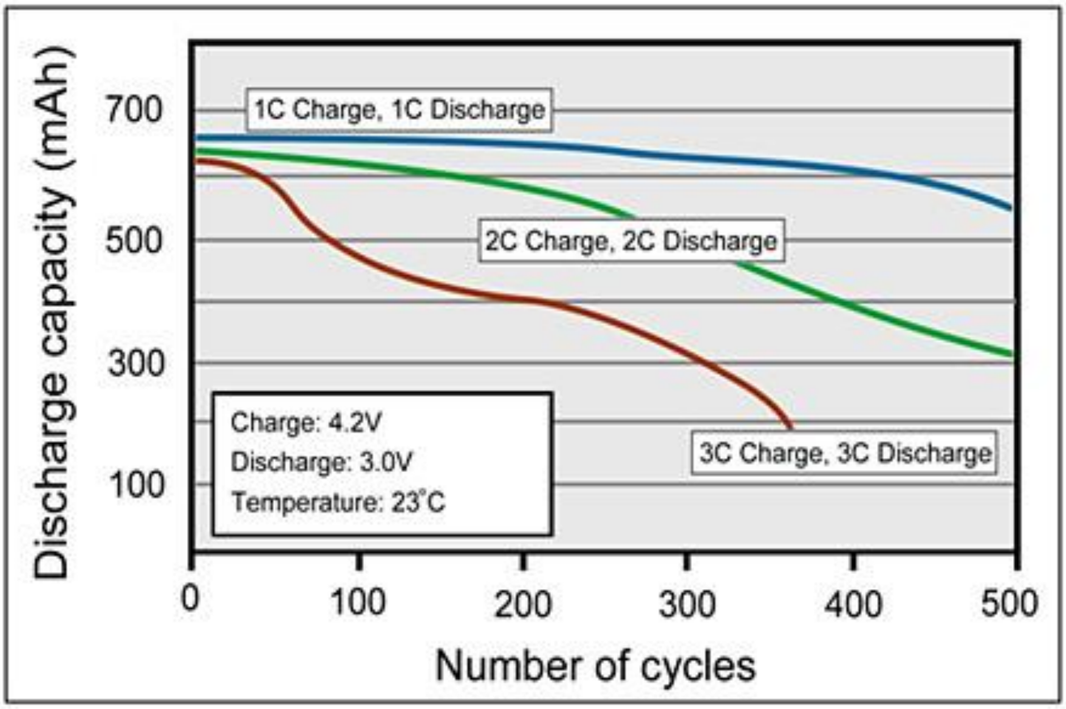

Figure 19: The effect of fast charging and discharging on a lithium-ion battery. (Buchmann, 2016)

Say the system was operational on a harvester and the system had two UAVs to keep complete coverage of the harvester. Each UAV has a 10Ah battery rated at 1C charge and discharge and has an optimistic flight time of 45 minutes. If our charger was 90% efficient the charge time on the battery would be 66 minutes at 1C, 36 minutes at 2C and 22.002 minutes at 3C.

If the work day of the harvester is around 12hours there will be 16 charge/discharge cycles during the day for each UAV. By 300 cycles or roughly two and a half weeks of continuous work the battery charged at 3C has lost over half of its discharge capacity and the 2C battery has lost a third. For this scenario this means that our battery is now down to 5Ah or 6.6Ah respectively.

The loss of capacity will impact the flight time of the drones meaning that the battery would have to be charged even faster before it could be properly replaced.

So what alternatives are out there to using lithium-ion batteries? Unfortunately, other types of rechargeable batteries are even worse off than li-ion. It might be possible to use primary (Non-rechargeable) batteries, but this creates a problem of supply and waste. In our modern sustainable world, it is unlikely that a system that could use over 122 disposable cells per week would be accepted by the consumer, never mind the cost of such cells.

Thinking further out of the box we could change to a petrol-electric system. A small petrol motor, similar to those used on model airplanes could be connected to a generator to supply power to the UAV.

One such proposal would give a flight time of around 60 minutes with a payload of 3kg. (airstier UG & Co.KG, 2016) Since our payload would only be around 1/6th of that we could use the rest of the payload for additional fuel. At this stage this design is still in the prototype stage. It is unlikely though that such a drone could be scaled down to a 2kg weight limit while still beating a lithium-ion system in range and capacity.

Since the purpose of the drone is to monitor the harvester during operation we could use the large electrical system on the harvester to our advantage. By tethering the UAV to the harvester using an umbilical cable to transmit power to the drone we can discarded any fuel system entirely.

The drone would need to have more sophisticated object avoidance in this setup as the umbilical tether could easily be tangled up in trees, power lines or machinery. The limiting factor is the weight of the cable; heavier cable would allow more current transmission enabling the use of more powerful motors which now have to lift a heavier cable. If CASA also decided that the umbilical cable was to be included in the overall weight of the UAV, it would almost certainly exceed the 2kg weight limit.

Proposals for drones using hydrogen fuel cells could have flight times in excess of 4 hours but at the time of writing these are all still in experimental stages of development.

Key Points

- The drone cannot have a weight of more than 2kg.
- The payload of the drone varies drastically with the quality of the video transmission.
- A petrol-electric or umbilical powered UAV is unlikely to meet the weight limit.
- To provide uninterrupted coverage multiple drones will be needed.
- Lithium-ion batteries degrade quickly when they are fast charged.
- A balance between charge time and charge capacity needs to be found in order to optimise the flight/recharge cycle.

3.1.7 UAV Flight Software

Apart from the software algorithm being used to detect the fire. The UAV itself will need sophisticated software to meet its performance objects. The UAV will need to be able to perform an automatic take-off and landing including a return to base function when the battery cells are nearing depletion. The drone will also need to have collision avoidance software to prevent crashes into paddock obstacles such as trees, powerlines and alike. Most critically it needs to be able to follow the harvester around the paddock and hold a position relative to the moving harvester.

This last task is known as a 'follow me mode' and is starting to become common on drones designed for the adventure market. These drones use an electronic device, most commonly a smart phone to provide a trackable moving point that the UAV can orientate itself around. This software solution could easily be transferred to tracking the harvester as it moved around the paddock.

While many off the shelf solutions exist for the problems described. Not too many have all the abilities needed wrapped into one package, those that do are proprietary solutions.

3.2 Project Software

3.2.1 Goals

The software is the vital component of the system. Without it, it is no different to any other video camera system. Based on the literature review, this project will be using cameras able to see both the visual and the NIR band. The software program will be tailored to detect fire based on the video information from these spectrums.

The program will need to reach a balance between processing speed and accuracy. The program must be speedy enough to process a live video stream. Preferably this processing should be done on every frame of the video output, giving an overall framerate of 24-30fps depending on the video camera used. As a compromise if the program is unable to process every frame, it could be possible to process every second frame (12-15fps). This would put the system in line with most CCTV systems.

Having said that, the majority of the testing will be done in a post processing fashion. By using post-processing, it allows testing of the hardware side of the system to be carried our independently of the software system. By recording the testing, we can test the fire detection program as it is developed and make changes on the fly rather than being forced to work sub optimal programing. It also keeps a record of testing that can be referred to later.

For the system to be considered reasonably successful in detecting fire, it should have an accuracy comparable to other types of fire detection systems. For example, the two types of commonly used household detectors, ionization and photo-electric have an accuracy of 45% and 96% respectively. Taking into account the much more difficult environment, I believe that an accuracy of 50% or better is a reasonable and achievable goal for the system.

3.2.2 Raw Video Feeds

The raw video supplied from the cameras is in an analogue format, before any algorithm processing can take place, the signal must first be converted into a digital format. This analogue to digital signal conversion takes place on the video-capture device before it reaches the processor proper. When the video feed reaches the processor it has a resolution of 720x576 pixels, the same as the Phase Alternating Line (PAL) TV format.

The colour format is somewhat more confusing as changes as the signal works its way from the camera to the video-capture device and finally into the processing software. The cameras record colour as a gamma corrected red, green and blue values (R' G' B'). These values are then used to calculate the intensity value Y':

$$Y' = 0.299R' + 0.586G' + 0.144B' \tag{1}$$

The R', G', B', Y' signals are then used to create the colour-difference signals (B-Y) and (R-Y) known as U and V respectively.

$$\begin{aligned} U &= B' - Y' = -0.299R' - 0.587G' + 0.899B' \\ V &= R' - Y' = 0.701R' - 0.587G' - 0.144B' \end{aligned} \tag{2}$$

For the PAL output signal, the U and V signals are combined into a single Chroma signal (C) for transmission.

$$C = U \sin \omega t \pm V \cos \omega t \tag{3}$$

Where $\omega = 2\pi f_{sc}$ and f_{sc} is the frequency of the colour subcarrier, for the PAL format this is 4.43MHz. (Margques, 2011)

Once this signal reaches the video capture device it is converted into either a UYVY or YUY2 colour space. Both of these are 16-bit colour spaces with 8 bits for black and white and 8 bits for the colour signal. In the processing software they are both treated as YUV video data. In this colour space the Y' signal is the luminance component, U and V are the Chroma colour difference signals as previously described. The difference is that now these video feeds are digital rather than analogue signals. Finally, the processing software will convert these values back into the RGB colour scheme that we started with. The convoluted method of handling colour in the video signal is due to the use of many different off the shelf parts.

The matter is made even more confusing by the NIR camera. The raw video feed appears to be a greyscale image indicating that only the luminance signal is being transmitted. But this is in fact not the case at all. The camera has had NIR pass filter added to the camera assembly, the particular wavelengths that the filter allows through appears to be a greyscale image. But the cameras sensor itself has not been modified and still supplies colour information.

3.2.3 Method of detections

To detect the fire, the algorithm will use three criteria to identify a fire in the video feeds. These are:

- The colour of the flame, defined as the **Colour** method of detection. This method will only be used on the visual light camera.
- The thermal emissions of the fire, defined as the **NIR emissions** method of detection. This method will only be used on the Near Infrared Camera.
- The motion of the flame as the reaction takes place, defined as the **Movement** method of detection. This method of detection will be used on both cameras.

The algorithm will run each of these methods of detection independently. Since the movement method of detection will be taking place on both video feeds, the program will be duplicated twice within the algorithm.

3.2.4 Colour fire detection

For the Colour fire detection program, I am using the theory from *Fire detection in video sequences using a generic colour model* by Turgay Celik and Hansan Demirel. This article was mention in the literature review and I believe that this method will be one of most effective ways of detecting fire in the colour video feed. I am interested to see how this theory can be applied to this problem with success.

Figure 20: A still image that will be used to explain the colour processing methodology. This image was chosen due to the large amounts of colour contrast between the red of the fire, the green and yellow of the harvester and the blue of the sky. The picture contents are also relevant for obvious reasons. (Strangefarmer.com, 2016)

Colour modelling in digital systems.

To understand how we can detect fire using colour information, first an understanding of how colour is handled by our processor and software algorithm is needed. Once we have the signal from the video cameras in our algorithm the colour region can be interpreted in a number of different ways. These ways are known as colour spaces (also called colour models or colour systems) are a specification of a co-ordinate system and within a subspace where each colour is represented by a single point. (Marques, 2011)

For our purposes there are two types of colour spaces that are of interest to us. These are

The **RGB** colour Space

The **YCbCr** colour space

The Red Green Blue (RGB) colour space is the default method in which our processing software handles the colour information of the video feed. The RGB colour model is a Cartesian co-ordinate system where each of the three primary axis (x,y,z) correspond to the three primary colours of light Red=x, Green=y and Blue=z. These values normalised to the range [0,1]. The resulting cube's eight vertices correspond to the three primary colours (RGB), the three secondary colours

(Magenta, Cyan and Yellow), pure white (When (x,y,z)=(1,1,1)) and pure black (When (x,y,z) =(0,0,0)). (Marques, 2011)

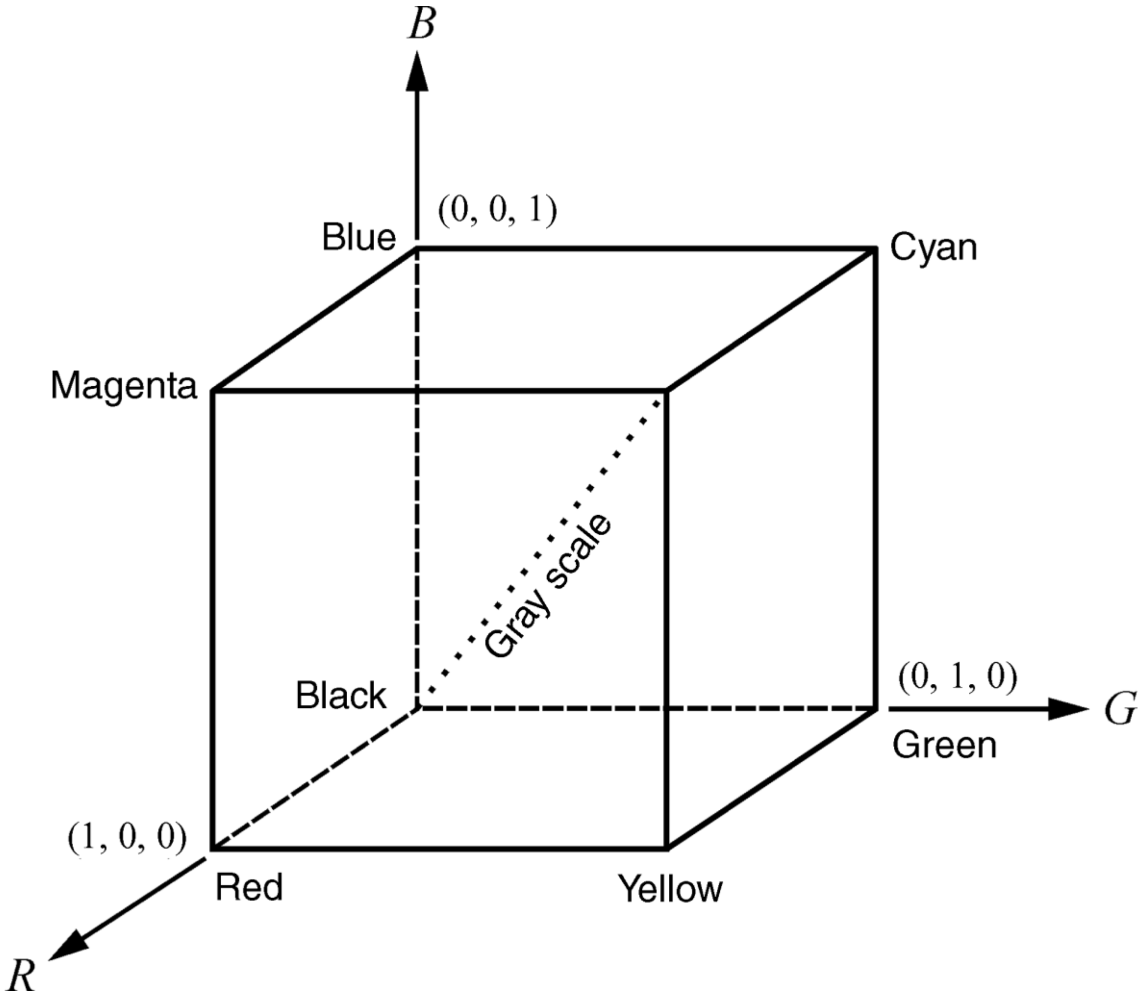

Figure 21: The RGB Colour Model (Marques, 2011)

The ability of the RGB colour space to display colour is dependent of the number of bits used to represent each pixel. The more bits used for each pixel the more colour combination that can be produced by the model. For 24 bit colour each pixel has a size or bit depth of 8 bits. This allows the reproduction of 16 million colours (Marques, 2011).

The processing software represents an RGB image as a 3D array of dimensions H x W x 3, where H and W are the image frames height and width respectively in pixels. The last dimension, 3, represents the three different colour planes or channels. Each channel contains the Red, Green or Blue colour information, each colour pixel is a triple containing the information of the three colour channels. For 24 bit colour each pixel has a range of [0,255].

Figure 22: The RGB Colour space. Notice how the greyscale images change through the different colour channels. In the red channel we can see the fire much more easily than in the Green and Blue Channel. We can see the same effect for the tractor and the sky in the Green and Blue channels respectively.

We can see from the above set of images that for region containing the fire, the red channel is greater than the green channel and the green channel is greater than the blue channel. From this we can start to develop a rudimentary rule set for the fire detection algorithm

$$\begin{gathered} For\ Fr_{(x,y,t)} \\ R_{(x,y,t)} > G_{(x,y,t)} > B_{(x,y,t)} \end{gathered} \tag{4}$$

Where Fr is the Fire Region and R,G,B are the Red, Green and Blue Channels respectively. x,y represent the pixels location within the image frame and t is the frame number.

Figure 23: Notice the difference in intensity across the three colour channels.

To get a better idea if this rule is correct I segmented out the fire region of the picture and calculated the mean values for the Red, Green and Blue Channel.

Table 3: The mean pixel values of the fire region

Mean Red	Mean Green	Mean Blue
215.3274	153.6781	80.7202

This supports the idea that in the fire regions the Red intensity is greater than the Green intensity and Green intensity is greater than Blue. Using the RGB colour space has significant disadvantages, namely how this colour space handles changes in illumination. If the illumination of the frame changes then the value of the pixel in each colour channel changes as well. This means that a potential fire detection program will be affected by changes in illumination, it may work better in shadow than in direct sunlight for example.

Additionally, it is not possible to separate the intensity of the pixel from chrominance in an RGB colour model. To understand these two values, think of an object under intense lighting and then the same object under a dark lighting. The object has not changed colour but the colour has appeared to become less intense or darker. The intensity is the change in lighting while the chrominance is the colour.

If we have the chrominance information of each pixel available separately to the intensity of each pixel, we can create a more robust fire detection algorithm. To do this we need to use a different colour space.

YCbCr Colour Space

The YCbCr colour space also uses three channels to store the colour information of the image frame. Unlike the RGB colour space, one of these is entirely dedicated to storing the intensity part of the image, this is known as the Y Channel. The other two channels contain the colour information, or chrominance. In this case the information stored is actually the colour difference. The Chrominance Blue(Cb) Channel contains the difference between the Blue Channel and a reference value and the Chrominance Red (Cr), the difference between the red channel and a reference value. (Marques, 2011)

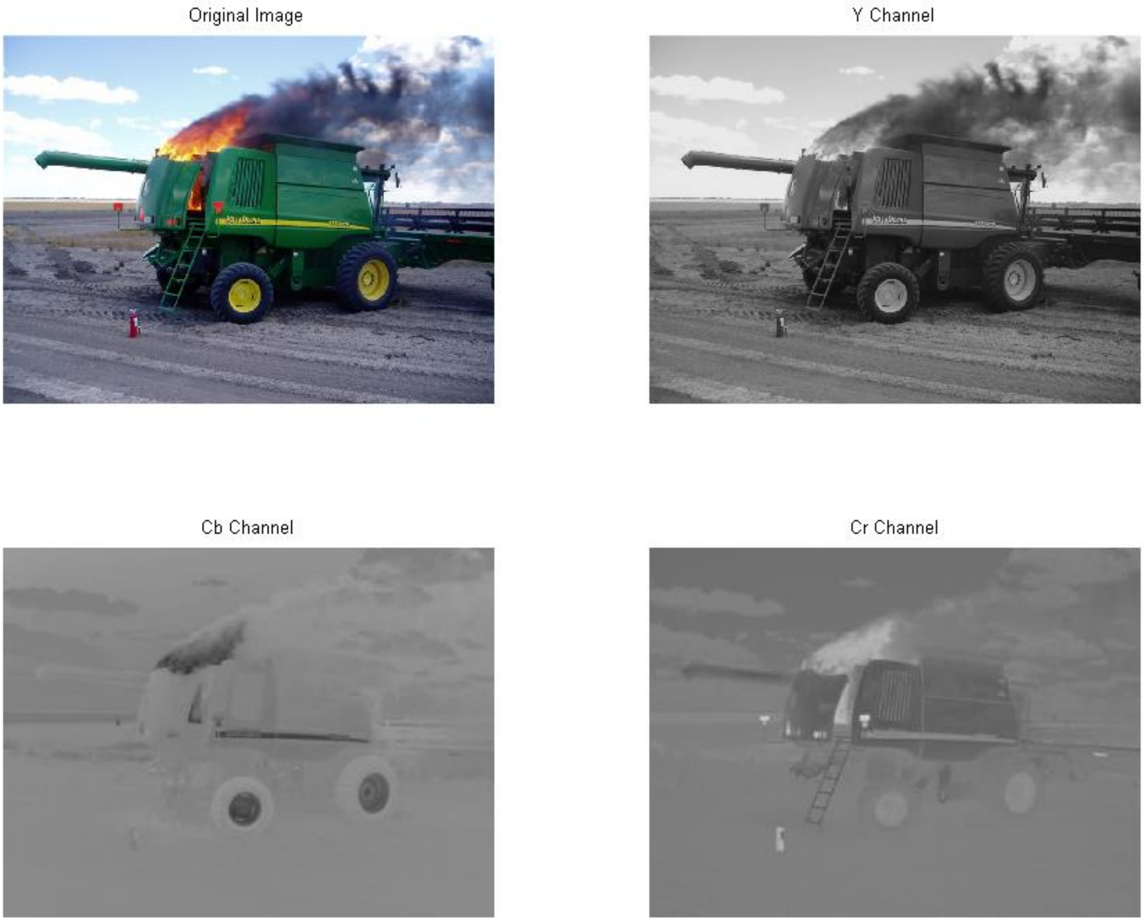

Figure 24: The same image as above, this time in the YCbCr colour space.

We can easily convert our RGB image into the YCbCr colour space using a linear conversion

$$\begin{bmatrix} Y \\ Cb \\ Cr \end{bmatrix} = \begin{bmatrix} 0.2568 & 0.5041 & 0.0979 \\ -0.1482 & -0.2910 & 0.4392 \\ 0.4392 & -0.3678 & -0.0714 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} + \begin{bmatrix} 16 \\ 128 \\ 128 \end{bmatrix} \tag{5}$$

Where Y,Cb and Cr are the intensity, Chrominance blue and Chrominance red respectively. (Turgay Celik, 2009)

Once again we can see that the fire has distinct characteristics in the different colour channels. We can see in the that the fire is defined by a dark area in the Cb Channel and light area in the Cr Channel. To see if the RGB rule can be translated into the YCbCr channel, once again we will take the mean values of the colour area.

Figure 25: Again we can see distinct differences between the colour channels in fire region of the picture.

Table 4: The mean pixel values for the fire region.

Y Mean	Cb Mean	Cr mean
154.7863	87.6760	157.1554

We can see a rudimentary set of rules emerge, our Y channel mean is larger than the Cb channel mean and the Cr channel mean is also larger than the Cb Channel mean.

$$\begin{aligned} Y_{(x,y,t)} &\geq Cb_{(x,y,t)} \\ Cr_{(x,y,t)} &\geq Cb_{(x,y,t)} \end{aligned} \tag{6}$$

To perform a rudimentary fire detection, the image is converted into a logical (1 or 0) array based on the two above equations, giving us two individual arrays. The program can then combine the two arrays into one to give a final result where the fire can be identified easily by its area. We can describe this operation as

$$Fr_{(x,y,t)} = \begin{cases} 1, & if\ Y_{(x,y,t)} \geq Cb_{(x,y,t)}, Cr_{(x,y,t)} \geq Cb_{(x,y,t)} \\ 0, & otherwise \end{cases} \tag{7}$$

Figure 26: Even with only this rudimentary rule set, we can already identify and isolate the fire area.

Figure 27: The result of Equation 1.7, while the fire region is identified, so are other unwanted regions such as the wheels and the vehicle branding.

This method is surprisingly effective for such a rudimentary equation but it still has quite a large amount of unwanted noise in the image, the wheels and the vehicle branding are both yellow in colour in the original image, being similar in colour to the fire, they are also picked up by the fire processing algorithm. Looking back at the original set of YCbCr images, the fire region is one of the brightest parts of the image in the Y and Cr channels. In a sharp contrast the fire region is one of the darkest regions in the Cb channel. The equation to describe this behaviour is

$$\begin{aligned} Y_{(x,y,t)} &> Y_{mean} \\ Cb_{(x,y,t)} &< Cb_{mean} \\ Cr_{(x,y,t)} &> Cr_{mean} \end{aligned} \tag{8}$$

To see if this idea has merit, we can compare the previously obtained mean values of the fire region with the mean values of the overall picture.

Table 5: The mean values of the fire region and the overall picture.

	Y mean	Cb Mean	Cr Mean
Fire Region	154.7863	87.6760	157.1554
Overall picture	122.6689	134.6142	121.7645

We can defiantly see that there is a significant difference between the fire region and the rest of the picture. Apply this set of equations onto our picture yields:

Figure 28: The parts of the image that have above or below average values.

Combine the three equations of 1.8 results in this rule

$$Fr_{(x,y,t)} = \begin{cases} 1, & if\ Y_{(x,y,t)} > Y_{Mean}, Cb_{(x,y,t)} < Cb_{mean}, Cr_{(x,y,t)} > Cr_{mean} \\ 0, & otherwise \end{cases} \quad (9)$$

First Method (Equation 7)

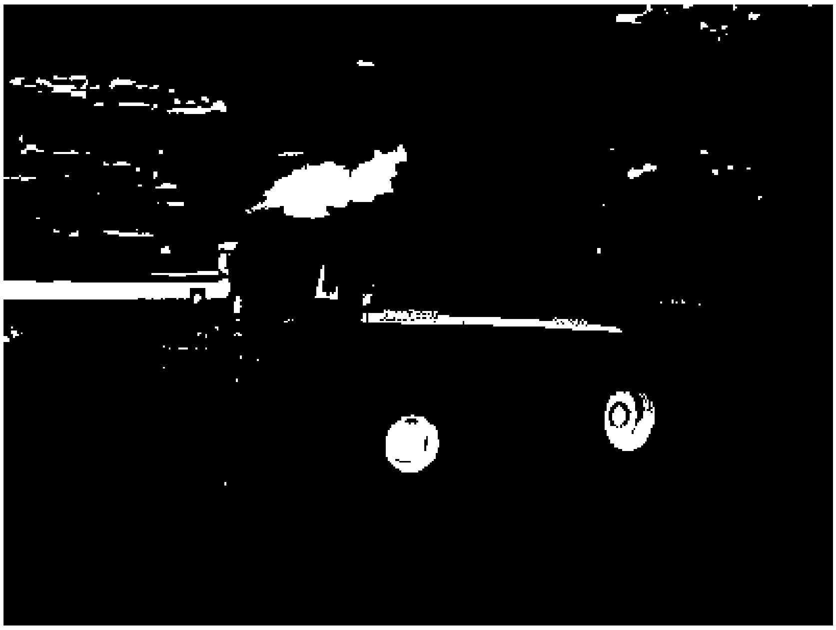

Second Method (Equation 9)

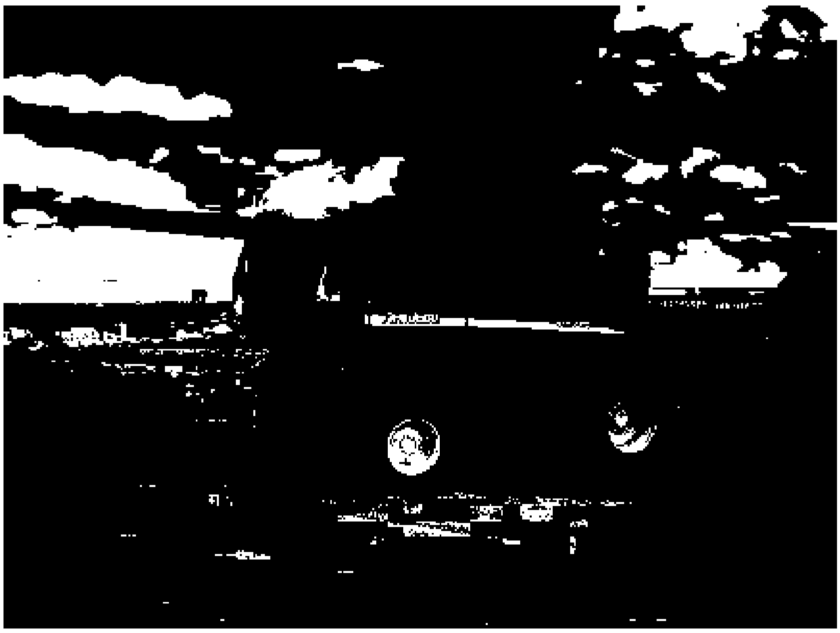

Figure 29: The first and second methods of fire identification.

In many ways, the second method is worse at identifying the fire region. The end result has much more noise, in particular parts of the sky are much more prominent than in the first method. We can still use this second method as an additional filter for our first method. Doing so gives the following result.

Figure 30: Methods one and two combined into one image.

Comparing this result to the result of method one, the wheels of the harvester are now much less prominent than they first were. There has also been a general reduction in the overall noise of the picture, albite only slightly. Clearly we need to improve the process before we can detect the fire with robust accuracy.

If we study the Cb and Cr channels of the image, we can see in the fire region that there is a significant difference between the two. The former is mostly dark or 'black' in colour while the latter is mostly light or 'white' in colour. Comparing the means of these two channels in the fire region there is a difference of approximately 70 between the two. For the overall picture this difference is reduced down to only 13. The rule that can be derived from this is

$$Fr_{(x,y,t)} = \begin{cases} 1, & if \ |Cr_{(x,y,t)} - Cb_{(x,y,t)}| \geq Tc \\ 0, & otherwise \end{cases} \tag{10}$$

Where Tc is a threshold value used to filter out pixels that have an insufficient difference between the Cr and Cb Channels. It should fairly obvious that the result of this equation is highly dependent on this threshold value.

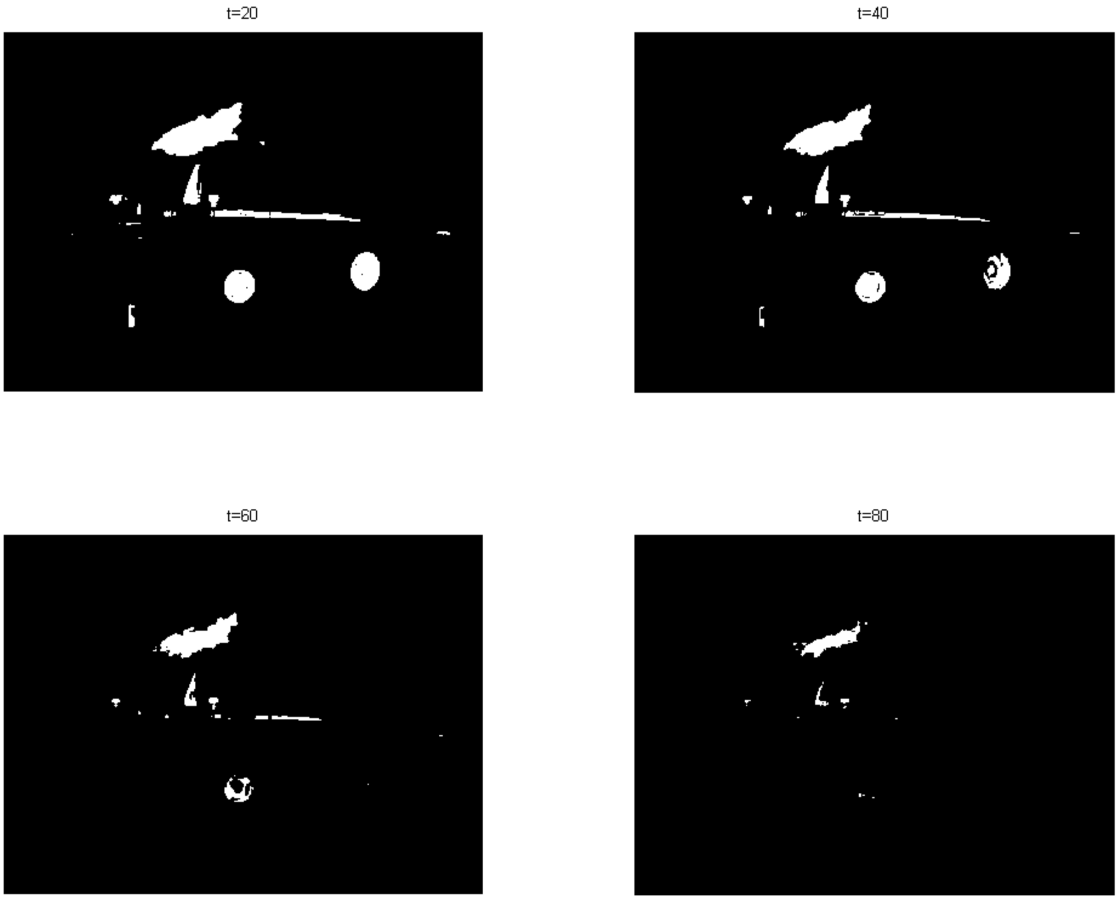

Figure 31: The effect of the threshold value t on equation 10.

The optimal value of t is one that does not miss any fires or true detections while not having any false alarms. It is not realistic to expect the system to be completely accurate in both parameters, so instead the optimal value of t is one that maximises the true detections while minimising the false positive detections. This t threshold is based on the Receiver Operating Characteristics (ROC). These characteristics can be shown as a x-y plot with the false positive rate on the x-axis and the true positive rate on the y-axis.

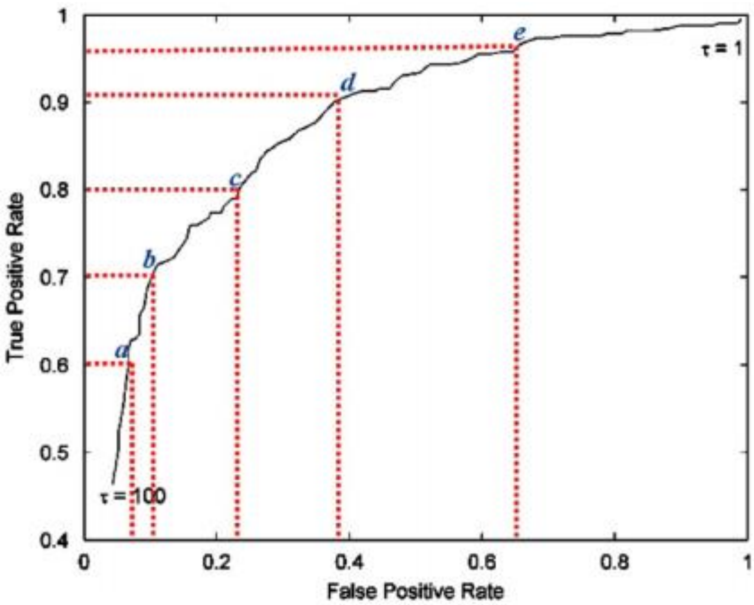

Figure 32:The ROC curve used to determine the threshold value t. (Turgay Celik, 2009)

Determining the ROC curve is a very involved process that would be highly dependent on the cameras used. To save time the ROC curve from *Fire detection in video sequences using a generic colour model* was used as the basis for selecting the t value as well as some basic experimentation with various images.

Combining the result of Equation 10 with that of equation 7 and 9 allows additional noise to be filtered out of the image. These set of equations will make up the basis of colour fire detection algorithm.

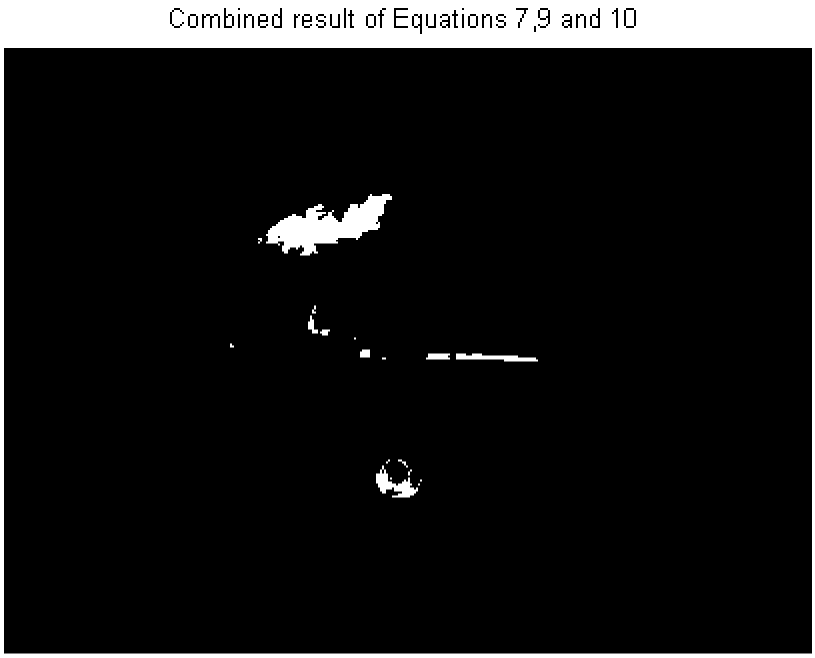

Figure 33: The top picture shows the result of all three equations on the image. The bottom image has had additional filtering to remove small objects.

Tagged Fire Regions

Figure 34: The bottom image of figure 34 superimposed over the original image. We can see that the algorithm has identified the fire area with reasonable accuracy.

The end result still has a small amount of noise in the image frame. Since these regions are much smaller than the fire region we can remove them on the basis of area. If the region does not meet the overall area requirement (In this case it was 800 pixels) it is removed from the frame. The resulting image only has the fire region left. This resulting shape can be used to calculate the size of the fire and its centroid. Both of these criteria will be used later on in the program to diagnose the overall size and severity of the fire.

3.2.5 Near Infrared fire detection

The NIR detection method is different from the colour detection method in that a large part of the processing takes place on the camera, and that this processing is done by the hardware rather than the software. When the thermal radiation from the fire reaches the camera it is composed of both visual and infrared light. The camera filter removes the visual light and only allows the infrared component to reach the camera sensor.

The thermal radiation that is emitted from the object is governed by Wien's displacement law and the Stefan-Boltzmann law.

$$\lambda_{\max} = \frac{b}{T} \tag{11}$$

Where λmax is the peak wavelength, T is the absolute temperature in Kelvin and b is Wein's displacement constant, equal to 2.898x10^(-3) m.K

$$\frac{P}{A} = \sigma T^4 \quad (12)$$

Where P is the total power emitted in watts, A is the of the emitting surface in m^2 , T is the absolute temperature in Kelvin and σ is the Stefan-Boltzmann constant which is equal to 5.6703x10^(-8) watt/m^2.K^4

Wien's displacement law tells us that the wavelength of the infrared light from thermal radiation is dependent on the temperature of the object emitting it. As the temperature of the object increases the peak wavelength decreases. Stefan-Boltzmann law tells us that as the temperature of the object increase the amount of thermal radiation emitted also increases.

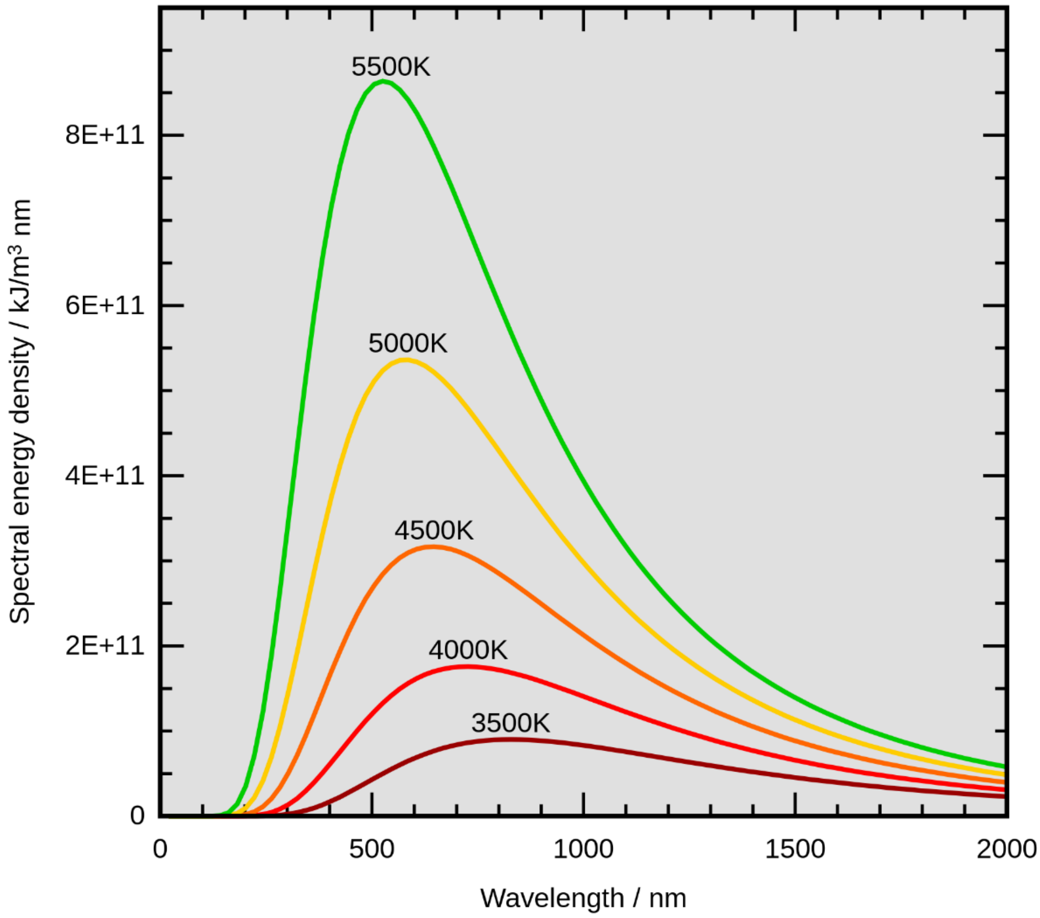

Figure 35: The relationship between wave length and energy density for different temperatures.

These to equations provide the basis for the rule set governing the ability of the NIR to detect fire and hot objects in general. The NIR part of the Infrared spectrum has the shortest wavelength of all IR radiation, therefore it can be understood that thermal radiation located in the NIR band will only be emitted from very hot objects. For the 950nm ± 100nm filter that is used, the minimum detection temperature is approximately 400 ^{o}C. (Y. Le Maoult, 2007)

The second part of this rule set is that hot object will be emitting thermal radiation at a rate well in excess of cooler objects in the same frame. This makes it very easy to filter out objects that are not of interest as a fire will almost always be the hottest and most luminescent object in the frame.

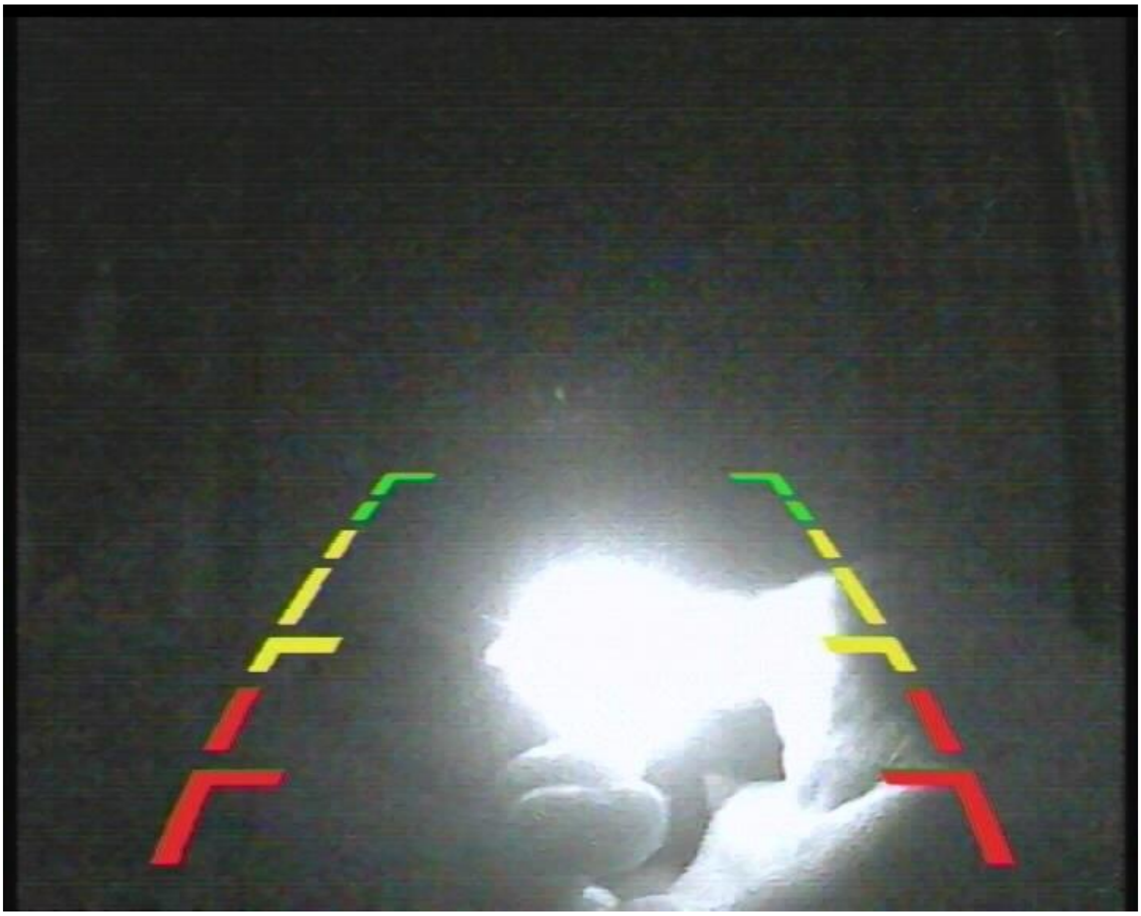

Figure 36: The image from the NIR camera before any software processing takes place.

The image in figure 37 shows a cigarette lighter being lit under NIR light. The lighter fluid was empty, the flash in the above image is purely the result of the flint-wheel ignition system. This bright white area is typical of hot objects viewed under NIR light. Also notice that the image appears to be greyscale. This is due to the lack of colour in the NIR spectrum. We know that the camera is still recording in the normal RGB colour space due to the presence of the hardcoded colour bars in the above image.

The hotspot region, or the region with the largest amount of NIR radiation, can be determined by this equation.

$$Hr_{(x,y,t)} = \begin{cases} 1, & If\ I_{(x,y,t)} \geq Tr \\ 0, & otherwise \end{cases} \tag{13}$$

Where Hr(x,y,t) is the resulting hot spot image, I(x,y,t) is the original luminance frame of the image and Tr is a threshold value use to filter out objects that are not hot objects.

This equation is normally applied after the image frame has been converted to a true greyscale image. This is done to reduce processing time as the data is reduced by 2/3rds. Since the image is already very close to greyscale anyway this does not reduce the accuracy of the algorithm at all.

The Threshold value is set well above the mean of the overall image as a rule due to the high luminance of the hot spot region. In this case it was set 80% of the maximum luminance.

Figure 37: The NIR image converted to luminance only.

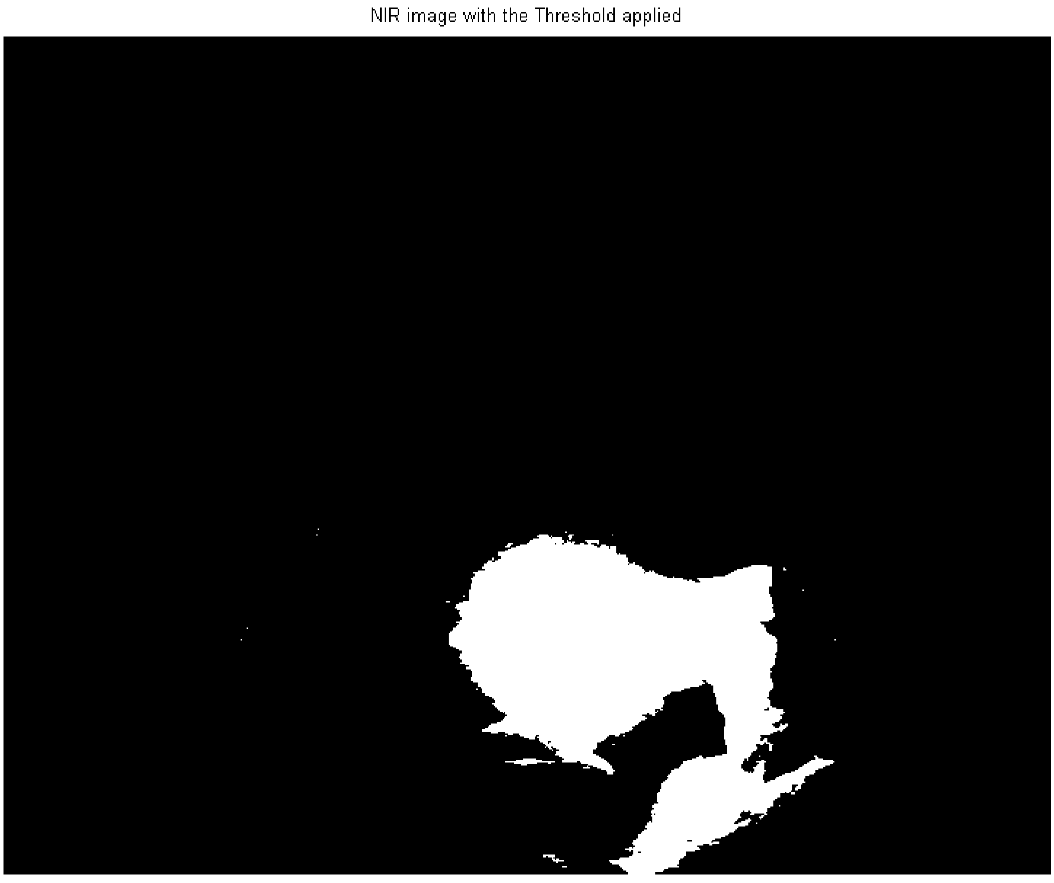

Figure 38: The NIR image with the threshold applied.

Most of the time additional morphological operations must be applied to the image frame to filter out unwanted noise. The above image having an almost complete lack of noise is the exception rather than the rule.

There are two additional operations performed on the image. The first operation removes objects under a certain area, this is the same operation that is performed on the colour video frame except that the size threshold is much smaller, only 200 pixels in this case.

The second operation is somewhat more complicated, this object filters the image by the areas shape and size. Based on initial testing, a fire region presents itself as an oval or circle shaped bloom. Unwanted noise such as reflections tends to be an irregular shape. Using this final morphologic operation, we can remove these irregular shapes and only leave the circle or oval shaped bloom of the fire.

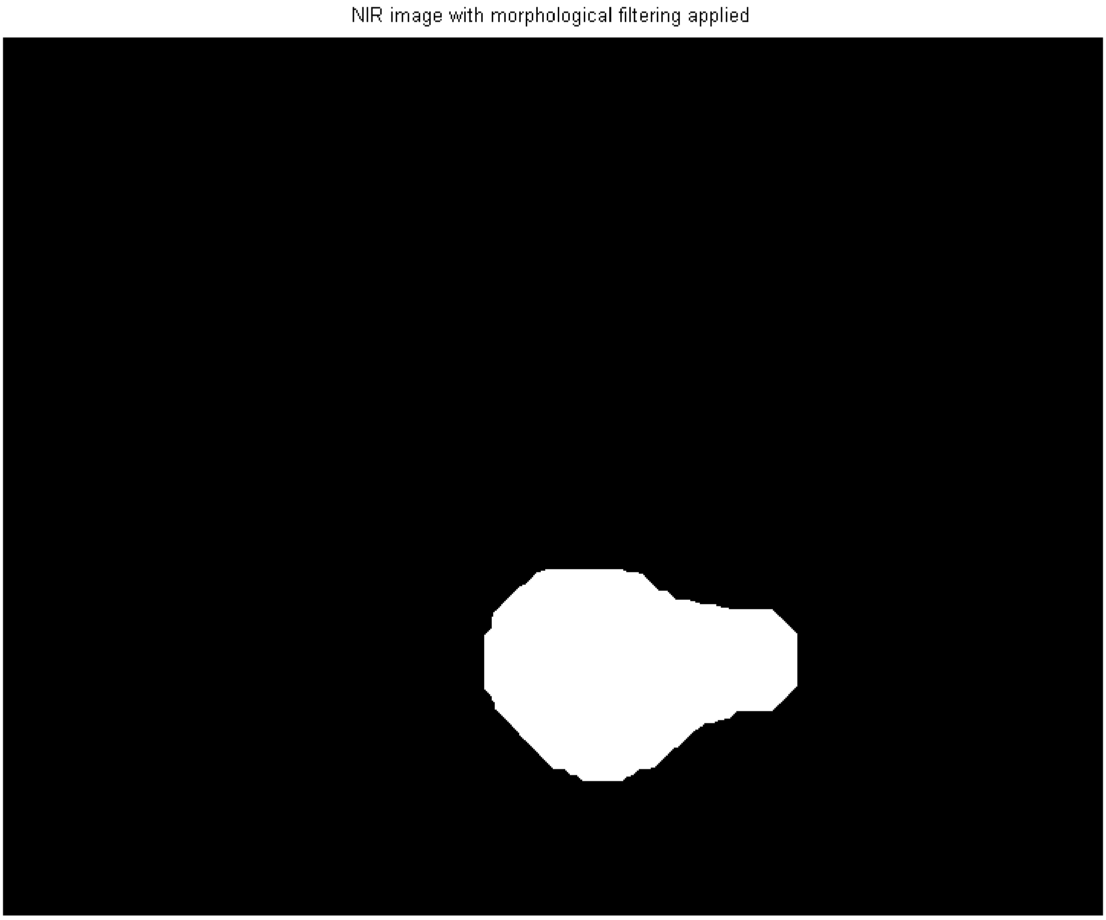

Figure 39: The image with the filtering operations applied. Notice that the first operation has removed all of the little dots around the main bloom. The second operation has removed part of the bloom that was a reflection off the hand and has rounded off the edges of the main bloom.

From here we can use this shape to determine the size of the fire and centre point of the fire.

3.2.6 Movement detection

Our movement method can be defined as the real-time segmentation of moving regions of the image frame (P. Kaewtrakulpong, 2001). In layman's terms in order to identify movement we separate the moving regions of the image frame, known as the foreground, from the static parts of the frame, referred to as the background.

One of the typical ways of performing this is known as background subtraction. In this method a reference background image is calculated and is subtracted from each image frame. Then a threshold value is applied across the resulting frame leaving only what is known as a foreground mask. (Tamersoy, 2009)

$$| I_{(x,y,t)} - B_{(x,y,t)} | > Th \tag{14}$$

Where I(x,y,t) is the current frame, B(x,y,t) is the background frame and Th is the threshold value.

The Threshold value is the absolute intensity difference between the current frame and the background image. The higher the threshold, the bigger the difference between the two must be for it to be considered a foreground object. (University of California, 2013)

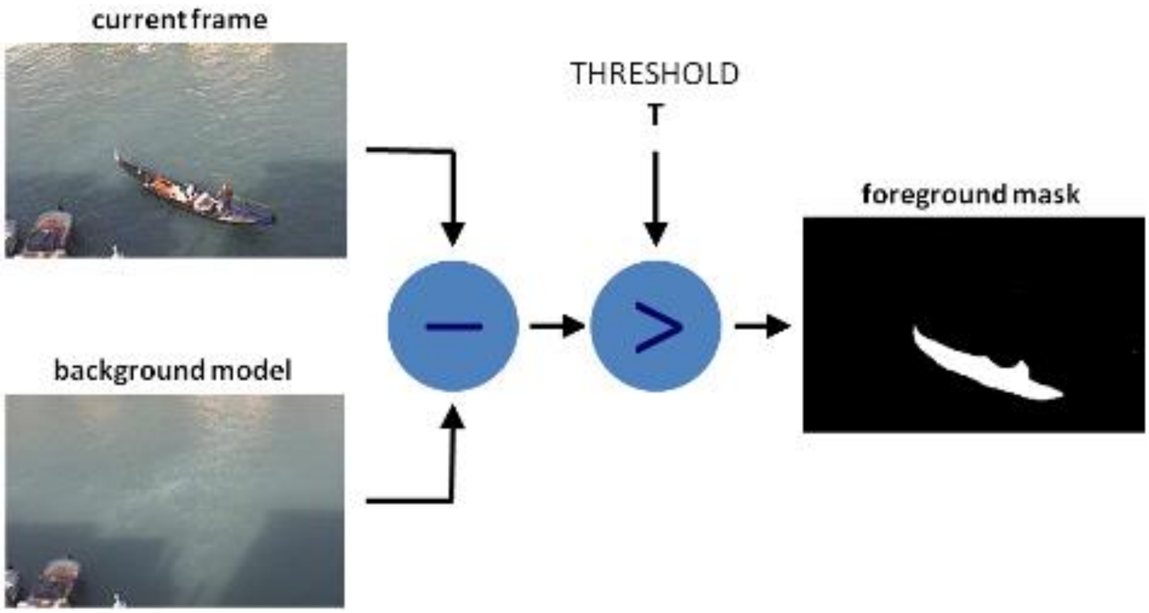

Figure 40:A very basic background subtraction mode using frame differencing (OpenCV, 2016)

The problem then becomes how to automatically calculate the background reference image. The simplest method creating the background image is to simply use the first image of the video or camera feed. So long as the moving object has a colour or intensity that is different from the background, we can identify it in the video feed. (University of California, 2013)

This simple method has a number of issues. Any object that enters the frame and stops will continue to be detected. If an object that was static at the start of the video feed starts to move, not only will it be detected a 'ghost' where the object original was will also be created by the program. On a longer term basis, slight changes to the background frame will create problems. Changes in illumination as the sun moves over, the movement of trees and other small changes will still be registered as movement. And the camera cannot be moved as this would also completely change the frame. (University of California, 2013)

The simplest way is a method known as **frame differencing**. In this method the current frame is compared with the previous frame of the video and has a global threshold applied.

$$B_{(x,y,t)} = I_{(x,y,t-1)}$$
$$| I_{(x,y,t)} - I_{(x,y,t-1)} | > Th \tag{15}$$

Where I(x,y,t) is the current frame, B(x,y,t) is the background frame and Th is the threshold value.

This methods detection ability is highly dependent on the object speed and shape, the frame rate of the video and the global threshold. If the object is moving fast and the frame rate is too slow then it will miss the object as it passes through the frame. The threshold value also has a dramatic effect on what the system detects. Too high and it will miss movement, too low and background noise will overwhelm the system. The advantage of this type of processing is the minimal computational requirements and its extreme ease of implementation. (Piccardi, 2004)

Instead we can use a background frame that is averaged over the length of the video. The two common methods are known as Mean filtering and Median filtering; Mode is also possible but much less common than the other two methods.

Mean Filtering

$$B_{(x,y,t)} = \frac{1}{n}\sum_{i=0}^{n-1} I_{(x,y,t-i)}$$
$$| I_{(x,y,t)} - \frac{1}{n}\sum_{i=0}^{n-1} I_{(x,y,t-i)} | > Th \tag{16}$$

Where n is the previous frames and $i \in \{0,1,...,n-1\}$ (Tamersoy, 2009)

Median Filtering

$$B_{(x,y,t)} = median\left\{I_{(x,y,t-i)}\right\}$$
$$| I_{(x,y,t)} - median\left\{I_{(x,y,t-i)}\right\} | > Th \tag{17}$$

Where $i \in \{0,1,...,n-1\}$ (Tamersoy, 2009)

Both of these methods are fast, but the memory required by the algorithm increases the longer the program runs. The memory requirement is $n * size(frames)$ (Piccardi, 2004)

To reduce the memory requirement, the background reference image can be a **running average**.

$$B_{(x,y,t)} = \alpha I_{(x,y,t)} + (1-\alpha) * B_{(x,y,t-1)} \tag{18}$$

Where α is the learning rate of the program. (Piccardi, 2004)

These methods are basic, computationally fast and easy to implement. But have significant drawbacks. The threshold value used in all of these approaches has no ability to change over time and is applied globally to all of the pixels in the frame. These approaches will produce sub-optimal results if the objects are slow moving, numerous and if the lighting conditions change. (Tamersoy, 2009)

None of these methods are particularly suitable for use in our fire detection program. We can expect there to be changes in the illumination of the scene as the harvester changes orientation to the sun as well as the sun moves across the sky during the day. We can also expect that the cameras maybe subjected to a slight amount of vibration from the harvester moving parts and we can also expect that there will be moving parts within the field of view of the camera.

For the cameras that may be possibly mounted on a UAV platform, the background will be forever changing as it moves around the harvester, making these approaches impractical.

For the movement detection method to be useful in detecting fire, the background subtraction algorithm must be robust enough to handle these repetitive motions and changes in the lighting of the scene.

One way of meeting this challenge is the use of adaptive background Gaussian mixture models. The values of each pixel are modelled as a mixture of adaptive Gaussian. A mixture of Gaussians is used as multiple surfaces appear in each pixel and it is adaptive to lighting conditions changes. (Tamersoy, 2009) In this method of background subtraction the history of any pixel at X0,Y0 for time n is

$$\{X_1,...,X_n\} = \{I_{(x0,y0,i)} : 1 \le i \le n\} \tag{19}$$

Where I is the video sequence. (Stauffer, 1999)

The recent history of each pixel $\{X_1,...,X_t\}$ is modelled by a mixture of K Gaussian distributions. The probability that a certain pixel has a value of Xn at time N is

$$p(X_n) = \sum_{j=1}^{K} w_j \eta(X_n;\theta_j) \tag{20}$$

Where w_k is the weight parameter of the k^th^ Gaussian component. $\eta(X;\theta_k)$ is the normal distribution of k^th^ component represented by

$$\eta(X;\theta_k) = \eta(x;\mu_k,\Sigma_k) = \frac{1}{(2\pi)^{\frac{D}{2}} |\Sigma_k|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-\mu)^T \Sigma_k^{-1}(x-\mu_k)} \tag{21}$$

Where μ_k is the mean and $\Sigma_k = \sigma_k^2 I$ is the covariance of the k^th^ component. (P. Kaewtrakulpong, 2001)

The K distributions are ordered based on the fitness value $\frac{w_k}{\sigma_k}$ and the first B distributions are used as a model of the background of the scene where B is estimated as

$$B = \arg\min\left(\sum_{j=1}^{b} X_j > T\right) \tag{22}$$

Where T is the minimum fraction of the background model, or the minimum prior probability that the background is in the scene. If a pixel is more than 2.5 standard deviations away from any of the B distributions, it is considered to be a foreground pixel and no longer part of the background. (P. Kaewtrakulpong, 2001)

The foreground pixels are grouped using 2D connected component analysis. This method is able to use a different threshold for each pixel. These pixel thresholds are able to adapt over time. Object that enter the scene and stop are allowed to become part of the background model and it has fast recovery. It is not perfect, this model still cannot deal with sudden, drastic changes in the scene, be it changes in lighting or otherwise. The Gaussians need to be initialised and there are many parameters to the overall model that all have an effect on how it will perform. (Tamersoy, 2009)

The end result is a background subtraction model that is able to adapt to changes in the background model in ways that the simpler models cannot.

3.2.7 Combined detection methods

Now that we have defined our basic methods of detection, we need to combine these into one singular program. We know that the software process starts with the Raw video feed, we know that it ends with some form of output to the operator. How the program gets from the start to the end of the process can be done in several different ways.

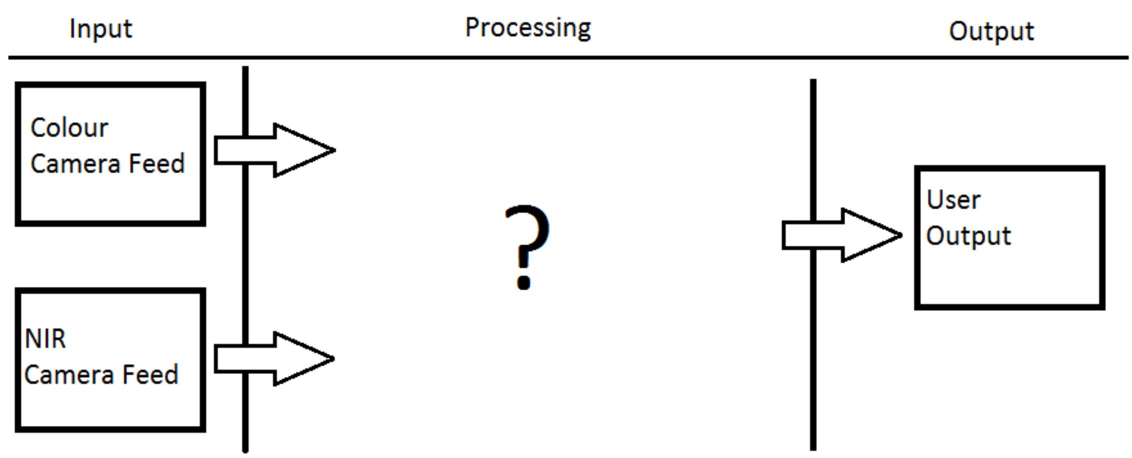

Figure 41: We know the input and output of the system. But how can the processing steps be configured for the optimum result?

We can treat each of the fire detection processes as individual processing blocks within the overall program. The movement processing block or blocks are able to either process the raw video feed for movement or the processed video from the other two blocks. Or we could reverse this and have the movement processing done first and then the Colour and NIR blocks process the feed from the two movement blocks.

There is also the need to combine these results into an easily understandable metric for the operator. Do we just supply him with the raw results from each process and let him make the judgement or do we introduce a level of judgement into the system allowing it to classify the potential risk factor based on the size, position or movement of the detected anomaly.

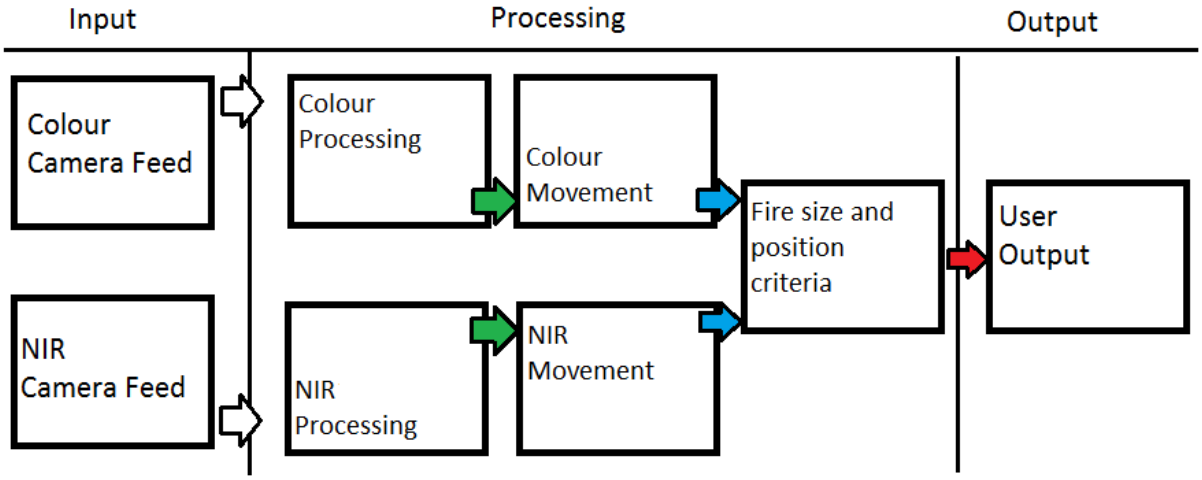

Figure 42: Dependent processing. Each set of blocks is fed the processed information from the previous step.

In the dependent processing method, each step is fed information from the previous step. In this case rather than feed the raw video footage into the movement states, it takes the processed video from the first stage and performs the background subtraction algorithm on it to detect movement. This result is then fed into the 'judgement' block where based on various criteria it decides if the fire detection is genuine or a false positive and alerts the operator if necessary.

The advantage of using this processing layout is speed. The data being fed into the movement blocks is a 1bit logical array (A black and white image) as shown in figures 34 and 39 for Colour and NIR respectively. The amount data needed to be processed is reduced, the movement block can work much quicker.

The disadvantage is a reduction in accuracy, if all of the blocks are working on independent data, there is essentially 4 sets of criteria that must be met to register a true fire detection, under this system, there is only two sets of criteria. The above system would be completely unable to identify hotspots, which a are major warning of a potential fire.

So instead we can improve the system by allowing the data from the first step to be fed directly into the fire criteria block.

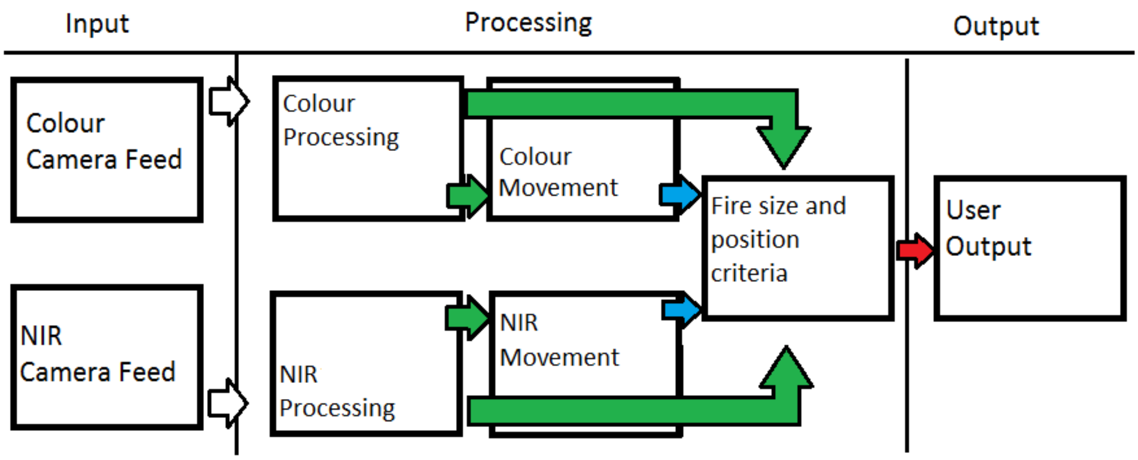

Figure 43: Improved dependent processing. The data from the first step is now allowed to be used as fire criteria.

Now the system can utilise the data from the fire from the Colour and NIR processing steps independently of the movement steps while still keeping the processing speed advantage.

3.2.8 Fire detection criteria

The final step of the fire software process before it gets to the end user is to determine how likely it is that a true fire is happening based on the data from the previous steps. The main criteria that we can use is the size of the detection and the position of the detection.

Size detection takes the end results of the NIR and Colour processing algorithms and calculates the size of the resulting 'blob' of fire in each of the images. The larger the blob the more likely that there is a fire event in progress.

Position detection utilises the centre point of the detection for each method and determines the distance between these points. If there happens to be a colour detection, NIR detection and movement detection all within a small region of the frame, then the system can make a reasonable assumption that a fire has begun and take steps to alert the operator. Conversely if there were detection events for all of the different detection methods at the same time but they were spread out across the image frame, it is much less likely that a fire event is occurring.

4.Experimental Set up and Procedure

4.1 Experimental Hardware

4.1.1 Hardware Requirements

Based on what has been discovered in the background information and the literature review, the system must have these required parts and abilities in order to best simulate a fire detection system for a combine harvester.

One or more digital video cameras. One of the cameras should be able to observe multiple spectral bands, particularly NIR and visible light. It is also preferable that the camera has no in built filter so that minimal modification is needed to use and external filter. Finally, the camera should be able to run of 12v, be rugged, lightweight, easily obtainable and as inexpensive as possible.

Light filter/s. These are attached to the camera and allow only a specific wavelength to pass through it. The filter is a NIR pass filter for Infrared only testing.

Video capture device. This device allows the signal from the video camera or cameras to be interpreted and stored as a digital video file on a computer. The device should have the ability to capture video from multiple sources at once, at the full resolution and framerate that the cameras are capable. It must also be compatible with the computer device that will eventually process and store the video.

Digital Signal Processing. Regardless of if the video is captured and stored or processed live. There will need to be some device capable of processing the video. Ideally this device should have the processing power to be able to process video frames at the same rate or faster than the frame rate of the cameras. It should also be able to process multiple video streams at once in order to simplify the overall configuration of the system.

User Feedback. The system needs to be able to provide either an audio or visual signal to a human operator in the event of a possible fire detection.

4.1.2 Parts List

CCD Reversing Cameras

The easiest source of a true CCD camera is a car reversing camera with an advertised ‘Night Vision’ capability. They are readily available online at a reasonable cost of approximately $30-40 and come in a durable and weatherproof metal housing. The ‘Night Vision’ capability is, in fact, a IR flood light combined with a lens without any Infrared filter. I have used this to my advantage as I am able to readily adapt the camera to only pick up NIR extremely easily by only adding a Pass filter.

The cameras that I currently have do have some drawbacks. The camera sensor is hardwired to show a reverse image and has reversing lines also hardwired into the image. So far my attempts to undo the hardwired image and reversing of the video has resulted in the destruction of one camera. So for the foreseeable future I will persevere with the camera as it is.

The other issue with the camera is the inherent ‘cheapness’ of it. It is not a HD camera and its image output is not the best quality. I believe that for now it should be adequate for what I need it to do, my initial testing after finishing assembly shows that the camera is able to detect a flash from a cigarette lighter under ideal conditions. Additional by using a camera with a lower overall resolution, the processing overhead is reduced meaning that a slower but cheaper DSP can be used.

The camera comes in a rugged, waterproof housing made out of aluminum and a mounting bracket. Four of these cameras were ordered during the project, two were accidentally destroyed.

Figure 44: CCD Reversing Camera (eBay,2016)

Specifications:

Image Device	1/3 CCD
Maximum Resolution	782*628
Video Format	PAL
Horizontal Lines	420 TV Line
Video Output	75Ω/1.0Vp-p
Signal/Noise Ratio	>48db
White Balance	>0.45
Electronic Shutter	1/50(1/60)-1/100,100sec
Backlight Compensation	AUTO
Operation Temperature	-10℃~+50℃
LED Switch	AUTO
Minimum Illumination	1Lux/F1.2
Power Adapter	DC 12V/500mA

Drift Action Camera.

During the course of the initial tests it became apparent that the CCD Cameras were not suitable for some of the planned tests. Specifically, the lack of any filtering on one of the cameras meant that Infrared light from heat sources and fires overwhelmed the visual part of the spectrum making any sort of detection process based on normal visual light useless. Therefore, for the rest of the testing process this camera was added to the system platform to make up for this short coming.

Unlike the reversing cameras the Drift Camera uses a CMOS sensor with a IR filter built into the lens. It also features a much higher quality sensor, for testing it was set to a resolution of 1080p, a frame rate of 25 frames per second and a field of view of 90 degrees. While the resolution is much higher than the reversing cameras, it has the same frame rate and field of view.

Unfortunately, the camera lacked the ability to output live video directly to the video capture device. Instead the footage was saved to a micro SD card within the camera and later transferred to the DSP.

Figure 45:Drift HD Ghost Action Camera (DRIFT INNOVATION Ltd, 2015)

Specifications:

Image Device	CMOS
Maximum Resolution	1920x1080
Video Format	H.264
Frame Rate	25 FPS
Field of View	90 Degrees

NIR Pass filter

To adapt one of the cameras to only interpret NIR, a pass filter was used. These filters are normally used on high end photography equipment to allow only a certain part of the spectral band to reach the camera sensor.

The filter used for this project was a 58mm diameter lens filter for use on a Digital SLR. The filter will block any light with a frequency below 850nm and above 1050nm. This type of filter was used as it allowed a true NIR image to be taken, to human eyes the image looks to be greyscale. If a filter with a ‘shorter’ wavelength was used, say 720 or 760nm, the image produced would still have elements of the visible spectrum in it.

Figure 46: An IR pass lens filter for use with a digital SLR. (eBay,2016)

Specifications:

Size: 58mm
Filter: 950nm ±100nm

USB Video Capture Cards

This device takes the RCA video signal input and converts it into a digital signal that can be read by a computer via a usb output. These devices go by the name 'Easy Cap' but the vast majority of the devices are in fact generic copies using multiple different video/audio controller chips.

The first device I bought used a Somagic SMI-2020CBE video/audio controller IC and was advertised to be able to receive and process multiple video streams at once. Once I got the device and started using it I found this not to be the case. The second device I used had a Easycap UTV007 video/audio controller. The UTV007 chip is alleged to be of better quality and produces superior images but I have not noticed any different in quality.

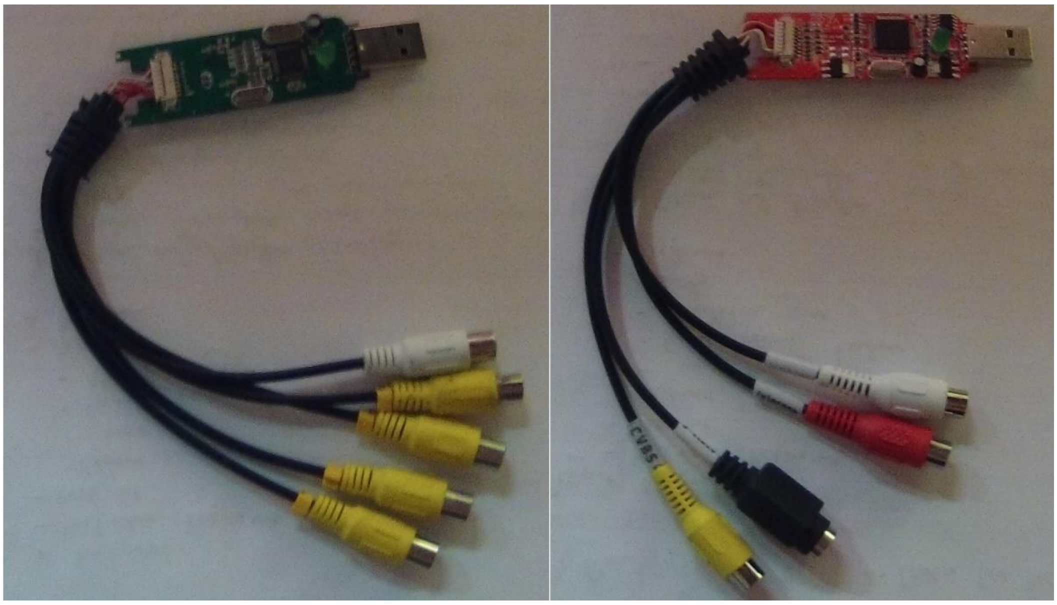

Figure 47: The SMI-2020 on the left and the UTV007 on the right. The SMI-2020 appeared to be configured for multiple connections when in fact only one of them worked.

Despite the different IC and external connections, both devices have worked relatively well and have been trouble free for the most part. There were some initial issues with the second device as the drivers would often cause 'blue screen' errors with windows. But this issue has been rectified since then.

Specifications:

First Device
Chipset: Somagic SMI-2020CBE
Maximum video capture: up to 720x480 at 30 fps for NTSC and 720 x 576 at 25 fps for PAL
Second Device
Chipset: EasyCap UTV007
Maximum video capture: up to 720x480 at 30 fps for NTSC and 720 x 576 at 25 fps for PAL

Digital Signal Processor

For this project, the DSP that will be used is nothing more than a normal consumer grade laptop. The laptop allows video footage to be stored and captured for processing at a later date or it can run the video processing program on a live feed. It also can be used to modify the program on the fly.

The laptop that was chosen on the basis of cost, portability and connectivity. It is by no means a top of the line laptop and is now around five years old. It has a 2.5GHZ Intel Core i5-2520M with 8GB of ram running Windows 7, but it proved to be adequate for the task.

For some of the testing a second laptop was used, this laptop was only used to capture video footage. This was due to the limitations of the video-capture software used.

Specifications:

Model: Lenovo ThinkPad X220
Processor: 2.5GHz Intel Core i5-2520M
Memory: 8Gb, DDR3 1,333mhz
Hard drive: 120GB SSD
Operating System: Windows 7 Professional

Figure 48: The Lenovo ThinkPad X220 (Lenovo,2011)

Miscellaneous Parts

During testing the cameras were mounted on a cheap tripod and were powered by a 12v Lithium-Ion battery. A PVC pipe fitting was used to make an adaptor for the filter to fit onto the Camera

4.1.3 Modifications

IR floodlight removal

The two CCD reversing cameras came with an IR floodlight built into the cameras to facilitate the 'night vision' function of the cameras. The floodlight consisted of an array of IR LEDs switched by a photoresistor. It was removed from the camera as it would introduce IR light into the field of view of the camera, potentially creating a false positive signal.

The removal was very simple as the board operated independently of the camera, it was only held in by four screws and the only connection was a 12v power and earth.

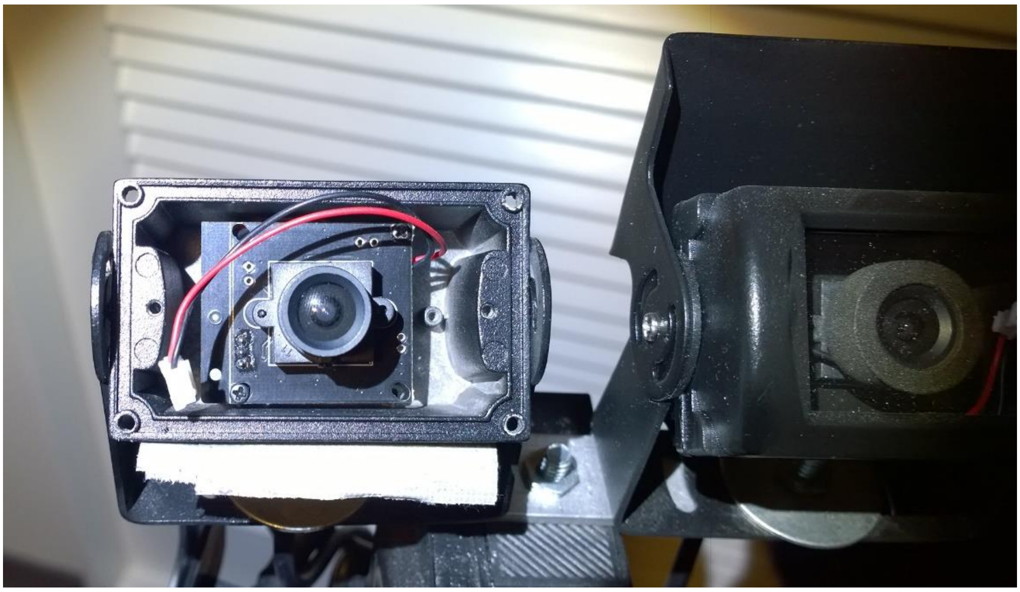

Figure 49: The Camera with its faceplate, sun visor and IR floodlight removed.

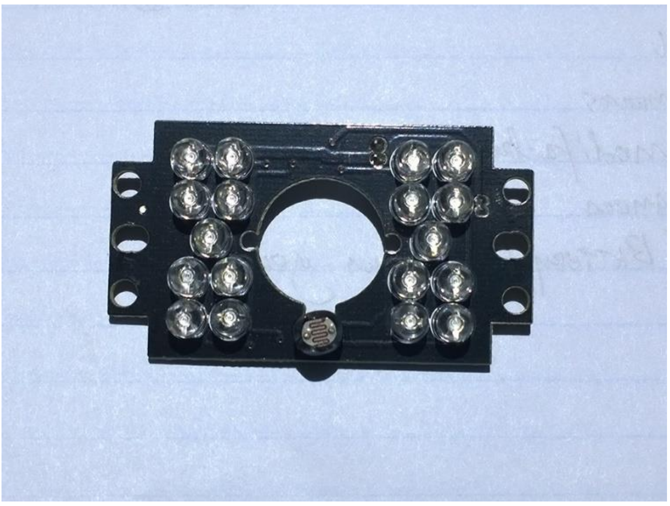

Figure 50: The IR Floodlight.

Install of the NIR Filter

In order to covert one of the cameras to NIR only, the NIR pass filter had to be adapted to fit. The easiest way to do this was to first remove the entire faceplate of the camera body and discard it. Next an adaptor was created using a 60mm to 120mm PVC pipe adaptor, the 58mm was attached to the end of the 60mm section using duct tape and the 120mm section had enough material that it could be cut down into a square shape to fit the camera body. Finally, holes were drilled into the adaptor so it could be secured to the camera body.

Figure 51: The NIR Filter with adaptor.

4.1.4 Sensor platform detail

General Configuration

Figure 52: The general configuration of the Sensor Platform. The IR Camera is on the left and the full spectrum camera is on the right.

Figure 53: Close up of the Cameras. The camera on the left has had the face removed and a NIR pass filter with an adapter has been attached to the body. Both cameras have had the original IR flood lights removed and are otherwise identical. The cameras can be pivoted with the handle.

Figure 54: The mounting configuration of the Drift Action Camera.

Test Images

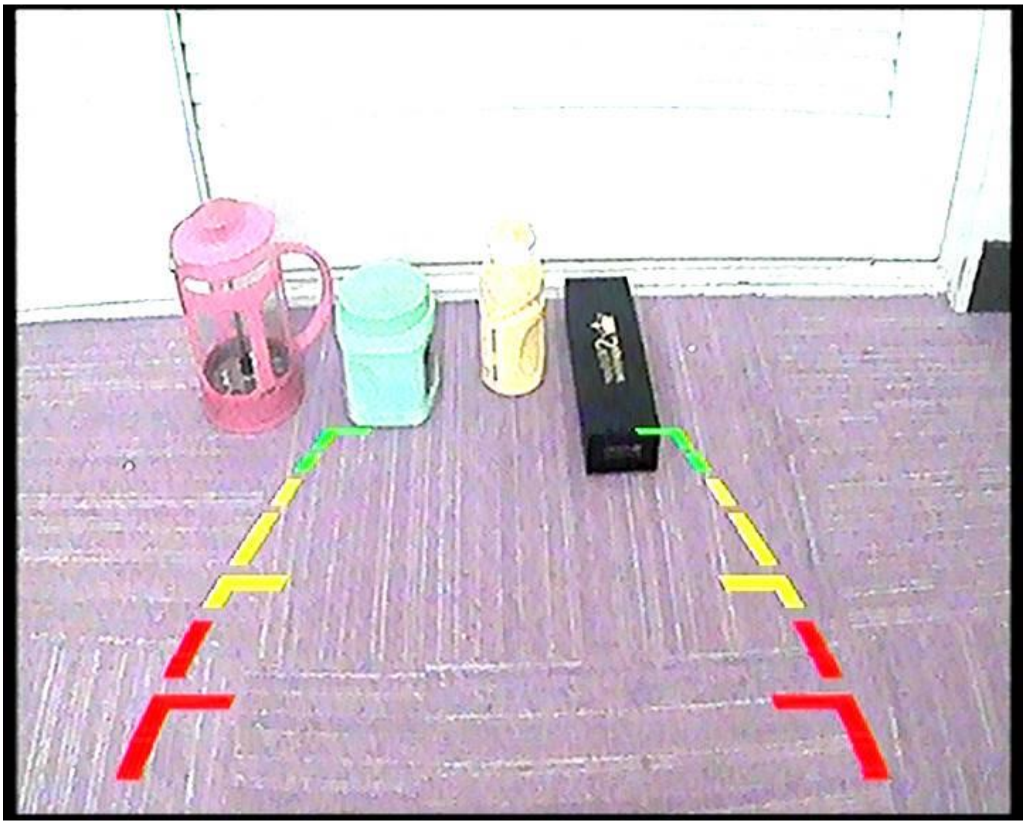

Figure 55: Still image from the unfiltered CCD camera.

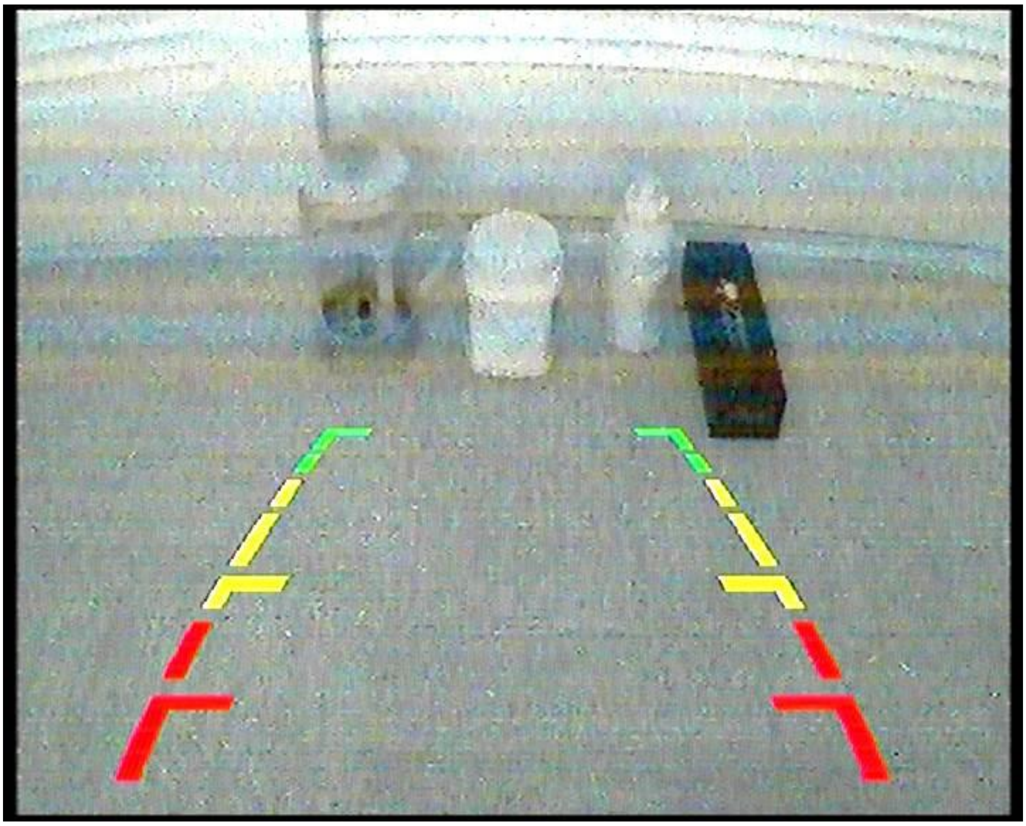

Figure 56: Still image from the NIR camera. The increase in noise is due to the lower total amount of light being received by the sensor.

Figure 57: Still image from the Drift Action Camera. The much more expensive sensor in this camera results in much better image quality.

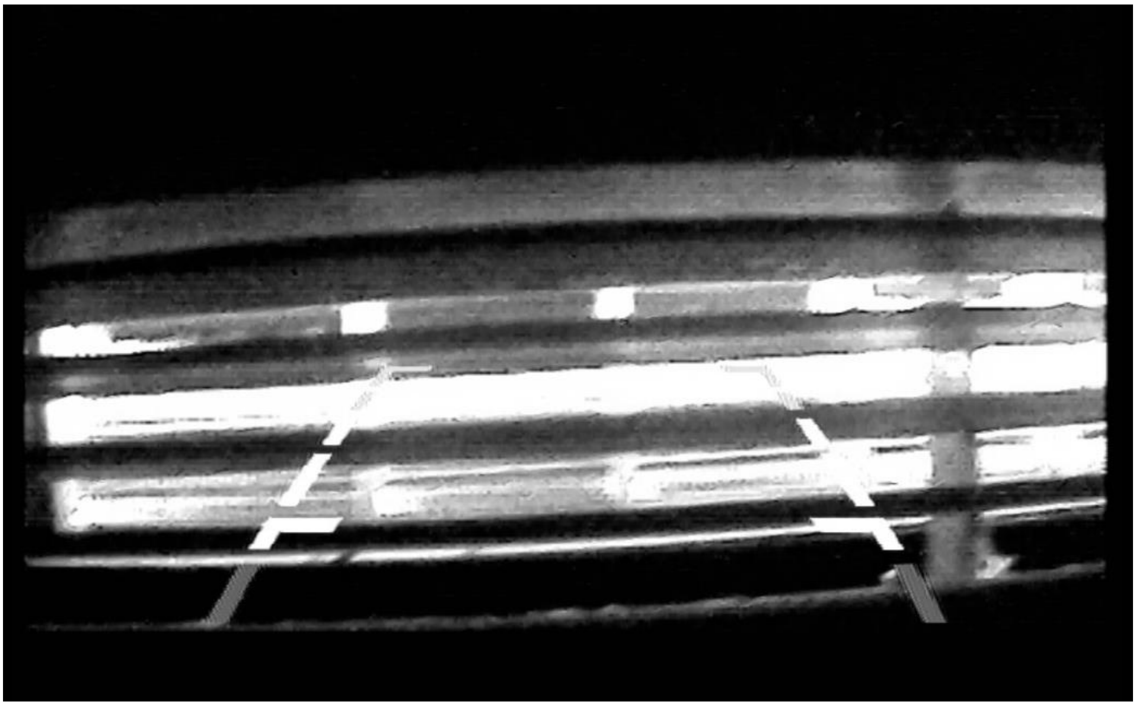

Figure 58: A view of a heater coil through the NIR camera. When this coil is viewed with the naked eye it only dimly glows red but when viewed in NIR it is very bright.

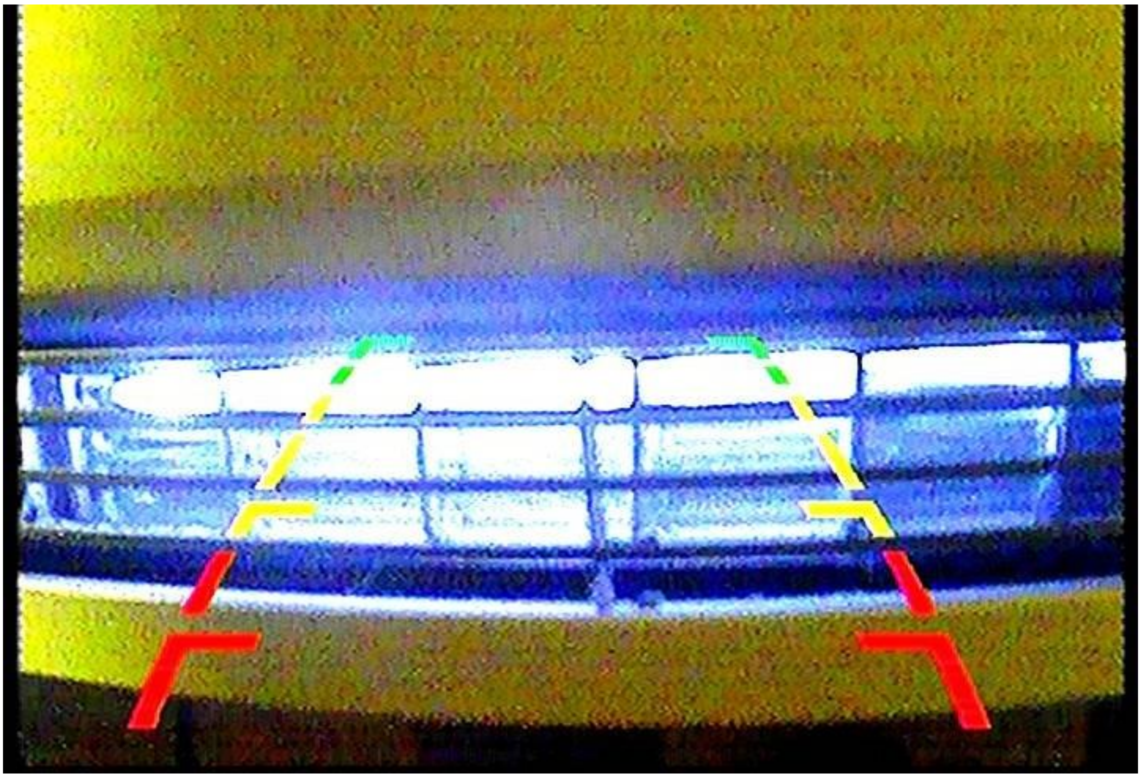

Figure 59: The heater coil view through the Colour CCD camera. Notice that the infrared light is overwhelming the sensor with its huge white glow.

Figure 60: The heater coil view through the drift camera. Only when the infrared light is cut can we actually see the coil.

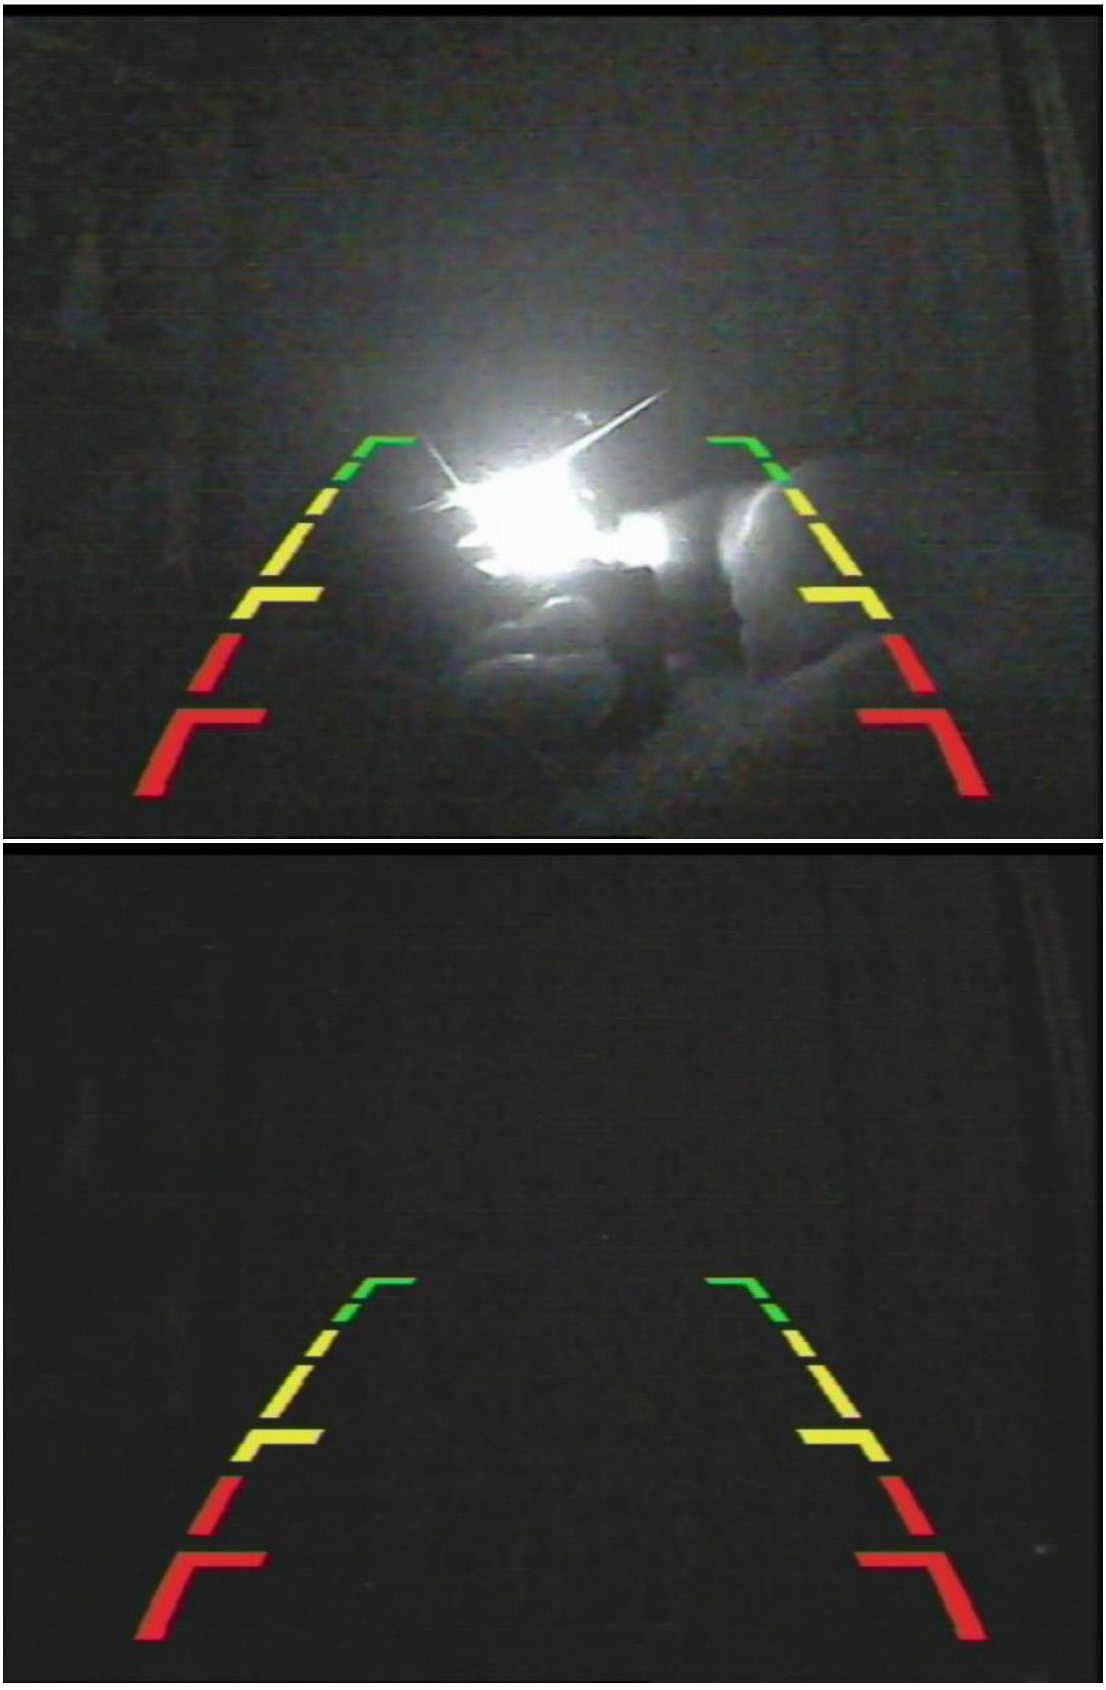

Figure 61: A two frame sequence of a cigarette lighter being struck in a dark room viewed by the camera with the NIR filter.

4.2 Hardware Testing

This phase of testing is designed to discover how the cameras and other parts of the systems hardware see the fire, smoke and hotspots under various environmental circumstances including high and low levels of ambient light, airflow, particulate and other factors that may affect the ability of the system to detect fire. The tests will be recorded using the video capture software detailed in the methodology section, and the footage captured during this test will be used during software testing.

4.2.1 Equipment and setup

Testing equipment

- The System Prototype, including the sensor platform, the laptop and software. For the vast majority of the tests the footage from the cameras will be recorded and processed at a later date.
- An Electrical Hotplate will be used to provide indirect heat in order to simulate a hot surface or part within the harvester.
- A Butane Torch for heating in the hotspot testing and as a backup means of ignition.
- An Industrial fan to simulate air-flow from radiator fans and wind.
- A metal pan for holding the flammable material during testing.

Testing Material

- Steel bar, for hotspot testing
- Diesel fuel, the most common type of fuel used in the engines of combine harvesters.
- Straw, serving as an approximation of the chaff material.

Safety Equipment

- Water hose
- Class D Fire Extinguisher
- Welders gloves, for handling hot objects.

General Set up

The test facility was a large industrial shed with a concrete floor and a concrete pad outside. The area was cleared of any flammable material before testing began. The majority of testing was done at night due to availability of the facility. Originally it was planned that the testing would be done inside the shed, however the lack of ventilation in the shed meant that it was more practical to do the testing on the concrete pad adjacent to the shed.

On the tripod there will be the two CCD cameras with the modifications detailed in the project methodology and the DRIFT action camera.

The tripod will be set up adjacent the tests objects approximately 1.5 meters from the center of the test area, with top of the camera at a height of 130 centimeters above the ground on the tripod. The camera will be looking directly at the center of the test objects.

In some tests the tripod was placed closer to the test object, approximately 40 centimeters away with the cameras being 70 centimeters from the ground.

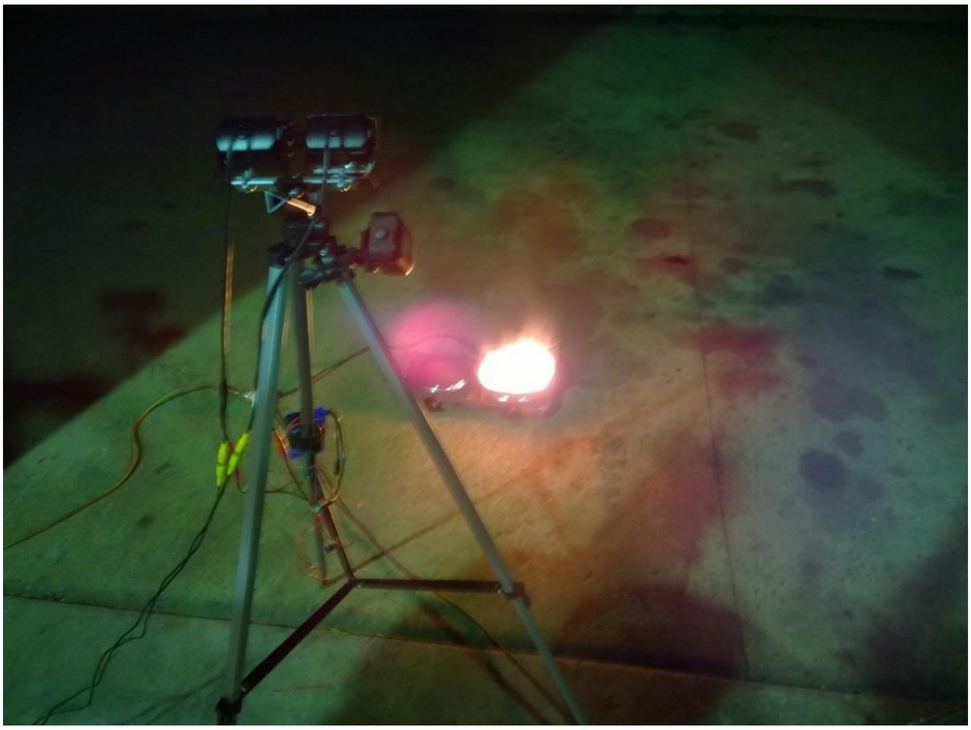

Figure 62: The general arraignment for most of the testing.

The laptops capturing video footage will be present at a safe distance away where the operator will also be present. Once the test has finished, the fire will be allowed to burn itself out if it is safe to do so. Otherwise the water hose or fire extinguisher will be used depending on the circumstances.

4.2.2 Testing Procedure

The first few tests are to establish a baseline for what the camera and the system can be expected to see. The remainder of the tests will introduce the environmental variables in a progressive fashion

The first test will be looking at the cameras ability to detect hotspots on steel and other materials. This test is design to simulate a bearing failure, or a possible fire that is concealed within the machine.

The second test will be looking at the ability of the camera to detect smoke and flame. This test is designed to simulate the start and consolidation of a fire within the engine compartment or on another part of the harvester.

A third test will combine both types of testing in order to best simulate a realistic environment.

The fourth, fifth and sixth test will introduce wind, sunlight and dust respectively. Since the engine bay will have high amount of positive airflow from the radiator fans,,ambient light from the sun and dust and chaff from the harvester operation it is important to understand what effect these environmental variables will have.

First Test: Hotspot Testing

- The sheet of metal will be heated up to a pre-defined temperature using the butane torch.
- This temperature will be confirmed with an infrared thermometer and noted down.
- The sheet will be held in front of the tripod at the same height as the cameras and approximately 1 meter distant.

- A screenshot will be taken with the image capture software.
- The plate will be heated again up to the next temperature.
- Once a range of temperatures have been recorded the test will end and the plate will be left to cool.

Second Test: Smoke/Flame Testing

- The test with two types of flammable material. Straw will be used to simulate the chaff/straw buildup in and around the harvester engine compartment. Diesel fuel will be used to simulate a possible fuel and/or oil fire.
- The diesel fuel it will need to be preheated to its flash point in order to encourage combustion.
- The material will be placed in the test area on top of a hotplate. The hotplate is to provide an indirect heat source of ignition so that the material may smolder before proper igniting.
- The straw will be piled loosely on top of the hotplate while the diesel fuel will be placed in a container on top of the hotplate
- In the event that the hotplate is unable to cause the material to ignite a butane torch will be used to ignite the material.
- The video camera will be left to run as the material goes from smoldering to the point of ignition to the consolidation of the fire to the peak of fire and the exhaustion of the fuel.
- The infrared thermometer will be used at pre-defined intervals on the fire to determine the actual fire temperature.

Third Test: Combined Testing

- The final test will combine the hotspot and smoke/flame testing in order to create the best and most realistic simulation possible.
- The straw will be soaked with diesel and preheated.
- The Steel bar will be preheated and the heated end will be used in an attempt to ignite the mixture.
- The cameras will be used to record the fire in its entirety.
- An Infrared thermometer will be used to record the temperature of the fire and of the steel plate at pre-defined intervals.
- Once the fire has exhausted its fuel, the fire will be extinguished as per safe work procedure.

Fourth Test: Wind Testing

This test is designed to introduce and simulate the effect of wind and air current movement on the ability of the sensor platform to detect fire. Under real world circumstances air current movement will almost always be present. The most common cause of these air current movements are wind in and around the harvester and radiator fans moving air for cooling purposes.

- The flammable material will be placed in the test area on top of a hotplate. The hotplate is to provide an indirect heat source of ignition so that the material may smolder before proper igniting.
- In the event that the hotplate is unable to cause the material to ignite a butane torch will be used to ignite the material.
- To simulate the air current a large industrial type fan will be used, preferentially with flow rate roughly similar to the radiator fans commonly used on harvesters.

- The fan will be placed on one side of the testing area. The speed of the fan will be varied from the slowest setting to the highest in order to help the material ignite.
- The video camera will be left to run as the material goes from smoldering to the point of ignition to the consolidation of the fire to the peak of fire and the exhaustion of the fuel. As the fire goes through each stage the camera will be moved around the test area so that it may view the fire with a head, tail and cross wind.
- The infrared thermometer will be used at pre-defined intervals on the fire to determine the actual fire temperature.
- The sensor package will record the test as per previous procedure. An infrared thermometer will be used to record the fire temperature.

Fifth Test: Sunlight Testing

Sunlight, specifically the thermal radiation from the sun, has the potential to create false positives in the fire detection system. In order for the system to be effective it needs to be able differentiate between infrared radiation from sunlight and the thermal radiation from the fire or hot spot on the machine.

Testing Procedure

- The roller door adjacent to the testing area will be opened, allowing sunlight to fall across the testing area.
- The system is set up on one side of the testing area that is not directly facing the sun.
- The sheet of metal will be placed within the testing area and heated up to a pre-defined temperature using the butane torch.
- This temperature will be confirmed with an infrared thermometer and recorded.
- While the camera is recording the orientation of the plate will be changed, the plate will be moved from 0 degrees (lying flat) to 90 degrees in the vertical plane and rotated about 180 degrees.
- The recording will finish at this point.
- The plate will be reheated to the next hottest temperature point and the test will be repeated.
- Once a range of temperatures have been recorded the test will end and the plate will be left to cool.

Sixth Test: Dust and Chaff Testing

This test was to try and determine the effect of dust and chaff on the effectiveness of the sensor. However, the test was not performed because of these issues

- The straw used to simulate the chaff could not be fed into the fan at a metered rate.
- The path of the straw was unpredictable.
- The straw was somewhat flammable and represented an unreasonable fire risk, especially since it was impossible to contain it to the test area.
- A non-flammable chaff simulant couldn’t be found.
- The dust particulate could not be held in the test area.

The dust/chaff test will be instead attempted during the field testing phase.

4.3 Experimental Software

4.3.1 Programs used

Video Capture software

By using Video Capture software, it enables testing of the hardware side of the system to be carried our independently of the software system. By recording the testing, we can test the alarm algorithm as it is developed and make changes on the fly rather than being forced to work sub optimal programing. It also keeps a record of testing that can be referred to later.

The software that was chosen was Debut Video Capture Software by NCH software. This software not only allows the recording of video directly from the cameras it also incorporates screen capture, meaning that during the later phase of testing, video recordings of live processing can be made. It is only capable of recording from once source at a time so during some tests an additional laptop with another copy of the software had to be used to capture both the feed from both CCD cameras.

The footage from the cameras is recorded at a resolution of 720x576 pixels with a frame rate of 25 frames per second. This footage will be stored on the laptops hard drive and backed up to an external hard drive in case of data loss.

Video processing software

For this task, MATLAB by Mathworks was used. It was chosen over other programming environments.

- **Familiarity.** The biggest reason I chose to use MATLAB was I was familiar with how it worked. I would not have to spend additional time learning a new programming language that would distract me from the goal of the project.
- **Ability**. The MATLAB program environment is ideally suited for the task of processing video and has been used for this task by previous projects. It also incorporates specialised toolboxes for image processing and computer vision that will speed up the creation of the eventual processing algorithm.
- **Adaptability**. Unlike other programing environments, with MATLAB I can run a program with having to compile it first. This allows changes to the algorithm to be made and tested quickly.
- **Compatibility**. MATLAB is compatible with all of the hardware used for this project.

It does however have its drawbacks.

- **Processing Overhead**. Because the program is not complied, it must be running through MATLAB. This puts an additional strain on the laptop and could result in the system not being able to keep up with a live video feed.
- **Future development.** By making use of it forces the use of MATLAB forever more. If this project was to ever progress to a commercial development stage, it's likely that the code would need to be re-written in an open source language such as C.

The version used was 2013b with the computer vision toolbox, image acquisition toolbox and the image processing toolbox.

4.3.2 Software Construction

The original program has branched out into two different but similar programs. The first program is designed to be used with live video feeds. It is able to the video feeds, process them and display the processed video in real time, subject to the processing speed of the computer its running on.

The second program was created to process recorded video for testing purposes. The first version of this program was able also display the process video. Due to issues with MATLAB features that called upon, this original version gave inaccurate results. It fixing this issue, this program lost that ability. It was also dropped to help speed up the overall processing time of each video.

The processing algorithms themselves are identical for both programs apart from minor differences in the way the frame is called to the algorithm. The testing program also has the ability to create and display statistics that will be used to compare the effectiveness of the program under different testing conditions.

Video feed acquisition

The live video program uses the imaq.VideoDevice function to call upon the video feeds. Two video device objects are created in MATLAB, one for the IR feed and one for the colour feed.

The original versions of the testing program used vision.VideoFileReader function to load videos into MATLAB. For whatever reason when this function was used the videos would stop playing roughly halfway through the tests, regardless of the length of the video. What was strange was that the rest of the program would continue to proceed normally but would only return results for the frozen video frame in every subsequent frame in the video.

The program was still useful for fine-tuning the algorithm and providing some of the images used during the testing report, particularly individual frames of interested. But another program would be needed for the testing part.

Instead the VideoReader function was used instead to load the videos into MATLAB. Unfortunately using the VideoReader function dramatically increased the processing time when the system was also attempting to procedure video at the same time. I made the decision to remove the video output component, so that the processing speed could be increased.

Algorithm processing

As discussed in the software methodology section, there were three different methods of fire detection used, colour, NIR emissions and movement. These algorithms were written into self-contained blocks within the overall program. We can see in the flowchart on the next page, the NIR and Colour detection methods were written into self-contained 'detection tracks' that only link up at the end of the program. The movement detection blocks used the already processed vision from the NIR emissions/Colour detection blocks before it. This was done as it improves the processing speed drastically.

The movement detection program used the MATLAB foreground detection program vision.ForegroundDetector, a program in the computer vision toolbox. I found that performing the background subtraction on the raw video feed was extremely inaccuracy as the entire screen would be detected as moving. Changing the settings on the foreground detector did nothing to help.

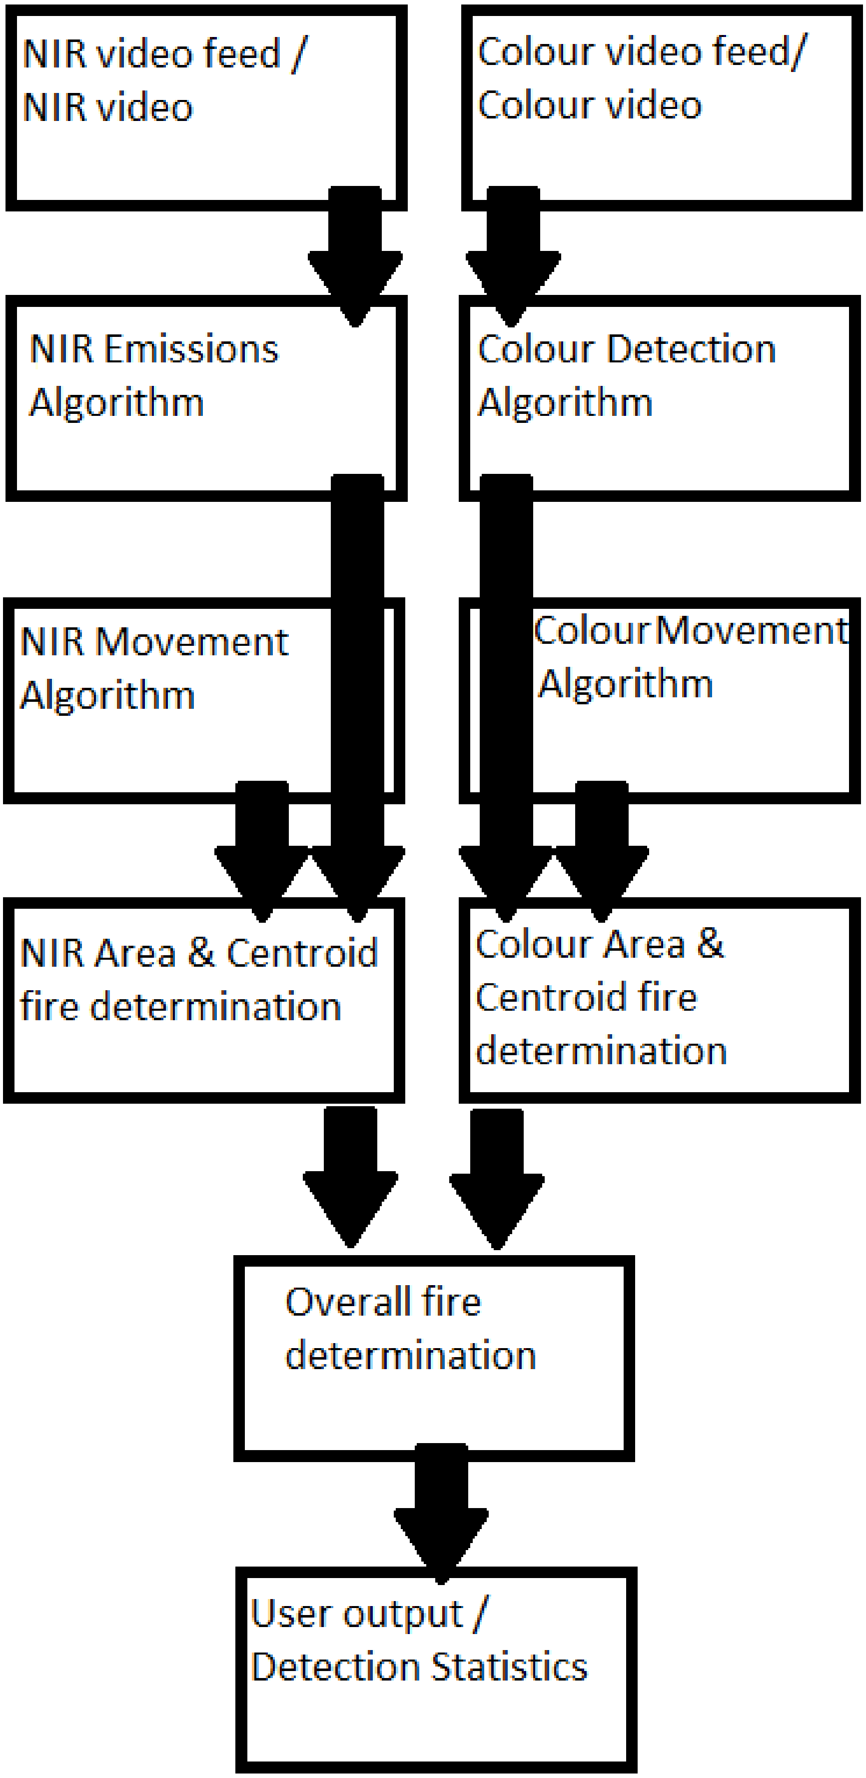

Figure 63: Program organisational flowchart

Filtering

To prevent false positives, morphological filtering processes have been used in each algorithm to remove unwanted noise. The two functions used are bwareaopen and imopen. Bwareaopen removes objects of any shape that do not meet a size criteria and imopen removes objects that also do not meet the correct shape criteria.

Bwareaopen is used in all four algorithm blocks, in the movement detection blocks it is set very low, only 20 pixels, since most of the noise has already been removed in previous steps. In the NIR emission program it is set at 100 pixels, mostly to remove noise in the image. The colour program is set much higher at 800 pixels in area. This done to attempt to filter out the fire reflections and other small objects that may cause false positives.

Imopen is only used on the IR emission algorithm block. It is set to filter out non-disk shaped objects. This to prevent large reflections of NIR radiation from sunlight from causing a false positive the IR emission algorithm.

Fire area and centroid

Once each algorithm block has performed its processing method the resulting image should have a white blob sitting somewhere in a black image. At the end of each algorithm block, the area and centroid of each of these blobs is calculate.

At the area and centroid fire determination, these variables are then used to determine if there is a fire event. There are three possible outcomes from these variable. No fire, possible fire and probable fire.

A no fire determination means that there is no anomaly detected by either method, or an anomaly or anomalies has been detected but they are not of sufficient size or close enough together to warrant a fire detection. A possible or probable fire detection is only differentiated by the size of the object being detected. In this situation, anomalies have been detected in both the emission and movement algorithms and they are close enough to each other and of a sufficient size to warrant further investigation.

If possible or probable fire detections are occurring on both the NIR and colour fire detection tracks, then an overall fire determination is made. This is known as a combined detection. The criteria for a combined detection is the distance between the calculated centre points of the NIR and colour detections. If these detections are occurring close enough to each other the system will set off an auditory alarm to warn the operator. If this is not the case, then it will only display a visual alarm.

NIR detection track

Probable fire: Distance between centroids of NIR emission anomaly and NIR movement anomaly must be less than 200 pixels. Detected fire anomaly size must be bigger than 40000 pixels.

Possible fire: Distance between centroids of NIR emission anomaly and NIR movement anomaly must be less than 200 pixels. Detected fire anomaly size must be bigger than 10000 pixels, smaller than 40000 pixels.

Colour Detection track

Probable fire: Distance between centroids of CC emission anomaly and CC movement anomaly must be less than 300 pixels. Detected fire anomaly size must be bigger than 20000 pixels.

Possible fire: Distance between centroids of CC emission anomaly and CC movement anomaly must be less than 300 pixels. Detected fire anomaly size must be bigger than 10000 pixels, smaller than 20000 pixels.

Combined detection

Must have both a CC and NIR probable detection in the frame. The distance between these probable fire detections must be less than 300 pixels.

User feedback

In the live video processing program, the original video feeds have the detected centre points of any anomalies superimposed back on the image frame before it is displayed to the operator. Text and audio warnings will also be created as needed.

Statistics

In the recorded video processing program. The total number of anomalies/detections is recorded along with when the first of these anomalies/detections occurred. This data will then be used to determine the effectiveness of the system.

4.4 Software Testing

Once the hardware tests have been concluded, the video footage from these tests will be used to assess the effectiveness of the system algorithm software in detecting fires.

The criteria that will be used for determining the effectiveness of the system are;

- Number of frames where a fire was detected
- How many of those detections were true detections of fire.
- Speed of detection.

From the first two criteria we can derive how accurate the system is in detecting fire. As previously stated in the software methodology, the goal is to achieve an accuracy of 50% or better. The last criteria will show how fast we can expect the system to detect any potential fire. Since time is such a critical factor, the algorithm must be able to identify the fire almost immediately. If the algorithm takes longer than a few seconds to identify the fire then it has almost certainly failed.

4.3.1 Equipment and setup

Equipment

- Laptop
- Video processing software

Test Material

- Video Files

Laptop

To run the detection program, the same laptop that was used for the hardware testing will also be used to run the software testing.

Video processing software

As described in the previous section. The processing software was constructed and run in MATLAB.

Video files

These video files will be a mixture of footage recorded from the hardware test phase and footage recorded of situations designed to induce false positives. For each video test there will be two video files, one from the NIR camera and one from the colour camera. Due to the issues that were discovered during the hardware testing, video footage from the colour camera will be substituted for footage from the DRIFT camera whenever possible.

For the footage recorded off the NIR and colour camera, video capture software was used. The software that was chosen was Debut Video Capture Software by NCH software. The footage from the cameras is recorded at a resolution of 720x576 pixels with a frame rate of 25 frames per second. This footage will be stored on the laptops hard drive and backed up to an external hard drive in case of data loss.

The DRIFT Camera was recorded directly to a microSD card in the camera at a resolution of 1920x1080 pixels and a frame rate of 25 frames per second. When this footage was used for these tests it as cropped to the same aspect ratio and resized to 720x576 pixels, to reduce processing time. Whenever possible footage was used from the drift camera in preference to the colour camera due to is superior quality and colour reproduction. The length of the videos varies from 6 seconds to 30 seconds.

4.3.2 Test procedure

The software testing will look at the ability of the system to detect fire and avoid false positive scenarios under various conditions. Two parameters of the program will be varied during testing to determine what effect they have on the ability of the program to detect fire. These parameters are, the colour threshold value Tc in equation 10 and, the NIR threshold parameter Tr from equation 13. The first four testing sets will vary the Tc value from 20 to 80 in increments of 20. The Tr value will be held at 80. The last four testing sets will vary the Tr value from 60 to 90 in increments of 10, the Tc value will be held at 60.

For each of these values eight different scenarios will be ran. The first four scenarios are of video from the hardware tests and as such show fires igniting. The last four scenarios are videos of situations designed to create false positives.

1. Footage of the hotspot testing
   - This video shows a metal bar that has been heated to approximately 400^{o} Celsius. It is being handheld but kept as steady as possible. Over the course of the video it cools. (NIR, Colour)
2. Footage of smoke and flame testing around the time of ignition.
   - A pan of diesel and straw is pre-heated by a hotplate. It is smoking at the start of the video, approximately halfway through it is ignited using a blowtorch. (NIR, Colour)
3. Footage of the wind testing around the time of ignition.
   - A pan of diesel and straw is pre-heated by a hotplate, with a fan blowing across the area. It is smoking at the start of the video, approximately halfway through itself ignites. (NIR, DRIFT)
4. Footage of the 1st sunlight testing around the time of ignition.
   - A pan of diesel and straw is pre-heated by a hotplate. It is smoking at the start of the video, approximately halfway through it is ignited using a blowtorch. (NIR, DRIFT)

5. Footage of a group of people walking.
    - A group of two people walking past the camera. Neither of them are wearing any fire coloured clothing. (NIR, COLOUR)
6. Footage of a reflective object on a sunny day.
    - Static footage of a roof on a sunny day. The camera is slowly panned around. (NIR, COLOUR)
7. Footage of persons in dancing.
    - Two people dancing in front of the cameras, one of them is wearing a fire coloured shirt.
8. Footage of a person in a red shirt with a fire in the back ground.
    - A person in the foreground wearing a red shirt standing still, she then moves, starts a small fire and then starts moving again.

Each test uses two videos, one recorded from the NIR camera and one either the drift camera or the colour camera. The videos are also synced to start and run at the same time and speed. These videos will be run and processed in parallel by the software program.

Prior to each test the video files will be played frame by frame and the total number of frames that have fire in them, as I am able to identify them, will be counted and noted down. The time in the video that the fire starts, in frames, will also be noted down.

For each test the three primary methods of fire detection, movement, colour and NIR emissions, will be used. When the test is running, the algorithms will run a number of counters inside the program. Each time a positive detection is made, the counter will increase by one. Each method of detection will keep a separate counter. For the movement detection method, a separate counter will be kept for each camera feed. In addition, a counter will be kept for the combined detection method. The algorithms will also note the first frame that a detection is made on for each method.

Once the test is complete the data will be compiled into a table for the test. The table will have

- The total number of frames
- The number of true fire frames as determined by a human.
- The number of fire detections from each counter.
- The derived accuracy of each method based on number of detected frames of fire vs number of actual frames of fire.
- The first frame that contained a fire as determined by a human
- The first frame that contained a fire as detected by the algorithm.
- The difference between detection times.

4.4 Field Testing

This final phase of testing will have taken place on a working harvester during the harvest period. The goal of these tests are

- Determine the optimal positions to mount the camera sensors
- Determine the ability of the cameras to function under normal operating conditions and if possible, how the cameras will operate under adverse conditions.
- Determine the feasibility of a UAV platform as a remote camera platform for the system.
- Determine if the system can warn the operator to a potential fire in a reasonable amount of time.

The testing will look at the various positions on the harvester that the cameras can be mounted to provide the optimum coverage of the engine bay and other potential sources of ignition. Specific areas of interest include the turbocharger, exhaust system and the engine. This test will also determine if a camera can be mounted outside of the engine bay on another part of the harvester and still provide an adequate view of the engine bay.

The second part of the testing will observe how the cameras perform under normal operation conditions. This test will not involve the lighting of fires on the harvester, rather it is to see if the system can monitor the engine bay without any false positive alarms caused by changes in illumination, movement of engine parts and other environmental factors.

The third part of the field testing will look at the feasibility of using a UAV platform as a remote camera platform for the system. The main questions that need to be answered are, can the drone keep up with the harvester? Does the dust created by the harvester impede or prevent the drone from flying? Can the cameras get an adequate view of the engine bay or other potential fire locations? Is it practical to have a UAV following the harvester? Will it get in the way of the normal harvesting operations?

The final part of the field test will look at the ability of the system to warn the operator in time. How long does the harvester operator need to get from the cab of the harvester to the fire? Based on this, is the ability of the operator to extinguish the fire improved? Does the extra warning time help him?

From here we should be able to gain an understanding of how effective the system is and if it is worthwhile continuing development of the system.

5. Testing Results

5.1 Hardware Test Results

5.1.1 First Test: Hotspot Testing

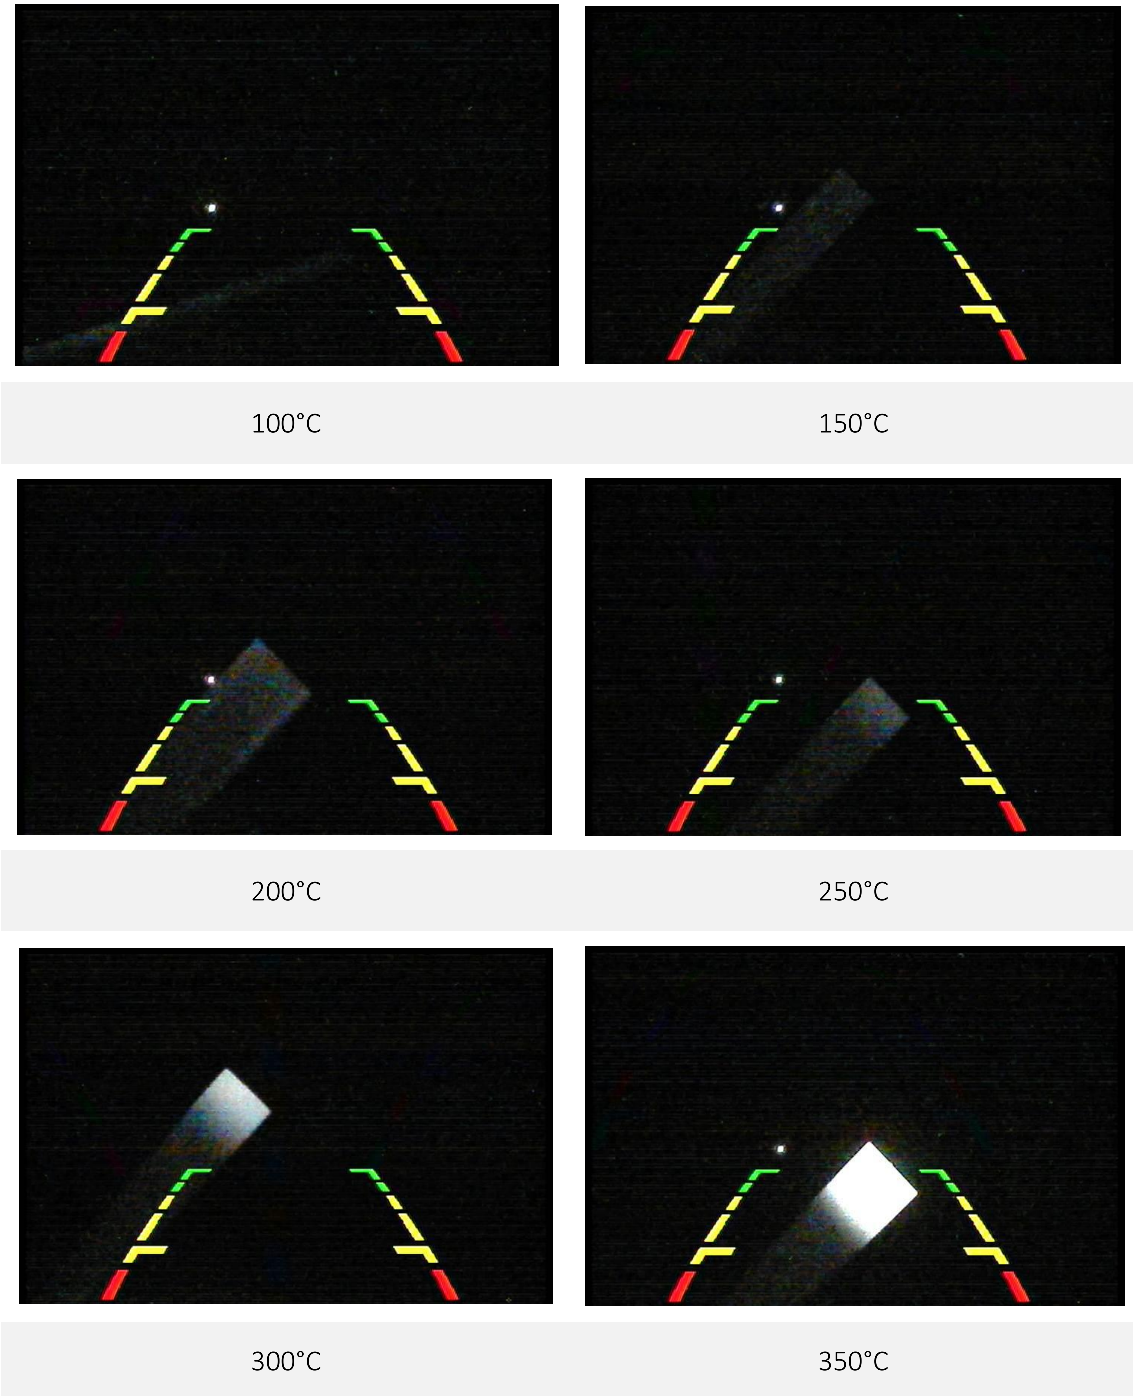

Figure 64: A six frame sequence of the bar being heater up.

This sequence of images captured during the hotspot test gave a good idea of the temperature sensitivity of the camera. We can set that the bit of steel bar used only really became visible once it was heated to the 250-300°C mark. Once this temperature was reached it became extremely luminous when viewed by the NIR camera.

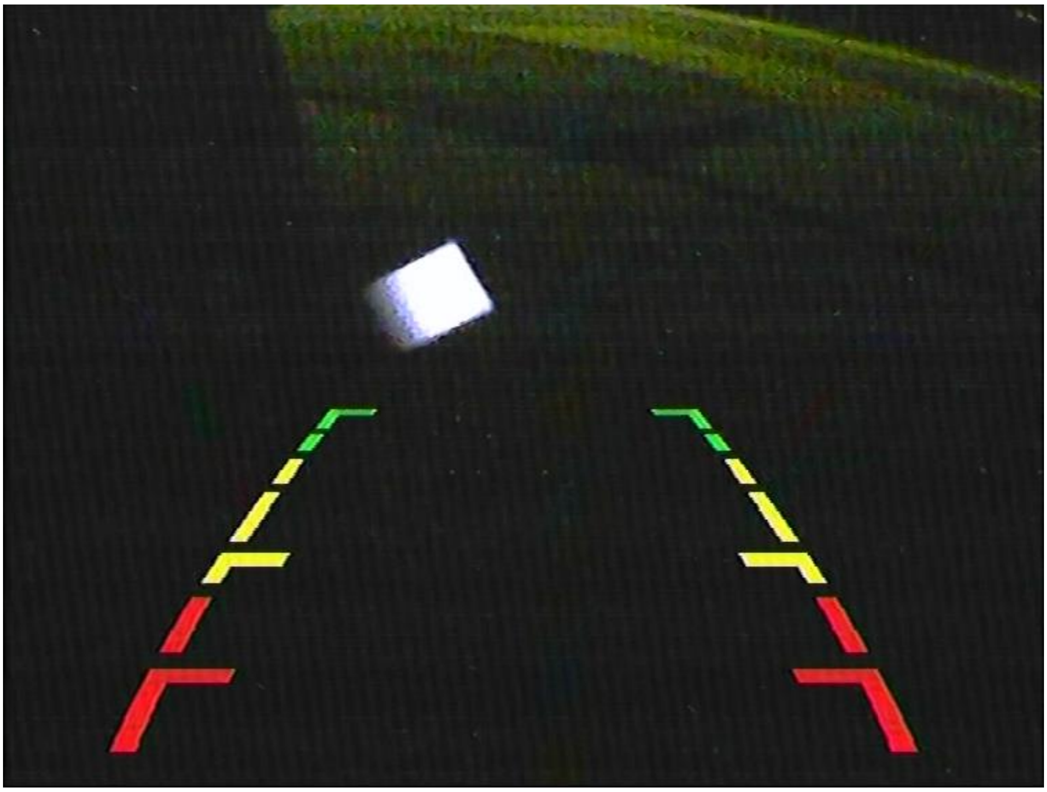

Figure 65: The same bar view through the unfiltered colour camera. We can clearly see the Infrared light being emitted from the heated bar.

The infrared light from the bar was also detected by the unfiltered colour camera. There is a purple colour at the end of the heated section that is not visible in the NIR camera.

5.1.2 Second test: Smoke/Flame Testing

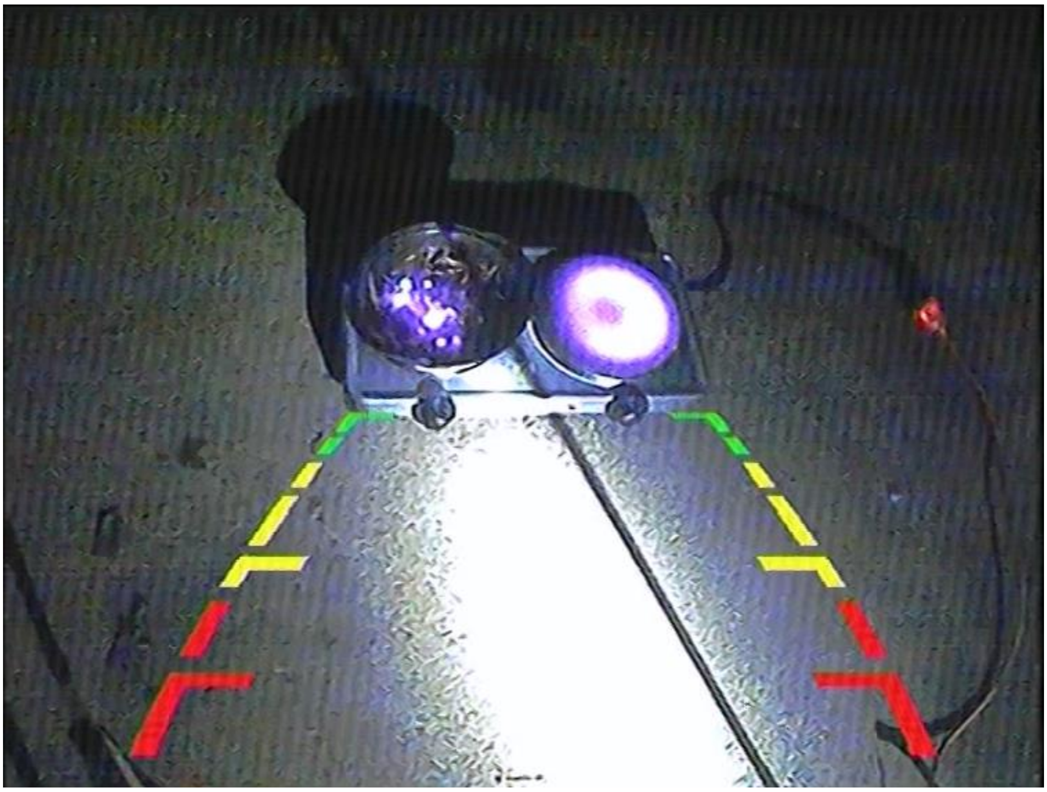

Figure 66: The 1st test attempt, pre-ignition viewed through the colour camera. The straw is starting to char and smoke at this stage. Ultimately the material smoked but did not ignite by its own in this test.

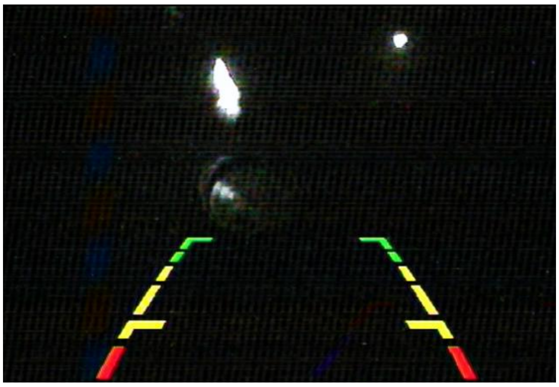

1

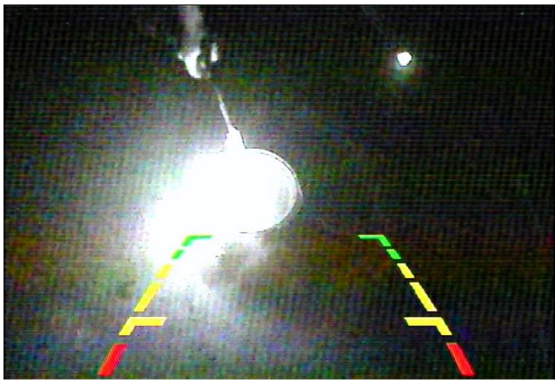

2

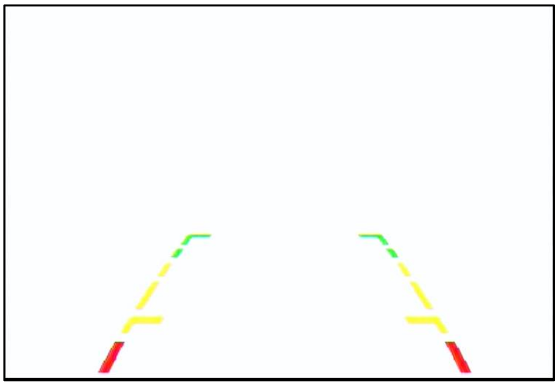

3

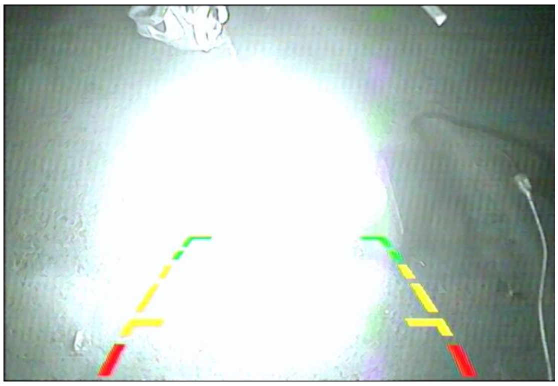

4

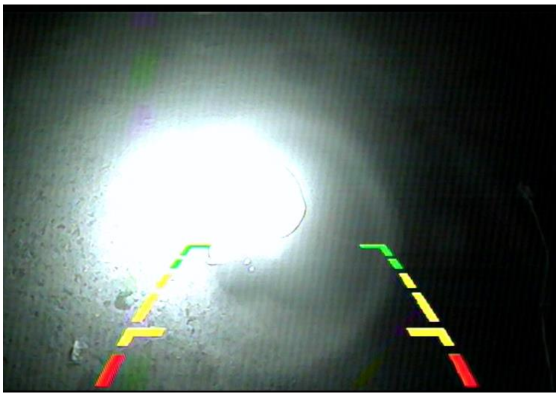

5

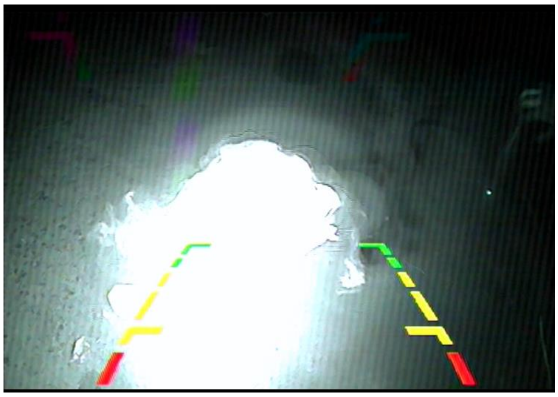

6

Figure 67: A six frame sequence of a pool of diesel being ignited. The first frame shows the flame of the butane torch reflected in the pool of pre-heated diesel. The second shows the very moment of ignition. Frame three shows the camera being overwhelmed by the release of light and energy. The fourth frame is almost 10 frames after frame three where the flame had started to subside. The fifth frame shows the fire subsiding into its normal size. Frame six is about three seconds after the rest and is one of the earliest frames where any sort of flame can be made out.

The infrared light from the fire was both a help and a hindrance. While it made any fire extremely easy to detect it also prevented the flame and smoke from the fire from being viewed. For the future tests, the drift camera was also used in an attempt to capture a view of the flame.

5.1.3 Third Test: Combined testing

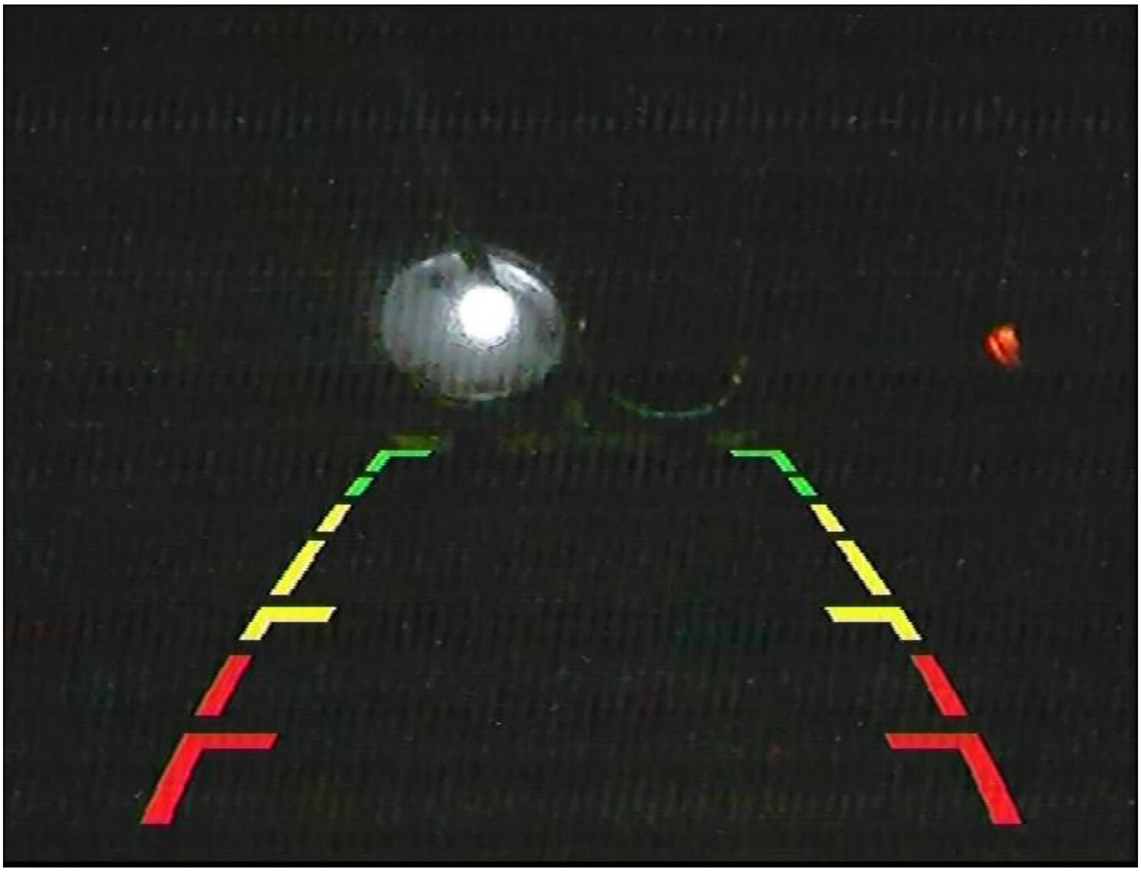

Figure 68: Attempting to light the diesel using a hot object.

Despite our best efforts, it was nearly impossible to get the pool of diesel to ignite using an indirect heat source. It would create a large volume of smoke, but it would refuse to ignite even when subjected to a glowing red hot piece of steel.

Once again we see that the infrared radiation being given off by the hot object overwhelms any visible light that might be captured. Eventually we gave up and ignited the diesel using a propane blow torch. Like in previous tests it ignited immediately when it was subjected to the flame.

5.1.4 Fourth Test: Wind testing

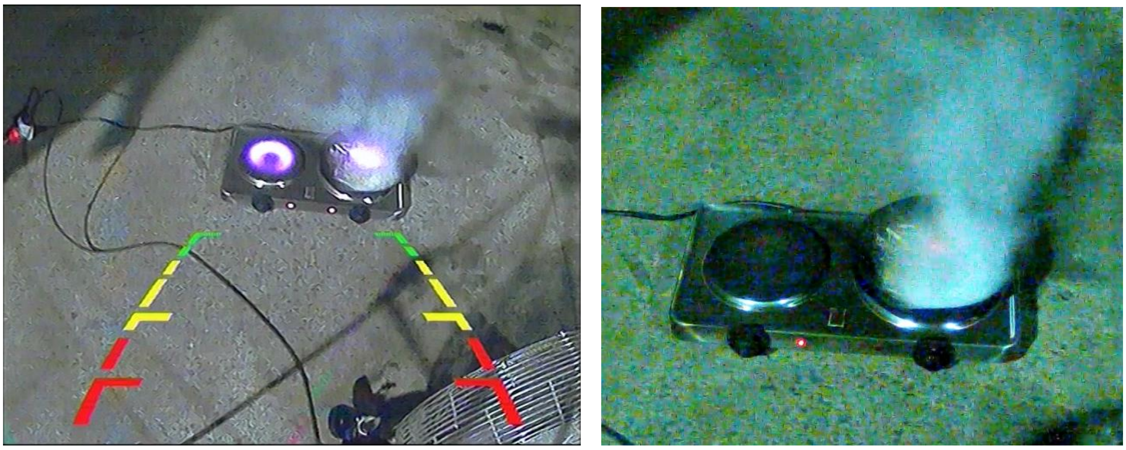

Pre-Ignition

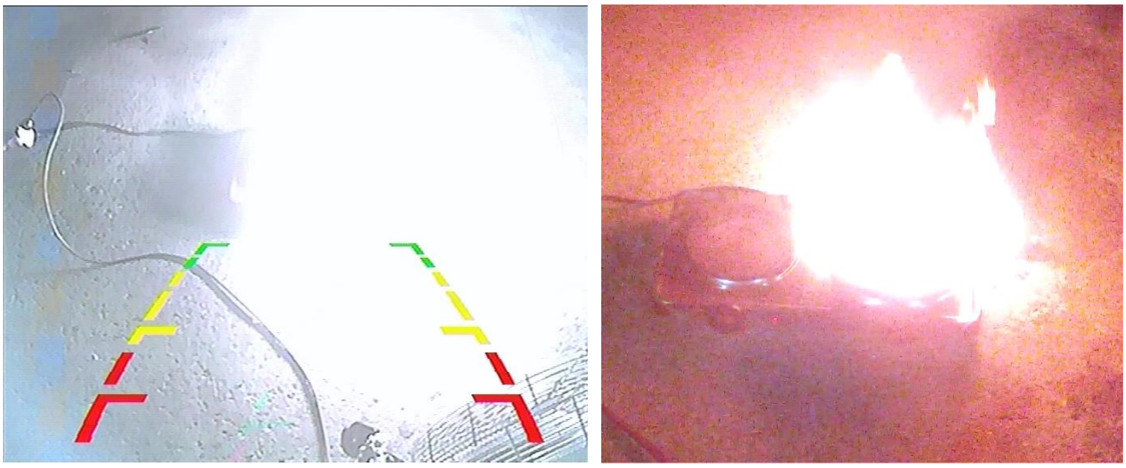

Post-Ignition

Figure 69: The frames on the left are from the colour camera with no filtering. The frames on the right are from the drift camera with a IR filter.

The results of the wind test were very similar to those of the previous tests. The cameras with no NIR filtering exited a large bloom when the upon the initial combustion of the material while the fire on the drift camera was much more subdued. The wind itself had virtually no effect on the ability of the cameras to see the fire.

The combustion material would not auto-ignite if the fan was left at full speed. My guess is that the fanning of the material prevented enough heat to build up to start the fire. In one case the fan was turned off until the material started to smoke profusely. Once the fan was restarted it only took moments before the material ignited and a fire started.

Due to the fan the fire was much more violent than it had been previously and consumed its fuel much more rapidly as a result.

5.1.5 Fifth Test: Sunlight Testing

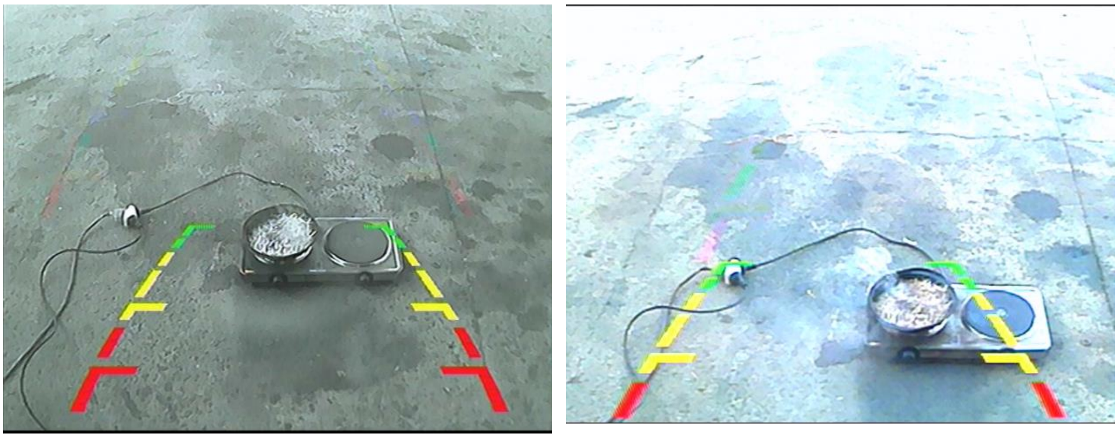

Pre-Ignition

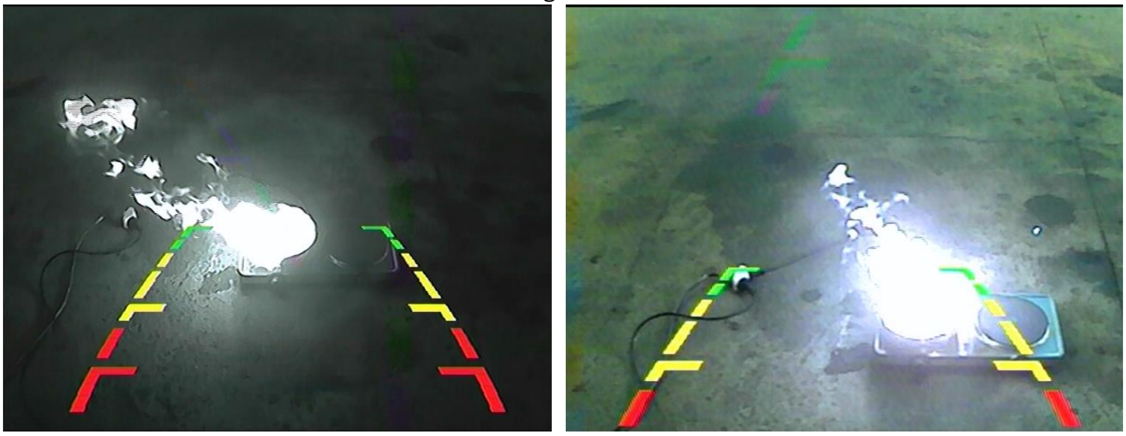

Post Ignition

The test was conducted in the late afternoon in the shade of the industrial shed. Ideally I would have preferred to test under direct sunlight but this was not possible due to time constraints. But the raised level of light across the spectrum seemed to have no effect on the ability of the camera to see the fire. Once the fire was lit it immediately darkened the rest of the frame and was easily detectable by human eyes.

One noticeable difference was the amount of NIR entering the camera. During the night time testing the camera was not able to distinguish any object that was not above 300°C. During the day however the camera was able to see much more detail.

5.1.6 Sixth Test: Dust and Chaff Testing

This test was to try and determine the effect of dust and chaff on the effectiveness of the sensor. However, the test was not performed because of these issues

- The straw used to simulate the chaff could not be fed into the fan at a metered rate.
- The path of the straw was unpredictable.
- The straw was somewhat flammable and represented an unreasonable fire risk, especially since it was impossible to contain it to the test area.
- A non-flammable chaff simulant couldn’t be found.
- The dust particulate could not be held in the test area.

The dust/chaff test will be instead attempted during the field testing phase

5.2 Software Testing

5.2.1 Scenario 1

- Total length (Seconds): 6
- Total length (Frames): 150
- Total fire frames: 0
- Total hotspot frames: 150
- Start of fire frames: N/A
- Start of hotspot frames: 1

This scenario uses video footage from the first hardware test. In this test a steel bar was heated to approximately 400^{o}C and was held in front of the camera sensors. The bar cooled down over the course of the video.

NIR threshold variance testing

Table 6: NIR detection data for scenario 1.

VIDEO 1	INFRARED THRESHOLD (TR)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	150	150	147	107
NIR MOVEMENT	38	52	64	90
NIR POSSIBLE	0	0	0	0
NIR PROBABLE	0	0	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	1	1	1	1
NIR MOVEMENT	112	99	74	43
COMBINED	N/A	N/A	N/A	N/A

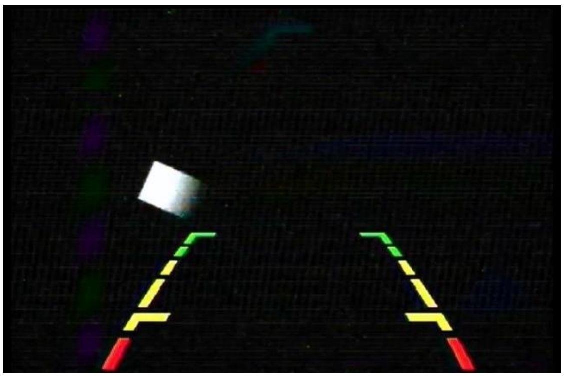

Frame 1

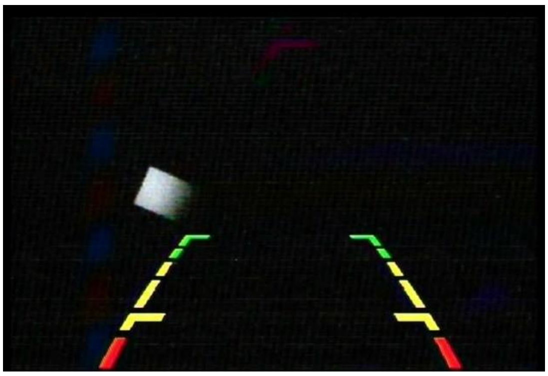

Frame 150

Figure 70: The first and last frame of the video.

The hot object was instantly detected on the first frame of the video in all cases by the NIR emission algorithm, but the length of time that the object was detected for decreased as the Infrared threshold was increased, this decrease is seen in particular on the 90% threshold setting.

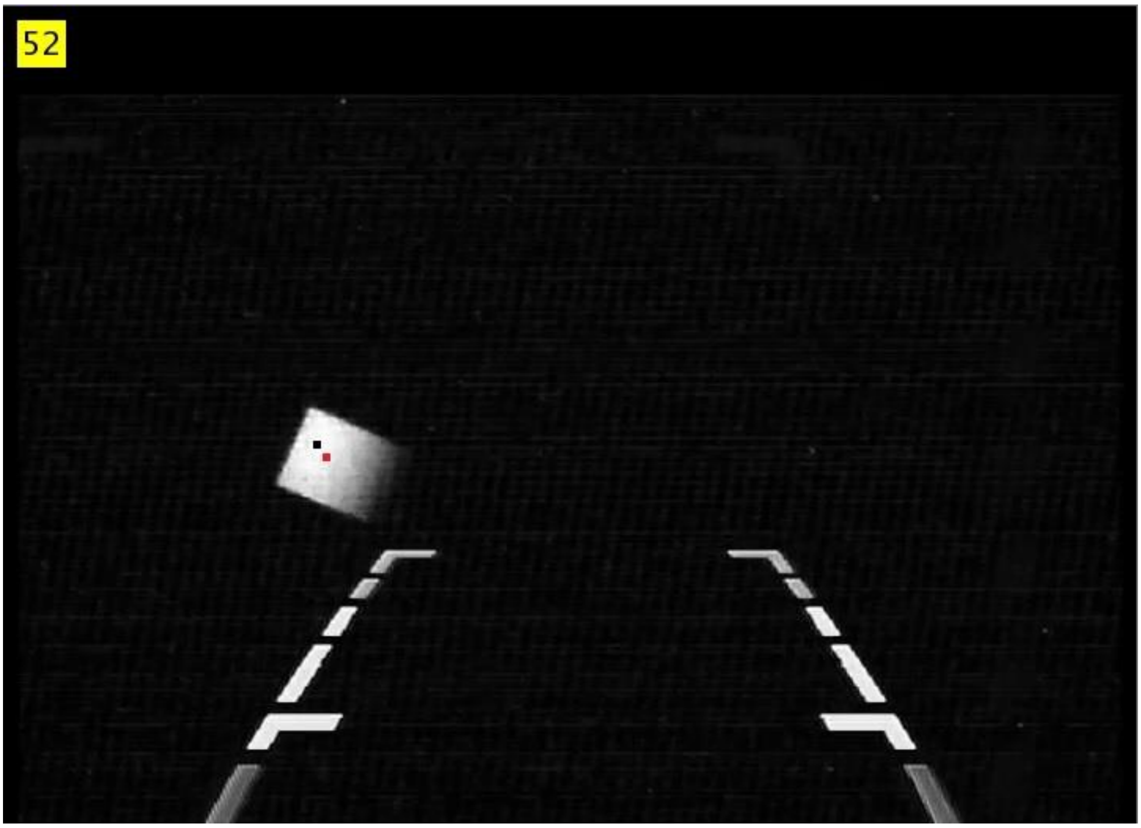

Figure 71: The black dot is centre of the hotspot as judged by NIR emission algorithm while the red dot is the centre as judged by the NIR movement algorithm. This was taken with the NIR threshold at 90% on frame 52. Tr=80, Tc=60.

We can see that the hotspot was located accurately by both the NIR emissions and NIR movement algorithms.

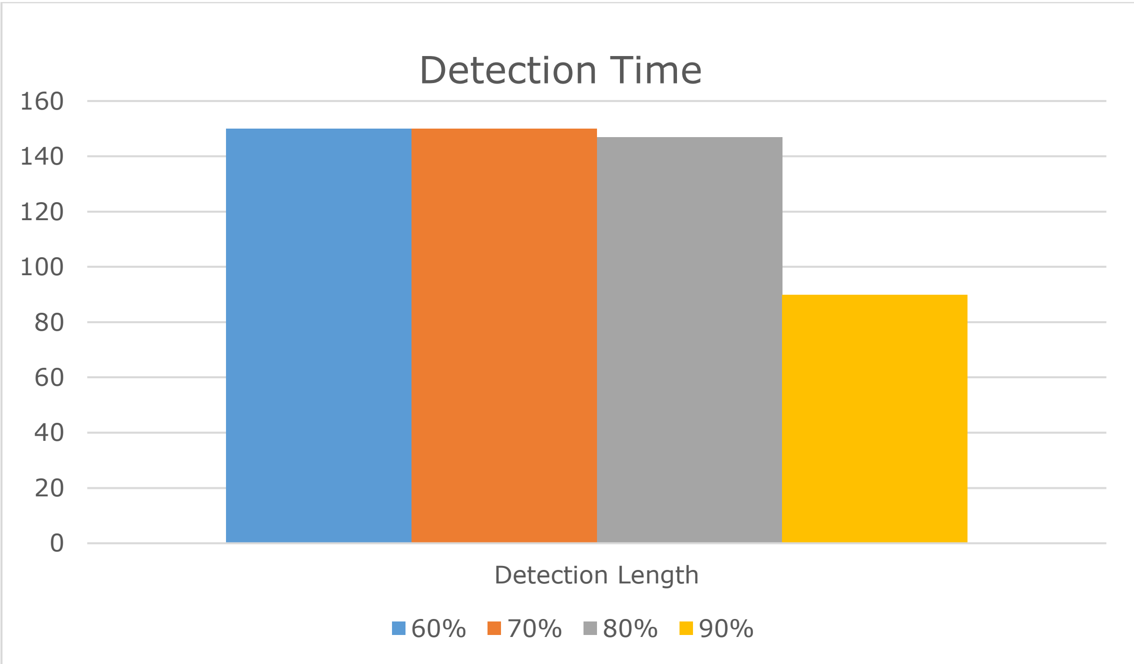

Figure 72: Detections times based on NIR emissions.

What was also very interesting was how the NIR movement tracking reacted to the change in IR threshold. As the threshold increased the amount of NIR movement detections increased and the first detection made was earlier on in the video.

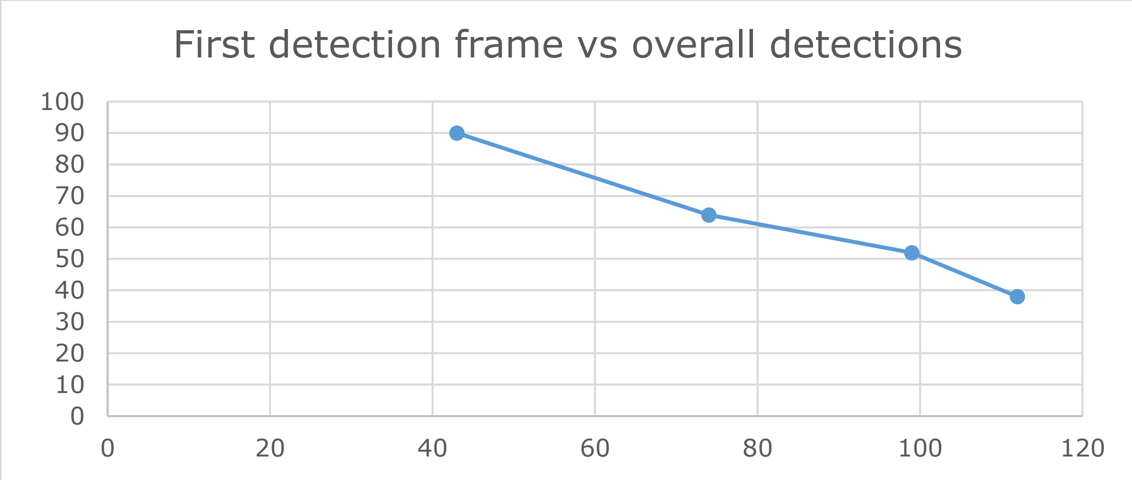

Figure 73: The first fame detected vs the overall frames detected by the movement method

My theory as to why this trend exists is that as the sensitivity of the system is increased the apparent size of the hotspot changes sooner due to the cooling of the hotspot. For a system with a high threshold, the outside parts of the hotspot that are cooling off rapidly in comparison to the centre will be rejected sooner for not meeting the minimum luminance set by the threshold value. Since the hotspot is cooling and decreasing in size during the entire video the hotspot appears to be moving as far as the movement detection system is concerned. For the lower threshold value, the cooling outer parts meet the threshold value for much longer and takes longer for the hotspot to change size.

Colour threshold variance testing

The colour algorithm failed to identify the hot bar. It made no detections at all. I believe that this is due to the camera used. In figure 73 a still image of the colour camera feed is shown. Notice that the piece of steel is glowing the same way as is it in the NIR video. This is because that particular camera did not have any IR blocking filters. Even if it did have these filters it is debateable if this would make any difference as the steel bar was not heated to the point of displaying any colour.

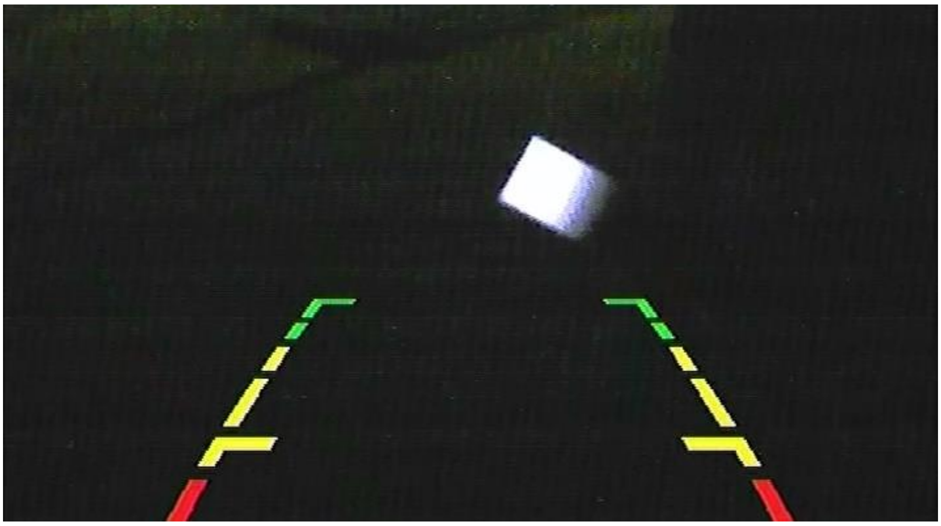

Figure 74: The steel bar viewed under the colour camera.

5.2.2 Scenario 2

- Total length (Seconds): 30
- Total length (Frames): 750
- Total fire frames: 460
- Total hotspot frames: 750
- Start of fire frames: 290
- Start of hotspot frames: 1

The video for this scenario is taken from the second hardware test. It shows a close in view of a plate of straw and chaff being heated on a hotplate. At the start of the video the hotplate has already heated the material to smoking point. Around the 12 second mark a slight blow of air causes the material to ignite and start to burn.

NIR threshold variance testing

Table 7: NIR detection data for scenario 2

VIDEO 2	INFRARED THRESHOLD (TR)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	750	750	750	750
NIR MOVEMENT	724	724	723	719
NIR POSSIBLE FIRE	236	235	247	240
NIR PROBABLE FIRE	468	455	443	400
COMBINED	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	1	1	1	1
NIR MOVEMENT	2	2	5	26
NIR POSSIBLE FIRE	2	2	35	49
NIR PROBABLE FIRE	177	269	273	277
COMBINED	N/A	N/A	N/A	N/A

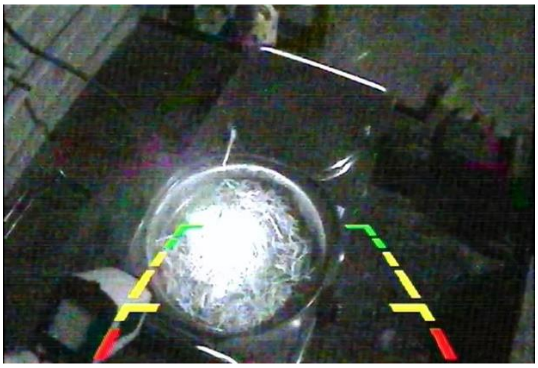

Frame 1

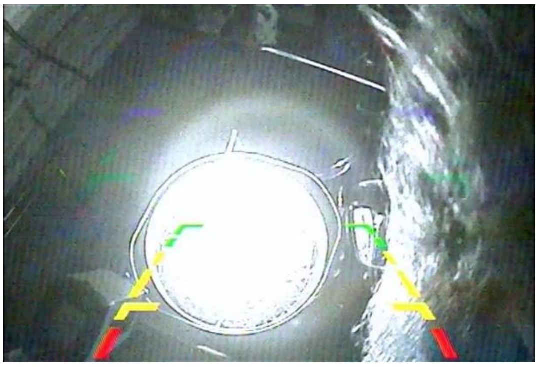

Frame 290

Figure 75:The start of the video and the perceived start of the fire as viewed by the NIR camera.

Since the video starts with an already hot object in it, we can see that there were NIR emissions detected throughout the entire video, we can also see that movement detection started very early on too. Both of these managed to located the fire with reasonable accuracy within the video frame.

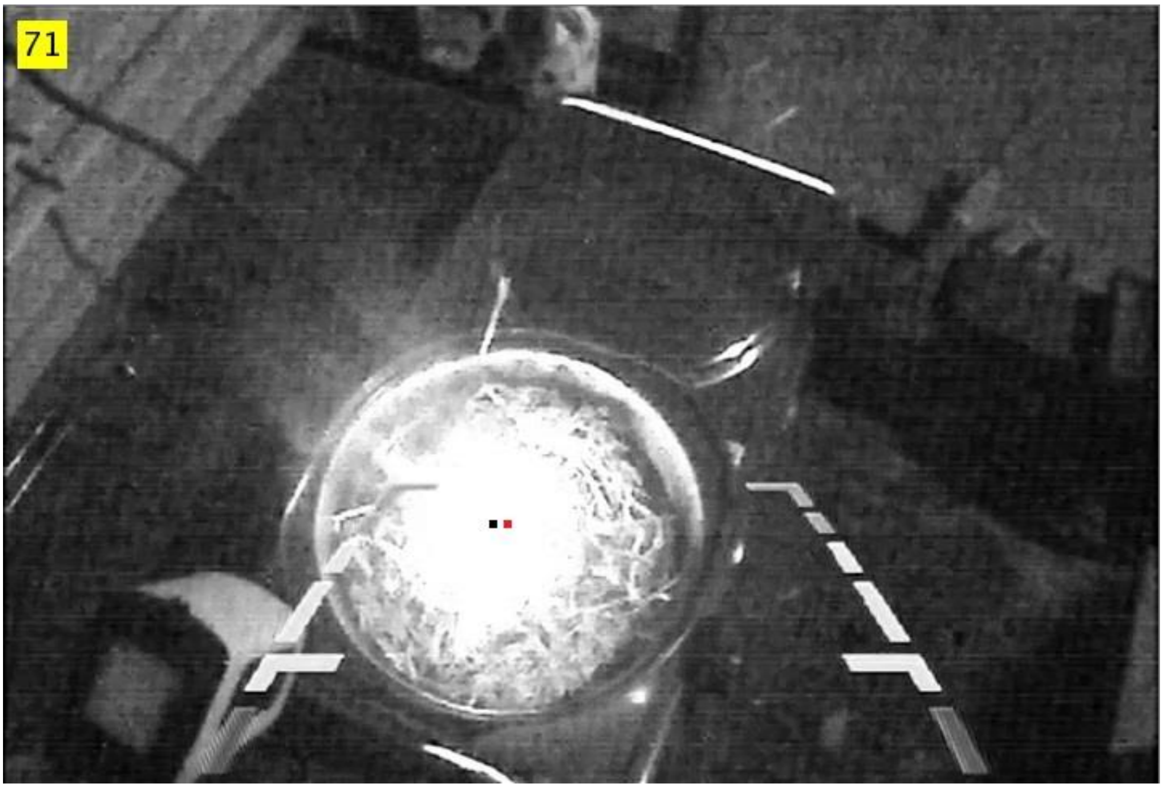

Figure 76: Frame 69 and 71 of the video. Tr = 80 and Tc = 60

The black dot marked the centre of the fire as determined by the NIR emission algorithm, the red dot marked the centre as determined by the NIR movement algorithm. Throughout the video, the red dot moved around the frame much more.

Since there was already a significant heat bloom in the first frame, the program registered that there was a potential fire right from the start in the 60 and 70% threshold tests. Even for the 80 and 90% Tr threshold tests, a potential fire was detected within the first 50 frames.

Probable fire detection in most cases was made just before the defined beginning of the fire in all cases except for the 60% Tr threshold which made a probable fire detection just over 100 frames before the fire started.

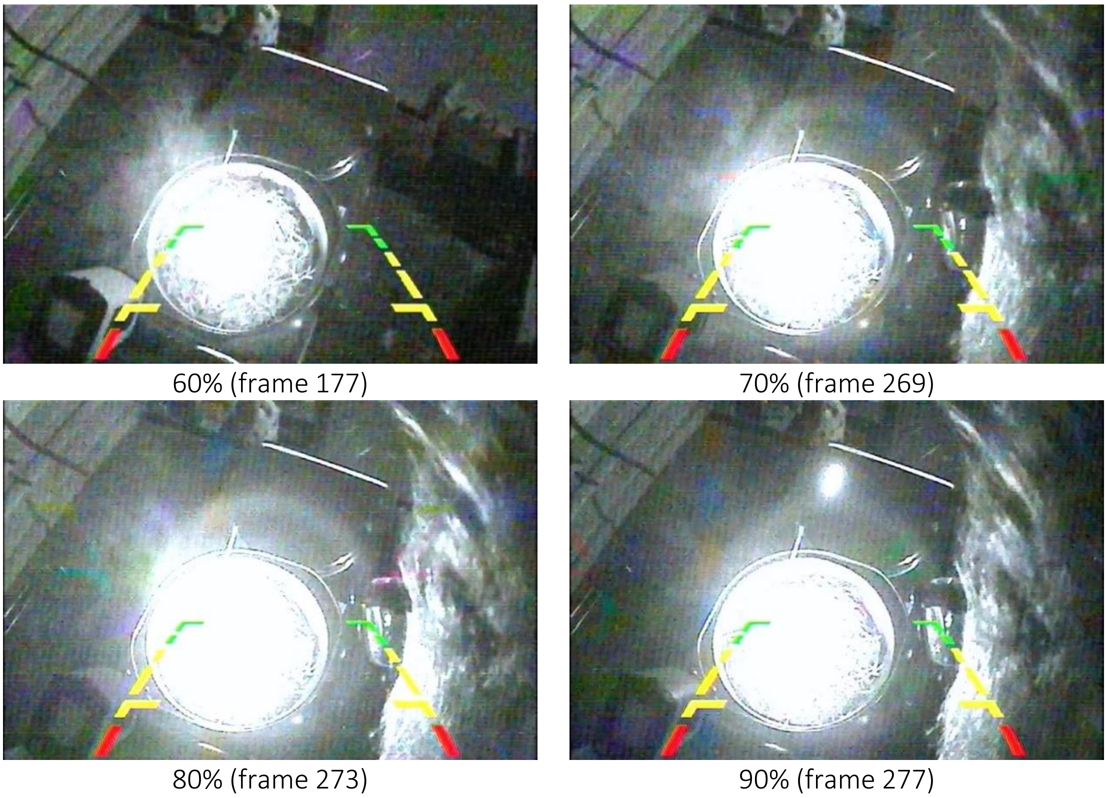

Figure 77: The first probable fire detection for each threshold.

Colour Testing

Table 8: Colour threshold results for scenario 2

VIDEO 2	COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	20	40	60	80
CC EMISSIONS	252	48	6	0
CC MOVEMENT	44	35	6	0
CC POSSIBLE	0	0	0	0
CC PROBABLE	0	0	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	20	40	60	80
CC EMISSIONS	1	34	35	N/A
CC MOVEMENT	31	34	35	N/A
CC POSSIBLE	N/A	N/A	N/A	N/A
CC PROBABLE	N/A	N/A	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

The colour camera seemed to struggle with this test, once again it was linked with IR emissions from the camera preventing the colour detection algorithm from seeing the fire.

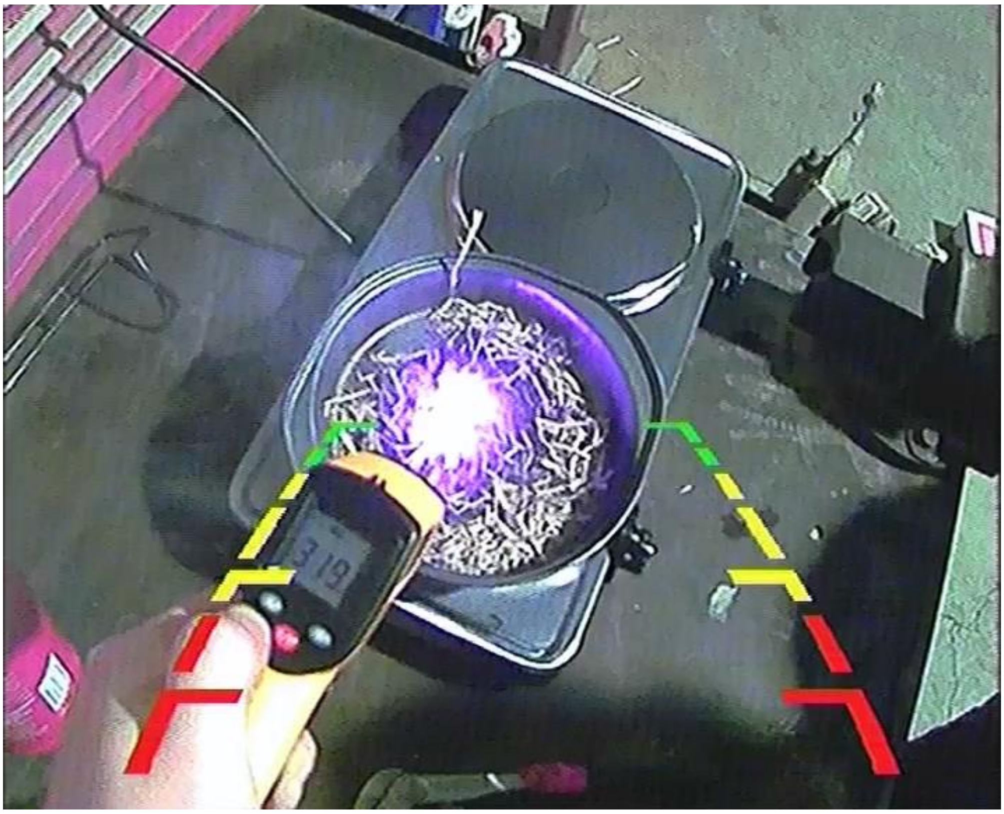

Figure 78: The start of the second scenario view on the colour camera. Once again the detected IR emissions overwhelm any fire colour.

Despite this limitation we can see that some detections were made especially on the lower threshold settings. None of these were big enough to cause a possible or probable fire detection. No combined detections were made during this scenario.

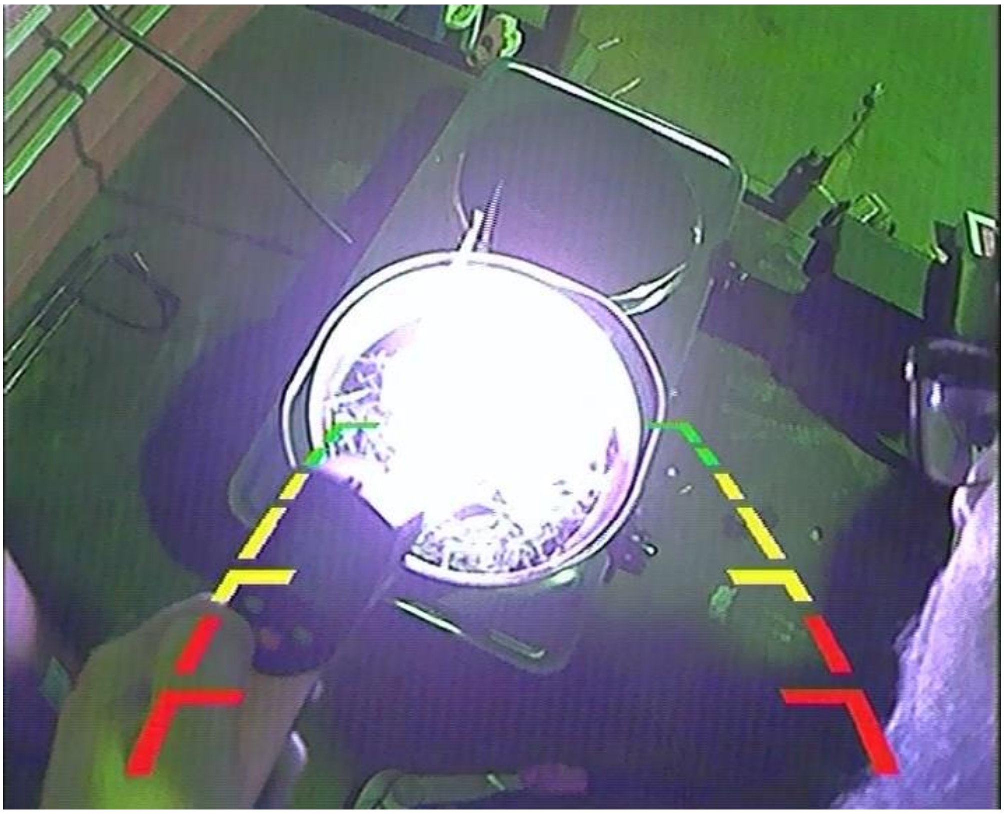

Figure 79: The point of ignition. Once again the IR bloom makes any colour detection method useless.

5.2.3 Scenario 3

- Total length (Seconds): 30
- Total length (Frames): 750
- Total fire frames:385(NIR)
- Total hotspot frames: 420
- Start of fire frames: 365 (NIR), 391(Colour)
- Start of hotspot frames: 330

The video for this scenario is taken from the fourth hardware test. It shows a hotplate heating up a pan of diesel and straw. There is a fan set approximately half a meter away blowing air across it. At the halfway mark the mixture is ignited using a blowtorch and burns for the rest of the video.

NIR threshold variance testing

Table 9 NIR detection data for scenario 3

VIDEO 3	**INFRARED THRESHOLD (TR)**			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	398	397	395	391
NIR MOVEMENT	393	391	391	391
NIR POSSIBLE	4	10	28	49
NIR PROBABLE	378	355	332	314
COMBINED	67	81	77	60
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	312	312	312	312
NIR MOVEMENT	44	312	312	312
NIR POSSIBLE	364	444	406	406
NIR PROBABLE	391	365	365	365
COMBINED	405	405	405	405

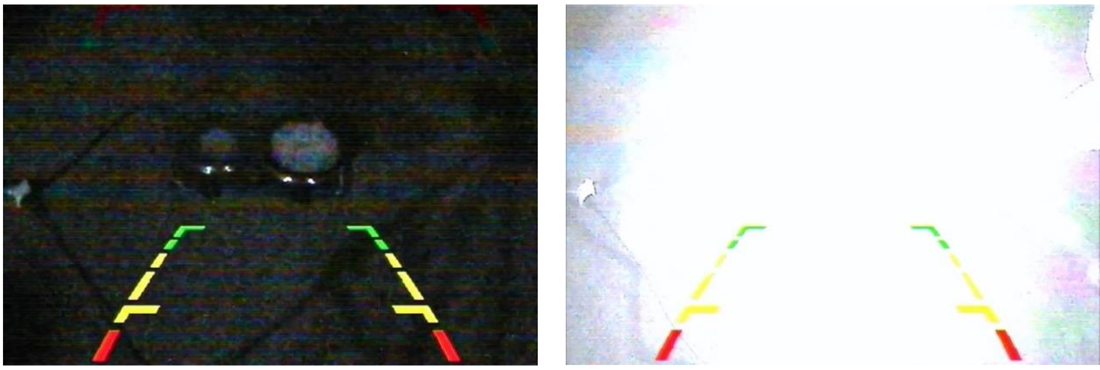

Frame 1 Frame 365

Figure 80:Scenario 3 at start and at point of ignition.

This scenario is different from the previous scenarios as it is further away from the fire, does not start with a large thermal hotspot in the frame and is ignited using a flame rather than under its own heat. Under these conditions the system behaves quite differently.

In this scenario, there was not an emission detection for every frame. The first recorded NIR emission takes place on the 312 frame, 53 frames or almost 2 seconds before the point of ignition. Apart from the 60% Tr threshold test, the first movement detection also took place on frame 312.

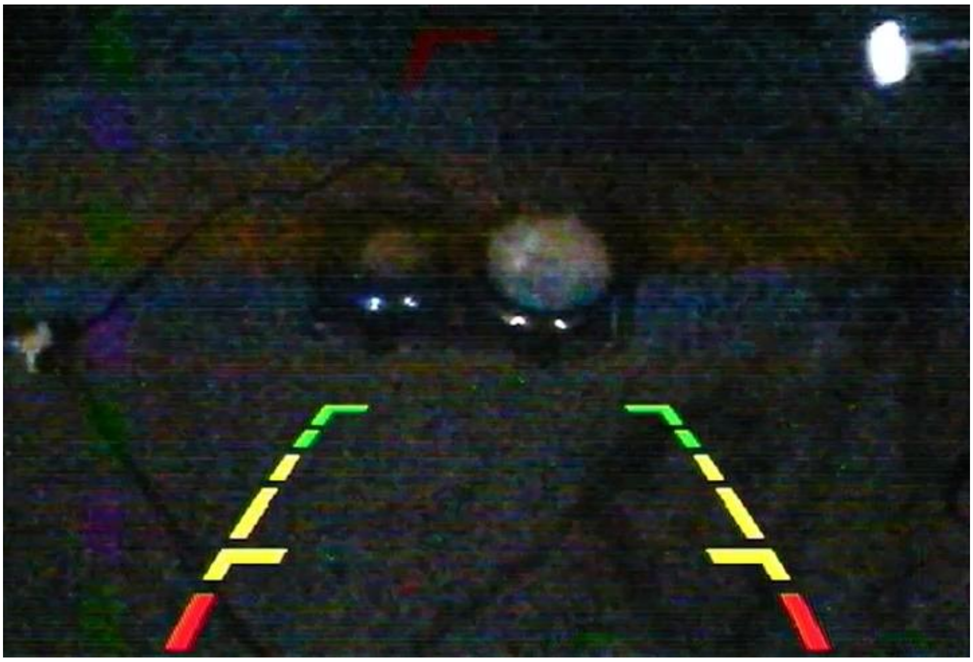

Figure 81: Frame 312 of scenario 3. Notice the flash of the blowtorch in the upper right corner.

We can see that this was cause by a flash of the blow torch as it is being ignited on the frame. This small flash was not enough to cause the program to register a possible or probable fire.

However, the not so small flash on frame 365 was more than enough to cause a probable fire tag for all of the modes except the Tr 60%

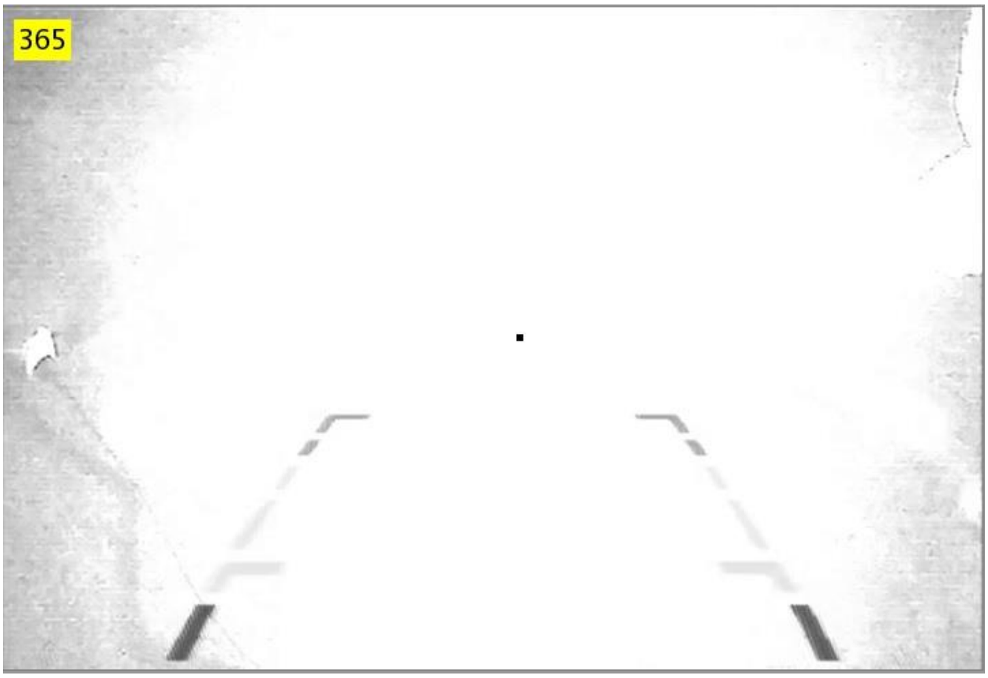

Figure 82: Moment of ignition as seen by the program. Both the NIR emission and NIR movement dots mark exactly the same point.

As far the amount of detections made both NIR emission detections and NIR movement detections varied very little as the threshold was changed. The possible and probable fire detections did vary somewhat, mostly as the threshold level was increase a number of frames were demoted from probable to possible detections.

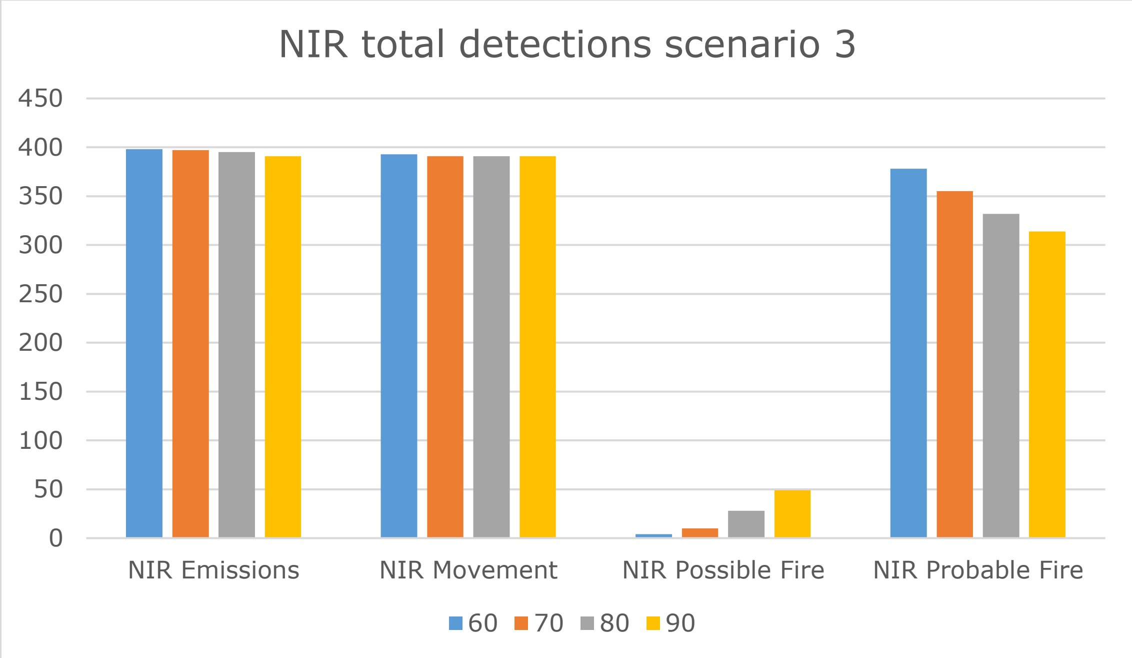

Figure 83: NIR total detections for scenario 3

Colour threshold variance testing

Table 10: Colour detection data for scenario 3

VIDEO 3	COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	20	40	60	80
CC EMISSIONS	543	338	232	105
CC MOVEMENT	325	179	144	99
CC POSSIBLE	27	20	12	17
CC PROBABLE	116	108	88	15
COMBINED	96	92	77	10
FIRST DETECTION FRAME	20	40	60	80
CC EMISSIONS	1	383	391	406
CC MOVEMENT	2	383	391	406
CC POSSIBLE	104	397	395	446
CC PROBABLE	391	391	405	447
COMBINED	393	393	405	480

This scenario marked the first use of the DRIFT camera instead of the cheap colour camera that had been used in previous tests. This camera with its IR cut filter had the ability to be able to see the flame colour, this is reflected in the results gathered from testing.

Frame 1 Frame 391

Figure 84 The start frame and ignition frame of the colour video. An error in editing meant that there was a syncing issue.

Unfortunately, these could have been better. An error was made when the videos were being edited down to the 30 second length and the two videos are almost one second out of sync with each other. We can see in the results of the colour testing that apart from Tc=80 test all of the Probable fire detections were about 25-30 frames behind the probable fire detections made by the NIR algorithm.

We can see in any case that with the new camera the colour algorithm was able to detect the fire only a frame or two after ignition. What was interesting was how the colour algorithm 'saw' the fire.

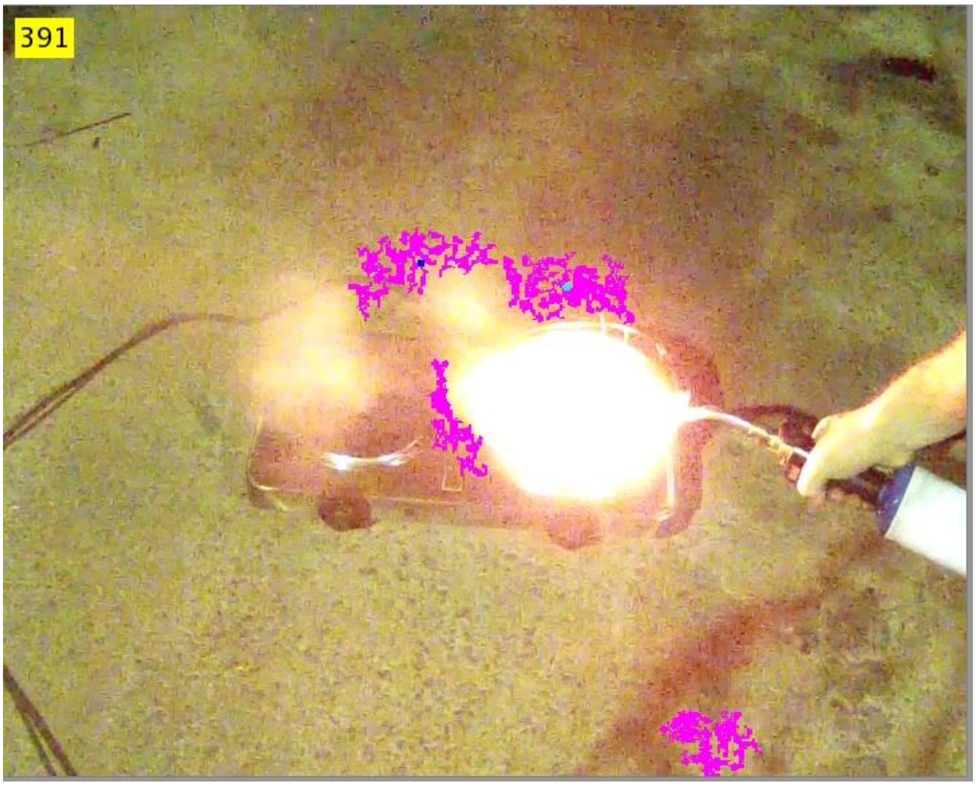

Figure 85: Frame 391, the shaded areas is where the colour detection algorithm has found fire. The blue dot is the centre of this while the black dot is the detected movement centre. In this instance Tc=60, Tr=80

We can see in figure 84 that even though a fire had been detected, what the algorithm was actually seeing was the 'halo' of the flame rather than the centre of the fire itself. It is also seeing light reflections off the concrete and off the hotplate. We can also see that the detection dots, based on area and centroid are quite far apart from each other. This is why it takes roughly ten frames for the system to detect a probable fire.

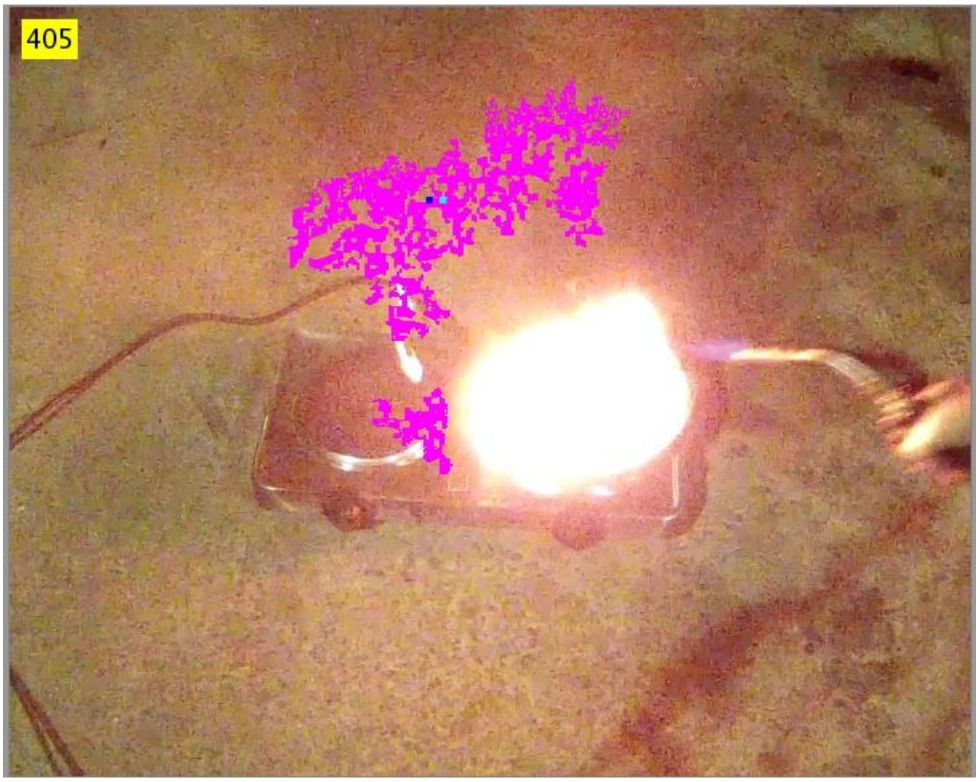

Figure 86: Frame 405, the dots are closer together and the 'blob' is larger, enabling a probable fire identification.

Otherwise the detection algorithm responds to an increased threshold level by reducing the amount of detected fire frames. The Tc=80 test did have a massive drop off in detections in comparison to the other methods used. It also detected the fire much later than the other thresholds, a full 80 frames or 3.2 seconds after the rest did.

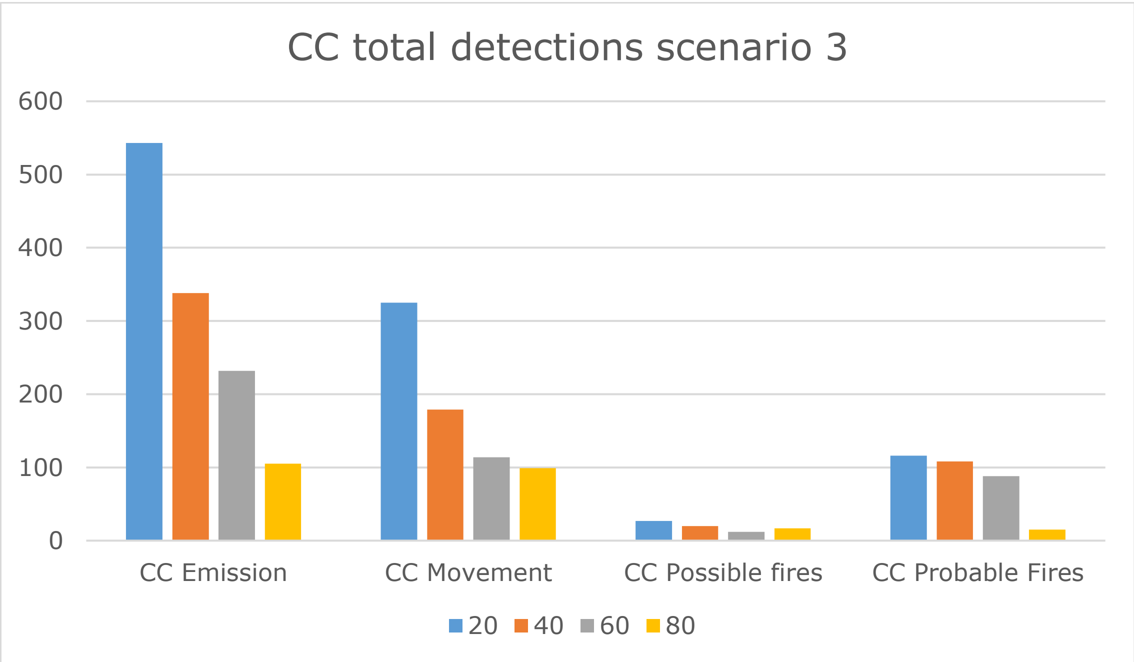

Figure 87: Total detections on the colour camera for scenario 3.

5.2.4 Scenario 4

- Total length (Seconds): 30
- Total length (Frames): 750
- Total fire frames:370
- Total hotspot frames: 370
- Start of fire frames: 380
- Start of hotspot frames: 380

The video for this scenario is taken from the fifth hardware test. It shows a hotplate heating up a pan of diesel and straw. There is a fan set approximately half a meter away blowing air across it and there is a large amount of ambient light from the sun. At the halfway mark the mixture is ignited using a blowtorch and burns for the rest of the video.

NIR threshold variance testing

Table 11: NIR detection data for scenario 4

VIDEO 4	**INFRARED THRESHOLD (TR)**			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	750	728	625	449
NIR MOVEMENT	748	745	578	433
NIR POSSIBLE	249	190	202	241
NIR PROBABLE	219	181	144	109
COMBINED	145	119	100	69
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	1	1	1	28
NIR MOVEMENT	3	3	3	104
NIR POSSIBLE	3	96	306	323
NIR PROBABLE	310	310	374	371
COMBINED	391	504	506	512

This is the first scenario that was filmed during the day. Because of the increased amount of ambient light, both visual and IR, it is no longer as easy to see hotspots. In the first frames we can see that the camera cannot see the hotspot in the centre of the pan where the diesel and straw is being heat.

But because of the increased about of ambient light, the quality of the video has improved dramatically there is no longer a large amount of noise in the image. Also we can now see the smoke rising off the hot plate much more clearly than before. In the fire ignition frame we can clearly see the large thermal bloom given off by the initial burst of fire. The Tr=60 and Tr=70 threshold tests were much more sensitive to the increased amount of ambient light than the Tr=80 and Tr=90 threshold tests were, making their first detections much earlier in the video.

But apart from the Tr=60 test, the first combined detection was made at around frame 500, almost five seconds after the start of the fire.

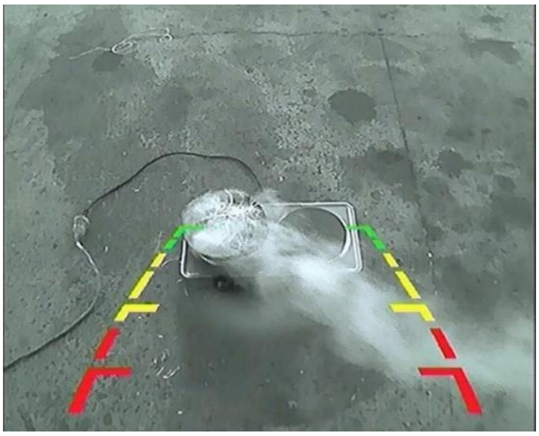

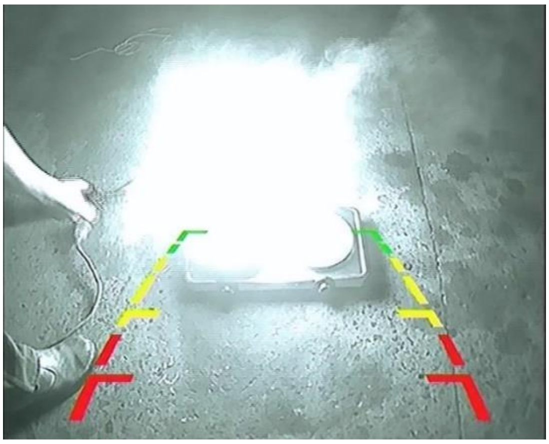

Frame 1 Frame 280

Figure 88: Scenario 4 at the start frame and at the ignition frame.

Figure 89: Scenario 4 at frame 512, no issue seeing the flame here.

Once again we see that the program has no issues with tracking the flame once it is large enough to meet the probable or possible fire criteria.

The total number of detections made exhibited a downward trend as the Tr threshold was increased except for the NIR possible fire criteria, there was an increase in detections for Tr=90.

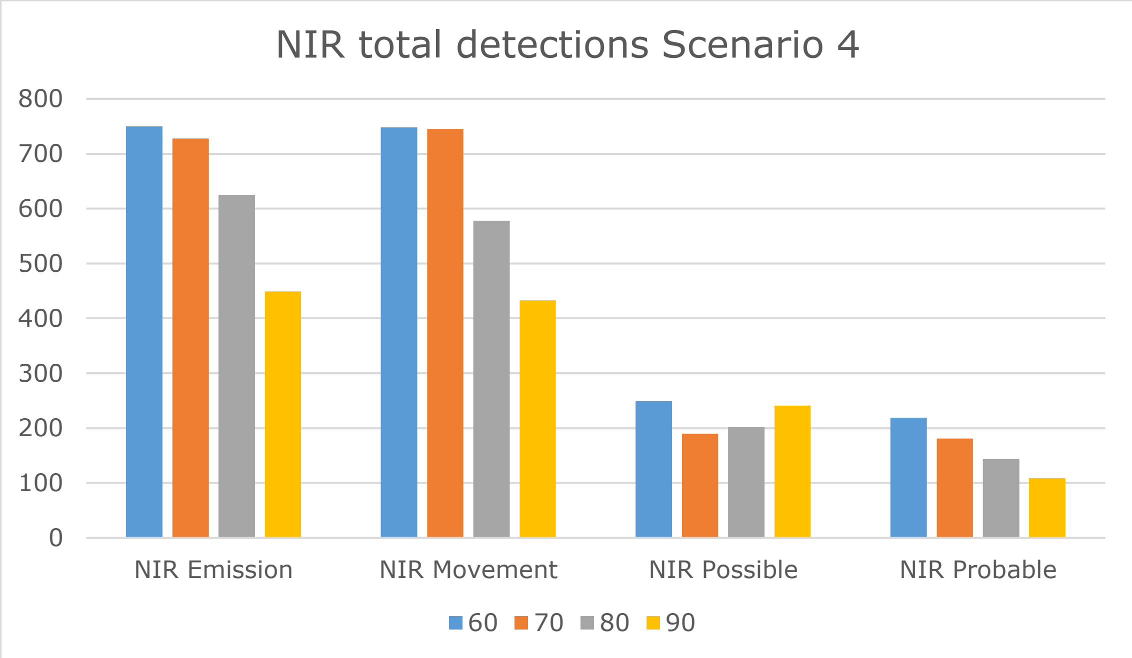

Figure 90: NIR total detections for Scenario 4

Colour threshold variance testing

Table 12: Colour detection data for scenario 4.

VIDEO 4	COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS	20	40	60	80
CC EMISSIONS	371	370	370	366
CC MOVEMENT	283	311	328	332
CC POSSIBLE	10	16	39	55
CC PROBABLE	247	266	253	221
COMBINED	86	95	100	97
FIRST DETECTION	20	40	60	80
CC EMISSIONS	379	381	381	383
CC MOVEMENT	379	381	381	383
CC POSSIBLE	381	383	387	394
CC PROBABLE	383	386	391	404
COMBINED	383	386	512	505

The colour detection and colour movement algorithms were able to identify the fire almost as soon as it started with the worse performing threshold value still taking less than a second to flag the anomaly as a probable fire. The Tc=60 and Tc=80 tests did not have a combined detection until the 500 frame mark unlike Tc=20 and Tc=40 which suggests that the Tc threshold was set a little too high for the NIR tests.

 

Frame 1 Frame 380

Figure 91: Scenario 4 at the start frame and the ignition frame.

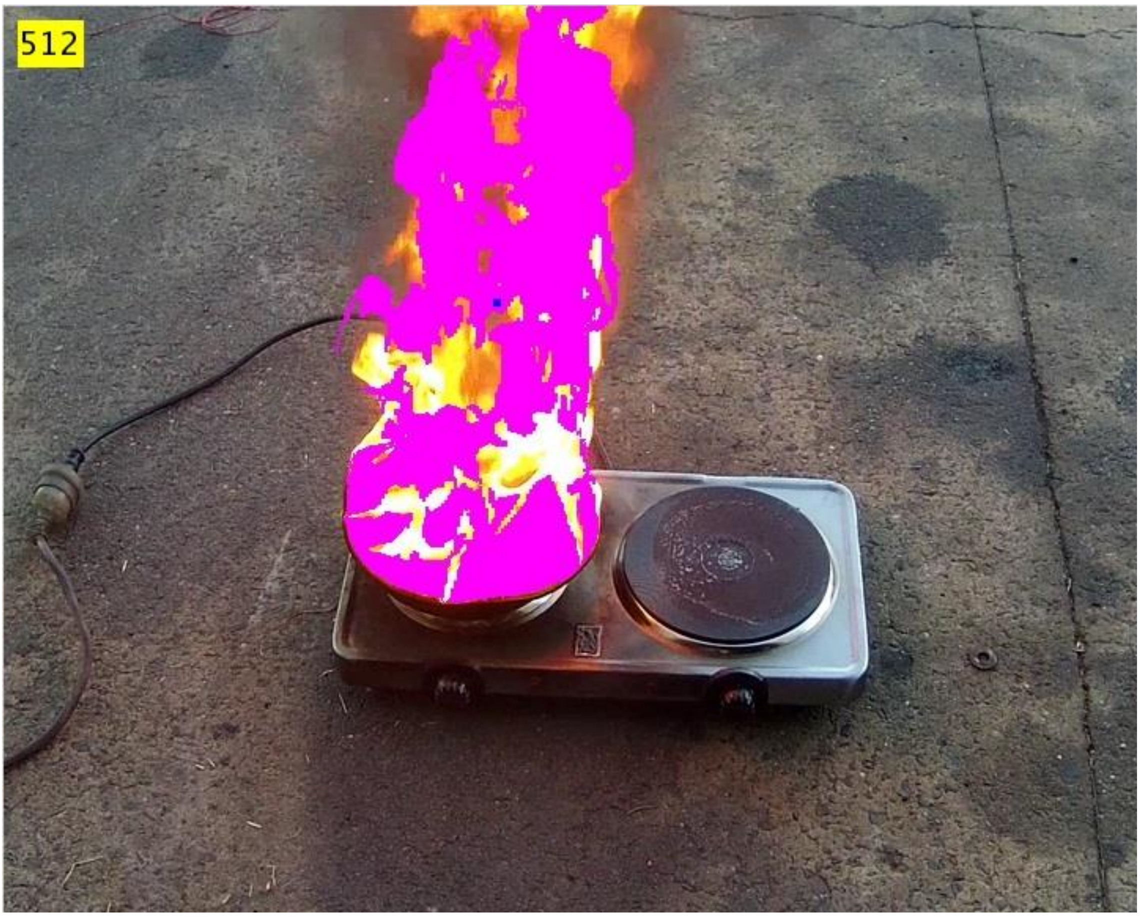

Figure 92: Frame 512 under normal light. The shaded area indicates where fire colour has been found and the blue dot indicates the detected centre of the fire. Settings were Tc=60 and Tr=80.

The performance of the colour algorithm improves quite significantly when it is used in day time conditions. The increased amount of ambient light allows the camera to see the colour of the flame instead of a bright white bloom. We can see in figure 92 that this time the algorithm has marked the flame proper instead of reflections around the flame and has placed the centroid do right in the centre of the flame.

The first detection values apart from the first combined detection remained fairly static throughout the testing. The total number of detections also varied somewhat during the testing.

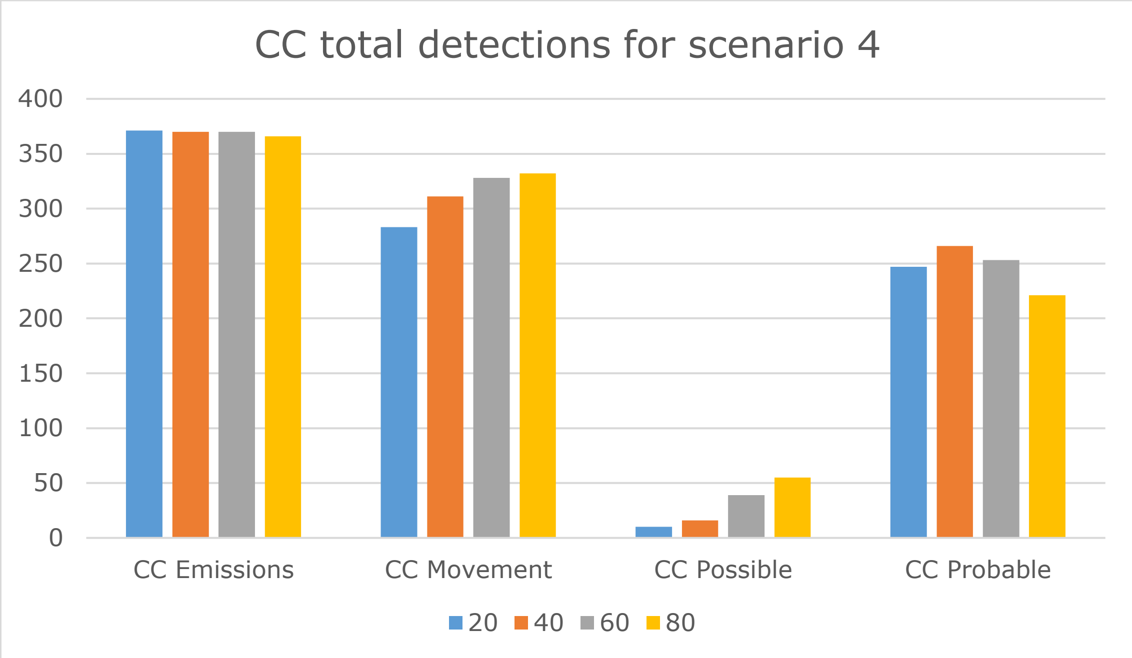

Figure 93: Total detections for the colour camera in scenario 4.

CC emissions detected remained mostly static, but movement detections increased with a higher Tc value. The possible and probable detection did a normal 'swap' of detections as fire frames were demoted for not meeting the probable fire criteria.

5.2.5 Scenario 5

- Total length (Seconds): 15
- Total length (Frames): 375
- Total fire frames: 0
- Total hotspot frames: 0
- Start of fire frames: 0
- Start of hotspot frames: 0

The video for this scenario is of two people walking in front of the camera. Neither of them are wearing fire-coloured clothing. This video was recorded during the day time.

NIR threshold variance testing

Table 13: NIR detection data for scenario 5

VIDEO 5	**INFRARED THRESHOLD (TR)**			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	375	375	375	375
NIR MOVEMENT	212	327	259	299
NIR POSSIBLE	0	0	17	107
NIR PROBABLE	31	30	62	14
COMBINED	0	0	0	0

FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	1	1	1	1
NIR MOVEMENT	3	2	2	2
NIR POSSIBLE	N/A	N/A	49	2
NIR PROBABLE	3	7	9	55
COMBINED	N/A	N/A	N/A	N/A

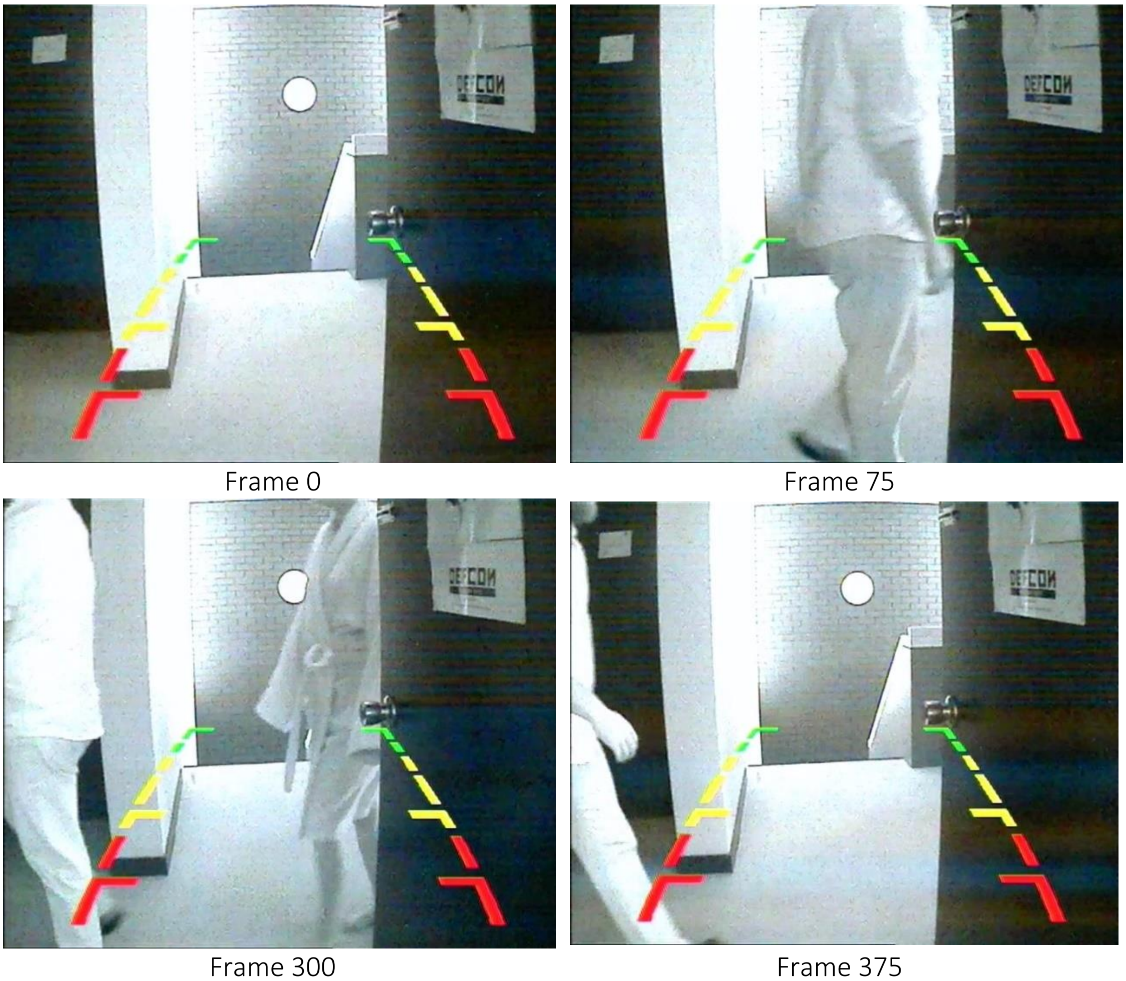

Figure 94: 4 frame sequence of the NIR video used for scenario 5.

We can see that there is a large amount of NIR light being detected by the camera, the light on the far wall is a concentrated source of this light. This is the source of the light along with the very bright reflection on the wall are the source of our NIR emissions. We can also see that the clothing worn by the two people also reflects a large amount of NIR light, adding to the NIR detections.

The NIR movement detections are mostly due to flickering noise from the camera at the edges of the bright reflections in the image. There is no discernible trend in the number of detections made in relation to the Tr threshold used. But in all cases the first movement detection was made within the first three frames of the video.

Frame 0 Frame 75

Figure 95: The processed NIR emission frames showing the reflections as detected by the algorithm.

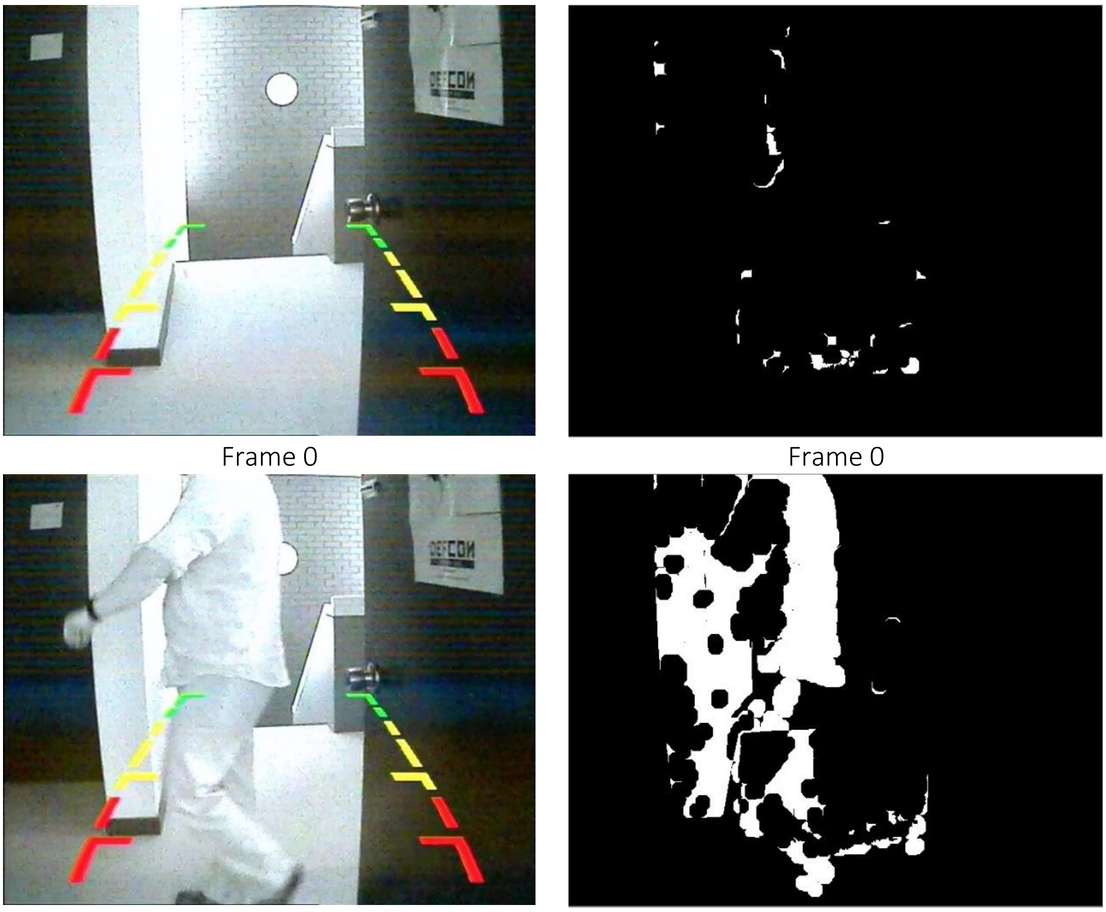

Frame 0 Frame 0

Frame 282 Frame 282

Figure 96: Movement detections, the left frames show the raw image while the right shows movement as detected by background subtraction.

Due to the size and movement of some of the reflections, possible and probable fire detections were made by the system since none of these conceded with any possible or probable detections on the colour detection track, no combined detections were made.

Colour threshold variance testing

Table 14: Colour detection data for scenario 5.

VIDEO 5	COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS	20	40	60	80
CC EMISSIONS	211	23	0	0
CC MOVEMENT	153	23	0	0
CC POSSIBLE	0	0	0	0
CC PROBABLE	0	0	0	0
COMBINED	0	0	0	0
FIRST DETECTION	20	40	60	80
CC EMISSIONS	2	123	N/A	N/A
CC MOVEMENT	2	123	N/A	N/A
CC POSSIBLE	N/A	N/A	N/A	N/A
CC PROBABLE	N/A	N/A	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

When a lower threshold was used a number of detections were made in the video. Upon further investigation these detections were made on the exposed skin of the two people walking through the frame.

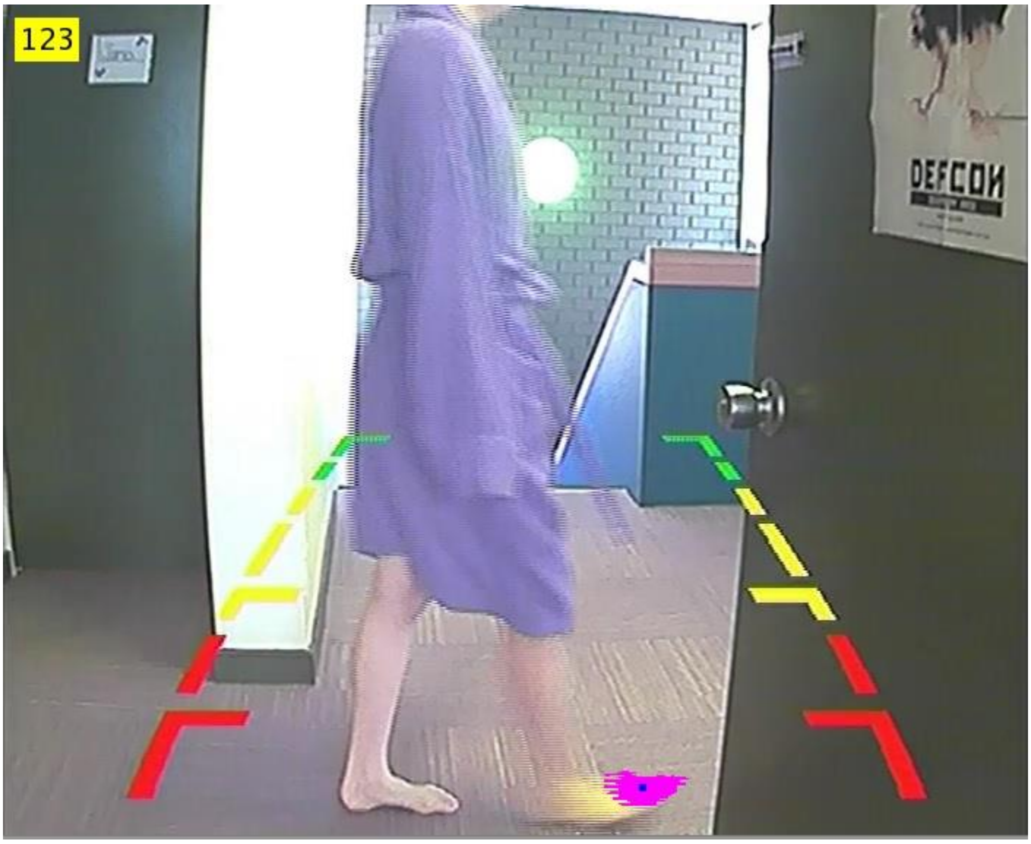

Figure 97: The colour detection made on the skin of the second person. Tc=40 Tr=80.

These tests had gone back to using the colour camera, and this may be the cause of the false positive detections on the skin. Despite these detections, none of these were big enough to cause a possible or probable detection.   On the higher threshold tests, no detections were made. It

5.2.6 Scenario 6

- Total length (Seconds): 20
- Total length (Frames): 500
- Total fire frames: 0
- Total hotspot frames: 0
- Start of fire frames: 0
- Start of hotspot frames: 0

The video for this scenario is of a rooftop on a bright sunny day. At the 10 second mark the camera is panned around.

Table 15: NIR detection data for scenario 6

VIDEO 6	INFRARED THRESHOLD (TR)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	500	500	500	500
NIR MOVEMENT	296	309	296	296
NIR POSSIBLE	21	7	2	1
NIR PROBABLE	0	0	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	1	1	1	1
NIR MOVEMENT	204	18	204	204
NIR POSSIBLE	229	322	371	227
NIR PROBABLE	N/A	N/A	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

The reason for this test may not be clear at first. In this video there is an extremely large amount of NIR light being reflected off the roof. So much so that the image from the NIR camera appears to be no different to that from the colour camera. Only when the image is inspected side by side can we see the lack of colour in the NIR image.

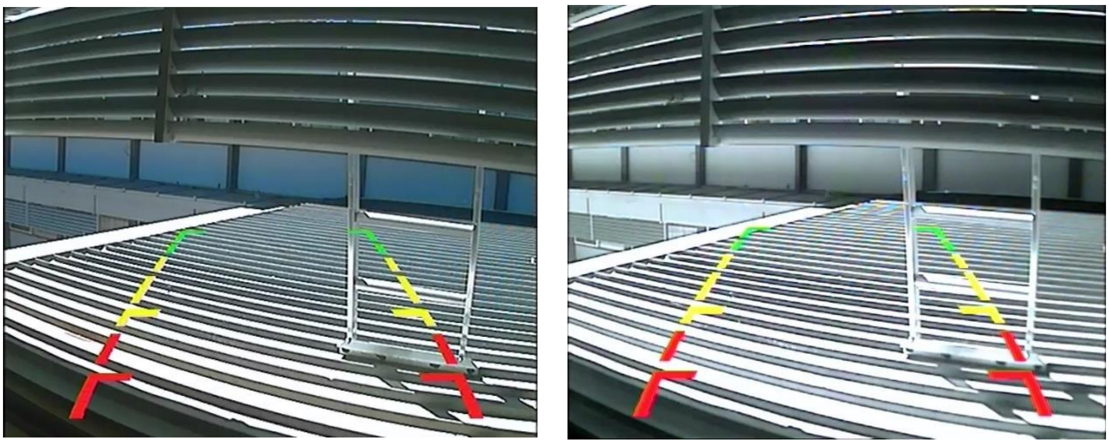

Colour Camera, frame 0 NIR camera, frame 0

Figure 98: The output from the two cameras, there is very little difference due to the large amount of NIR being detected.

This could mean that on a bright and sunny day, the NIR camera may have a large number of false positives.

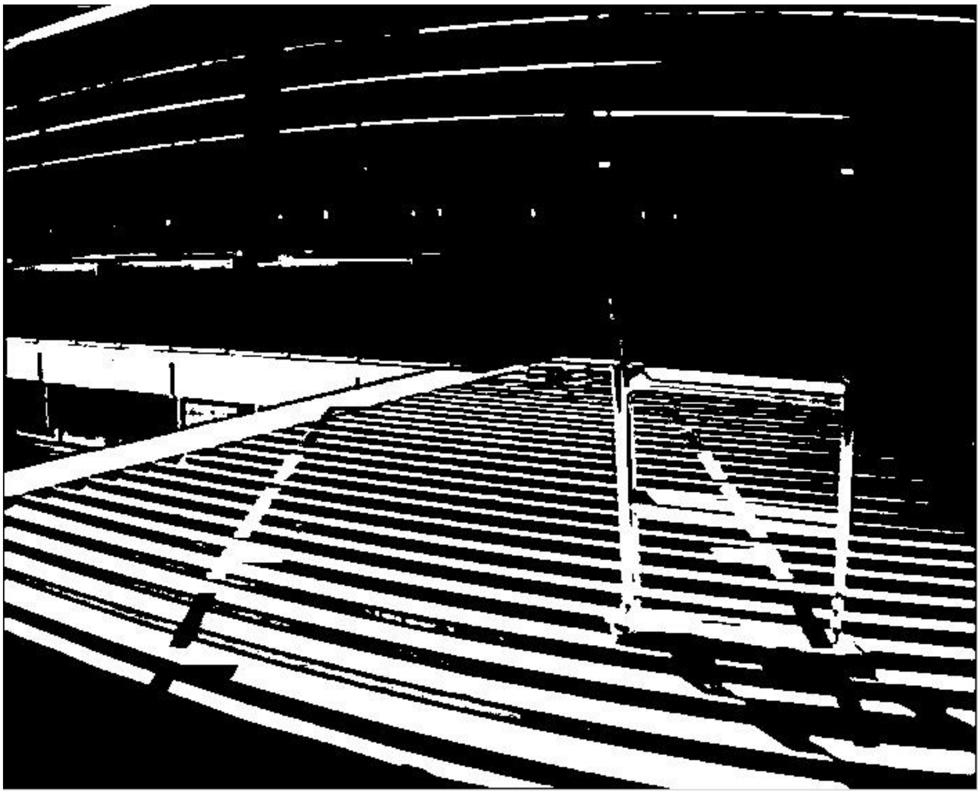

Figure 99: The NIR emissions detected by the algorithm at frame 0. Tr=80, Tc=60.

We can see in figure 99 that large amounts of the NIR are detected by the algorithm and none of these features are fire related. This is where our morphological filters come into their own.

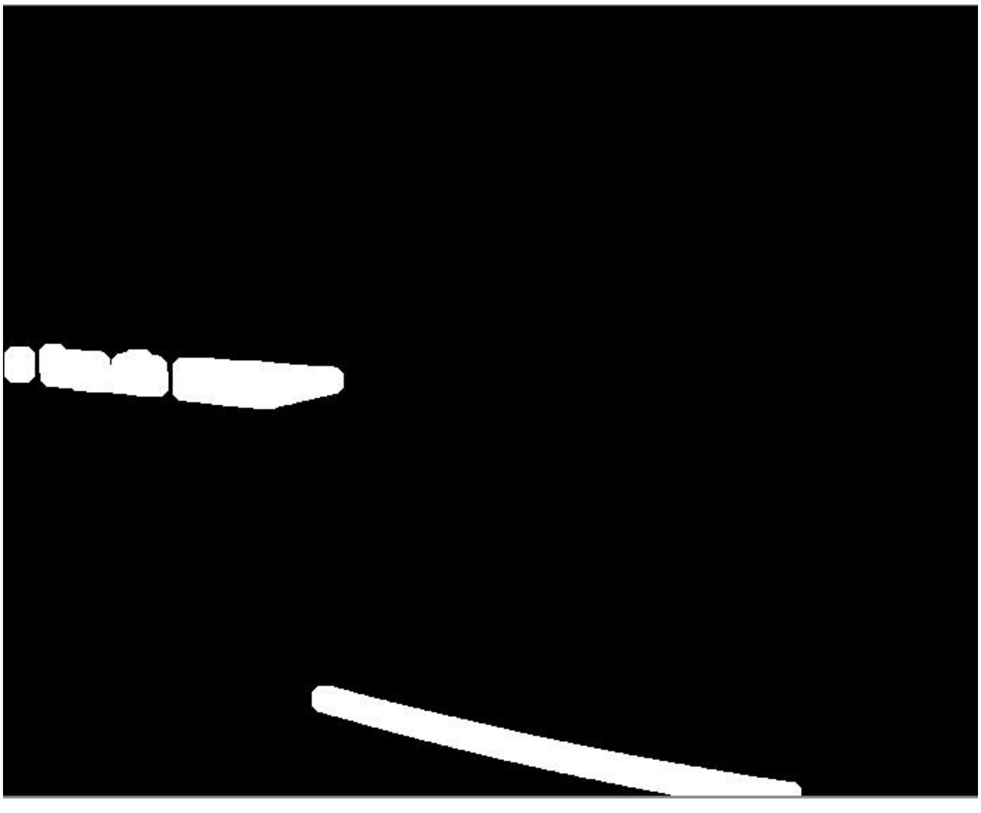

Figure 100: The same frame, after morphological filtering is applied.

The filters remove the vast majority of the false positive detections, leaving the two large objects seen in figure 100. But these objects still do not cause more than a couple of possible or probable detections. The amount of possible and probable fire detections decreases as the Tr threshold increases.

Since not detections at all were made in the colour detection track, no combined detections were made.

5.2.7 Scenario 7

- Total length (Seconds): 30
- Total length (Frames): 750
- Total fire frames: 0
- Total hotspot frames: 0
- Start of fire frames: 0
- Start of hotspot frames: 0

Two people are moving and dancing in front of the camera. One person is wearing a bright red jumper. This video was recorded during night time.

NIR threshold variance testing

Table 16: NIR detection data for scenario 7.

VIDEO 7	INFRARED THRESHOLD (TR)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	637	507	352	206
NIR MOVEMENT	544	411	251	173
NIR POSSIBLE	32	19	3	0
NIR PROBABLE	0	0	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	7	145	149	181
NIR MOVEMENT	11	159	227	311
NIR POSSIBLE	315	541	717	N/A
NIR PROBABLE	N/A	N/A	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

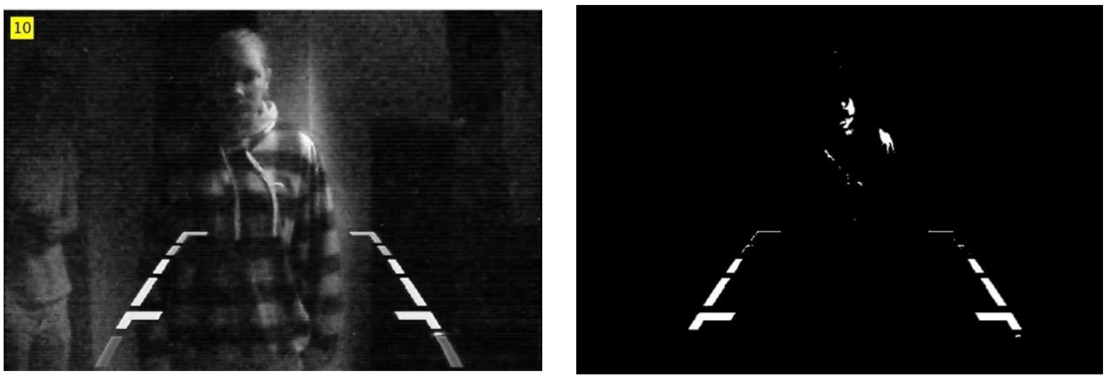

NIR camera feed, frame 10 NIR emissions, frame 10

Figure 101: Comparison of the NIR video feed and the NIR emissions detected. The hardcoded reversing lines are the largest object in the video.

This scenario has many similarities to scenario 5. But since this scenario video was filmed during night time there is a dramatic reduction in the overall amount of NIR emissions being detected by the algorithm. In particular, the Tr=80 and Tr=90 had a dramatic reduction in the amount of NIR emissions detected. This has a flow on effect for the rest of the detections made.

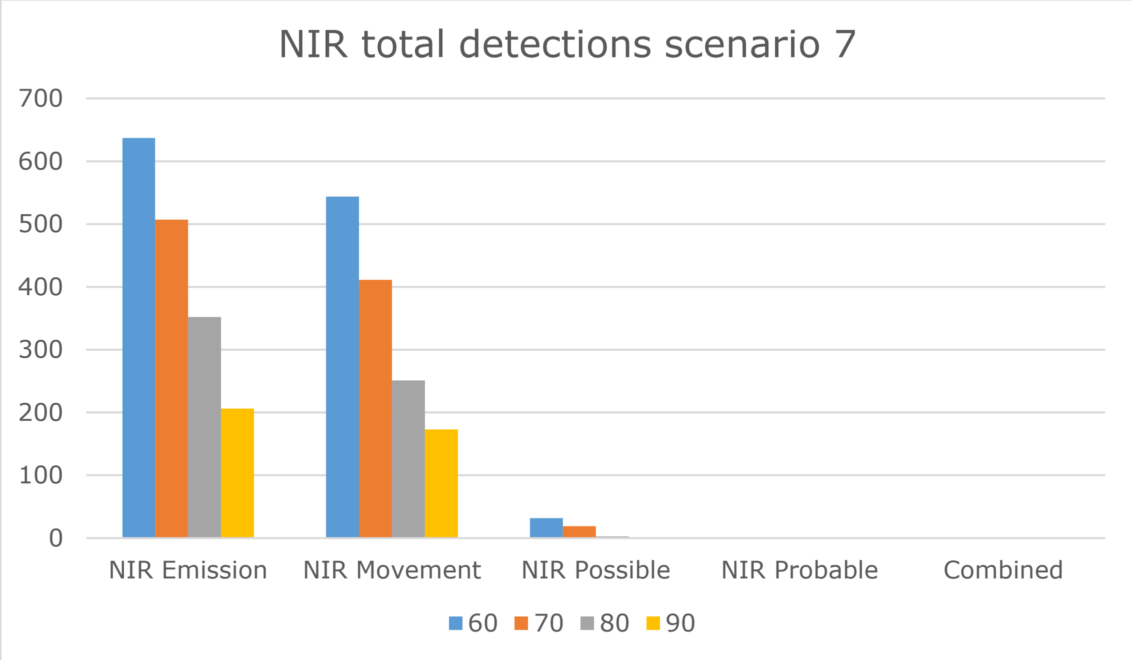

Figure 102: Chart of the total NIR detections made in scenario 7.

Colour threshold variance testing

Table 17: Colour detection data for scenario 7

VIDEO 7	COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	20	40	60	80
CC EMISSIONS	723	522	432	181
CC MOVEMENT	543	402	338	129
CC POSSIBLE	55	48	1	0
CC PROBABLE	5	6	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	20	40	60	80
CC EMISSIONS	1	1	1	4
CC MOVEMENT	2	4	5	4
CC POSSIBLE	4	15	117	N/A
CC PROBABLE	102	86	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

For this scenario, a fire coloured element was deliberately introduced in an attempt to cause false positive detections. We can see that the fire coloured element was detected in numerous different frames.

Figure 103: Frame 10 of the colour detection algorithm. We can see that it has detected a fire coloured object in the frame. Tc=60, Tr=80.

We can see in figure 103 that the colour detection algorithm has not detected the entire fire coloured object, only a small section of it. If we look closer at the detected area, we can see that it is has higher luminance than the rest of the object.

Figure 104: Frame 65 of the colour detection algorithm. It has not only detected multiple fire regions but also movement in the frame. Tc=60, Tr=80.

In figure 104 we can see as the fire coloured object moved to the right of the frame the original fire region becomes too luminescent to be detected as a fire region.

We see a trend of decreased detections as the Tc threshold is increased. For Tc =60 and Tc=80 we see no probable fire detections made at all. Since the NIR detection track made no probable fire detections either, no combined fire detections were made for this scenario.

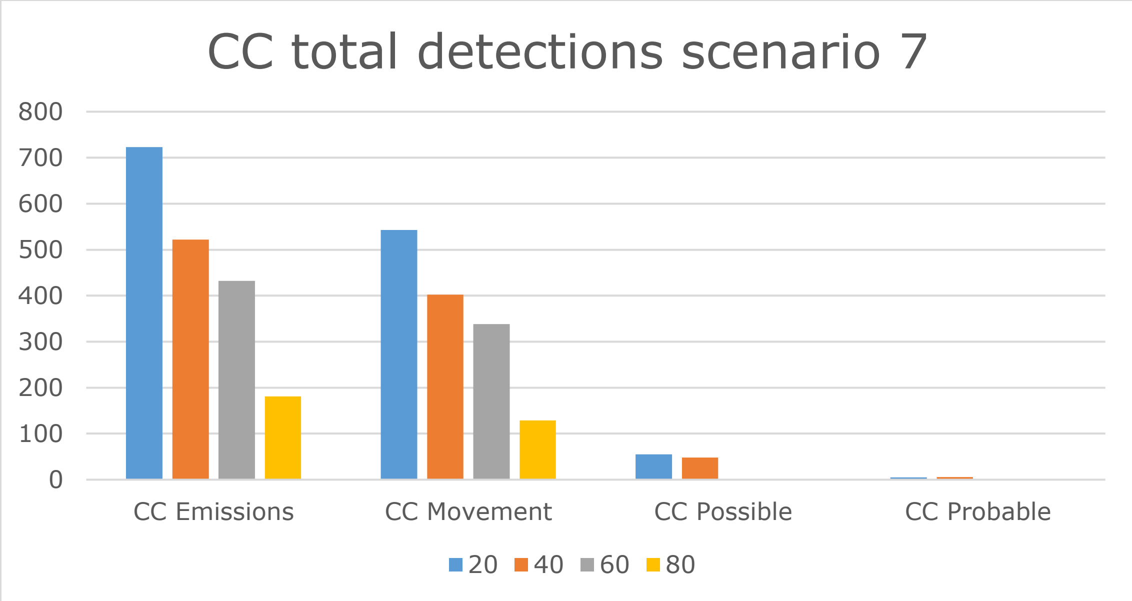

Figure 105: Chart of the total detections made by the colour detection track in scenario 7.

5.2.8 Scenario 8

- Total length (Seconds): 30
- Total length (Frames): 750
- Total fire frames: 529
- Total hotspot frames: 529
- Start of fire frames: 221
- Start of hotspot frames: 221

The video for this scenario has a person wearing a fire colour jumper who walks across the frame, starts a small fire and starts moving.

Table 18: NIR detection data for scenario 8.

VIDEO 8	**INFRARED THRESHOLD (TR)**			
TOTAL DETECTIONS (FRAMES)	60	70	80	90
NIR EMISSION	284	197	107	63
NIR MOVEMENT	118	63	31	14
NIR POSSIBLE	0	0	0	0
NIR PROBABLE	0	0	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90
NIR EMISSION	206	206	208	208
NIR MOVEMENT	210	213	213	213
NIR POSSIBLE	N/A	N/A	N/A	N/A
NIR PROBABLE	N/A	N/A	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

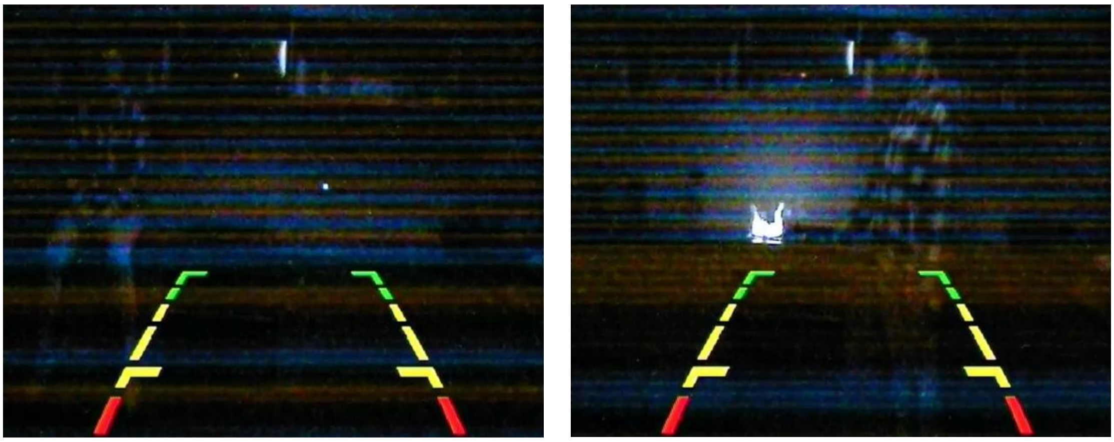

Frame 0 Frame 220

Figure 106: The first frame and the ignition frame of the NIR video.

This scenario is the most difficult that the system was presented with. It was deliberately designed to induce false positives and miss the actual fire in the frame. A person dress in a fire colour moves around the frame, this is designed to confuse the colour detection track. The fire itself used a different fuel to the other tests, natural gas, and has a blue colour unlike the red-orange flame of the other tests. We can also see that the fire is much smaller than it has been in previous tests.

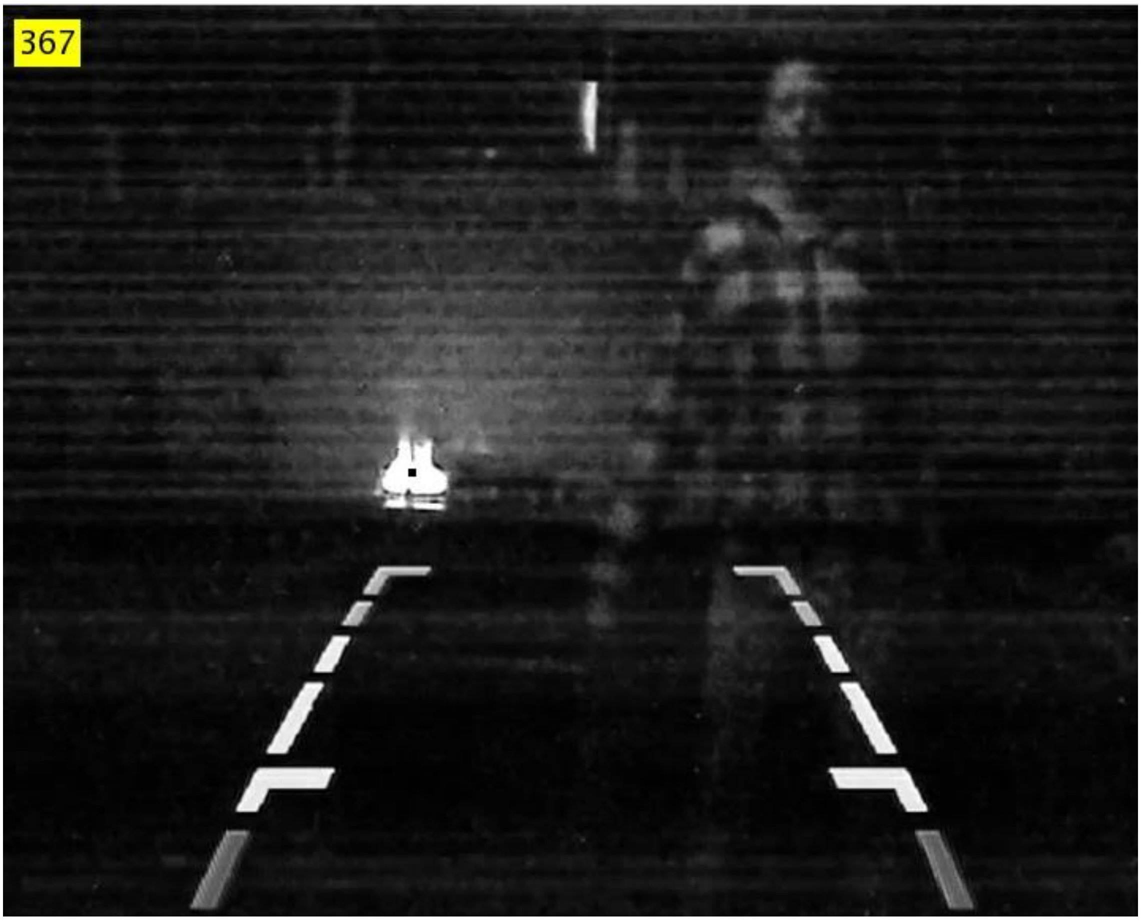

Figure 107: Frame 367, the black dot represents the detected centre of the fire. Tc=60, Tr=80

Despite these disadvantages, the NIR system was able to detect the fire in some of the frames of the video. But because of how small the fire was it did not cause a possible or probable fire detection in any of the tests. Once again we see a downward trend in detection as the Tr threshold is increased.

Colour threshold variance testing

Table 19: Colour detection data for scenario 8.

VIDEO 8	COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	20	40	60	80
CC EMISSIONS	750	750	750	121
CC MOVEMENT	246	216	201	198
CC POSSIBLE	101	77	35	2
CC PROBABLE	3	1	0	0
COMBINED	0	0	0	0
FIRST DETECTION FRAME	20	40	60	80
CC EMISSIONS	1	1	1	131
CC MOVEMENT	4	2	128	131
CC POSSIBLE	154	167	548	573
CC PROBABLE	573	574	N/A	N/A
COMBINED	N/A	N/A	N/A	N/A

Frame 0 Frame 220

Figure 108:The first frame and the ignition frame of the colour video.

The performance of the colour detection algorithm was average at best. None of the fire detections took place on the fire. The colour algorithm either marked the red jumper worn by the person or it marked the red shelves in the top of the frame. The red shelves were added accidentally into the frame but no the less they did provide additional objects to confuse the system. Apart from the highest Tc value tested, all of the tests had colour detections in every frame but even on the lowest Tc threshold there was only three probable fire detections.

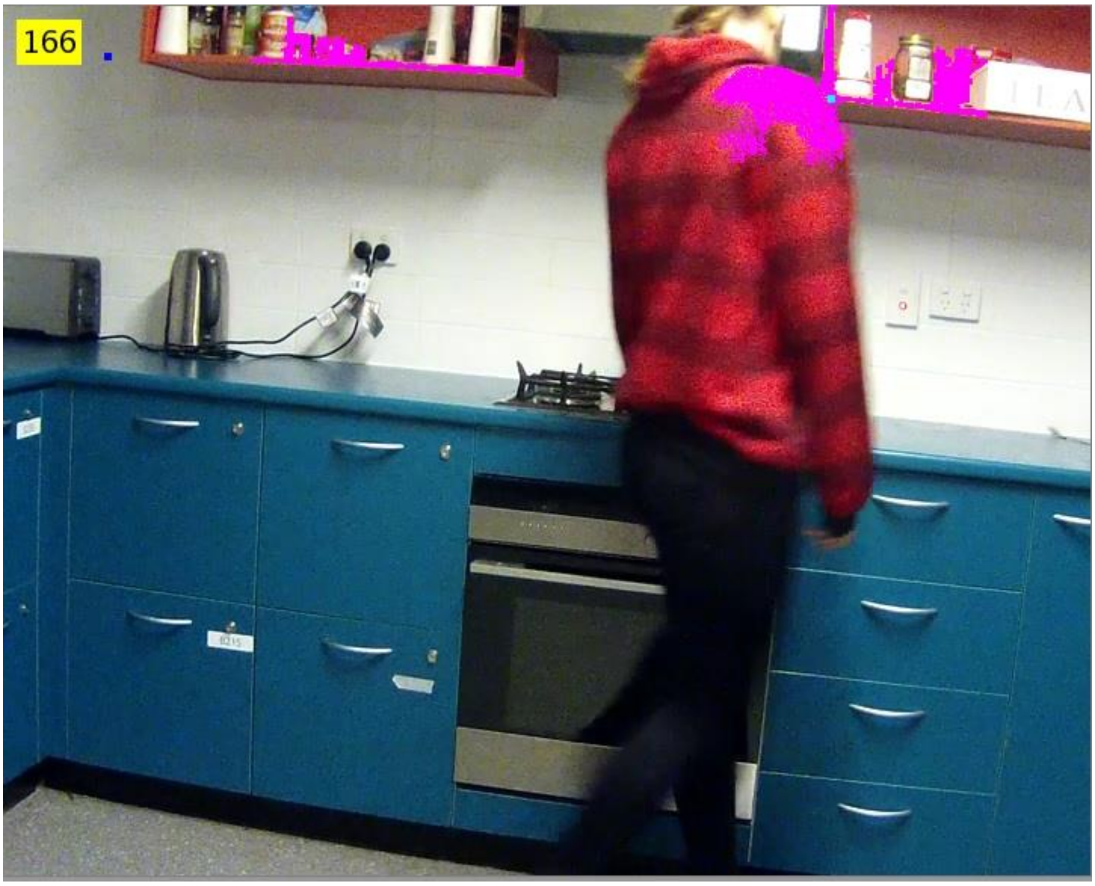

Figure 109: Frame 166 of the colour detection feed. The pink areas are detected fire regions.

Figure 110: Frame 280. Pink areas are detected fire regions; the blue flame is undetected.

5.3 Field Testing

Originally the aim of the field testing was to trial the system on a working harvester during harvest. For several reasons, mostly time related, this was not able to be done.

I was able to get a look at three different harvesters, including the engine bays and underneath other panels in the harvester. I was able to assess where the cameras for the fire detection system could potentially be mounted in the harvester. Only one of these harvesters will be used in this section in order to summaries the findings. The other two harvesters could have cameras mounted in similar positions.

5.3.1 Assessing potential camera mounting positions

Figure 111: The engine bay of the first harvester.

The first harvester was a John Deere 9760STS. It has an 8.1 litre 253 Kw diesel engine located at the rear of the harvester as seen in figure 111.

Figure 112:The induction system on top of the radiator

The engine induction system and exhaust system is routed across the top of the engine bay. We can also see the radiator behind the open panel in the same figure.

Figure 113: The exhaust muffler on the harvester. This is located on the opposite side to the induction system.

Figure 114: The radiator is on the right side of the image. On top we can see the air cleaner.

The engine is mounted transversely, with the front of the engine at the radiator (right hand side). The induction side of the motor is facing the rear of the harvester and the exhaust side facing forward. The engine output powers hydraulic motors and is also directly linked to belt and chain drives.

Figure 115: The induction side of the motor. The air cleaner can be seen on the right hand side of the frame.

Figure 116: The top of the motor can be seen with the cover removed. The air cleaner can just been seen in the bottom of the frame.

There are several mounting positions that could be used for the systems cameras. The first such place is the edge of the grain tank in figure 111. It overlooks the engine bay and provides an ideal place to survey the exhaust system. The second such position could be the engine bay wall at the rear of the harvester. It can just be seen at the extreme right of figure 114. A camera mounted in this position would have a view roughly similar to figure 115, but lower down. This would give an idea view of the induction side of the motor.

Figure 117: The output of the engine.

Ideally a camera would also be mounted to view the exhaust side of the engine, it was hard to see if there was any position there that would allow this.

A final position would be looking at the output of the engine. Mechanical failure is the most common source ignition and any system at did not monitor the complex array of belts and chains would not be a complete system.

6. Analysis

6.1 Hardware performance and observations.

The hardware used was mostly successful at detecting fire and fire like phenomena. The main issues with the hardware used during this test were the low signal to noise ratios of the camera under some circumstances and the unwanted NIR detections made by the colour camera.

The CCD cameras that were used were not of the best quality. Under low light conditions large amounts of noise were created by the camera. This noise could possibly false positive detections or prevent a true fire detection from being properly recognised by the software. This was noticed in particular on the NIR camera. Since the camera was fitted with a NIR pass filter, the camera sensor was essentially operating in low light conditions under all circumstances. The noise is very noticeable during the night time tests, scenarios 1,2,7 and 8. In scenario 8 the NIR video (figure 107) is grainy and it is hard to make out any detail. The colour video (Figure 109) from scenario 8, which used the higher quality DRIFT camera, is clear and detailed making it much easier to see what is happening in the frame.

The second main issue was the colour camera. Since the camera did not have any sort of NIR blocking filter in the camera lens, when it was exposed to fire the NIR emissions from the fire the camera was unable to detect any fire colours. Instead the fire appeared to be a bright white/purple bloom. To rectify this problem, a different camera was used with a much better quality CMOS sensor and most importantly, a IR cut filter built into the lens.

When this camera was used the colour of the fire was now detectable. Under lowlight conditions this new camera still had a tendency to show the fire as a white bloom (figure 86) but there was still enough residual colour information left for the colour detection algorithm to positively identify the fire. Under normal daytime conditions the flame colour was easily seen (figure 93)

We can see from the hardware tests that the NIR camera is able to detect objects with a temperature about 350-400 ^{o} C. (Figure 64) Any object above 400^{o}C is extremely visible to the NIR camera. The hardware seemed to perform adequately under several different conditions. There is a lack of data about how well the cameras will perform under dusty conditions with a large amount of chaff in the air. With the testing materials and facilities I had at my disposal there was no safe way of properly testing this. I was hoping to gain some data in the field testing however, I ran out of time to be able to do these tests.

Once the program was properly optimised the laptop used to run the program was able to run it at a reasonable speed with the frame rate varying from 8 to 15 frames per second. An informal endurance test was done on the live video system to see if it would crash from any instability in the program, after six hours there was no sign that the program would stop working.

I was disappointed that I was not able to do any testing on a harvester during this project. This mostly came down to timing, as harvest takes place over the summer months, by the time the initial research was done it was too late to gather data in the 2015 harvest and it is too early now for get data from the 2016 harvest.

6.2 Software performance and observations.

The overall performance of the software can be summed up by two parameters, the number of true fire detections made and the number of false alarms. We want the former to be as high as possible with the latter as low as possible.

But I believe that for a fire detection system the most important factor is detection speed. If the system has high overall accuracy but takes 10 or 15 seconds to detect the fire and raise the alarm it is not any better than having a simple low tech system. The system only needs to detect one or two frames for it to alert the harvester operator, from there he can quickly make the final decision as to if it is a true fire detection.

6.2.1 Scenarios 1-4

The first four scenarios of the software testing involved a fire or hotspot. Only two of these scenarios, three & four, made confirmed fire detections. Scenario one did not have a fire, only an extremely hot object. The colour algorithm was not able to detect this hot object at all since it was not the right colour. The NIR algorithm detected and tracked the hot object for the entirety of the video

In scenario 2, it was once again the colour algorithm that let down the program. This time though it did make some detections except these were not of sufficient quality to trigger any possible or probable fire detection on the colour track.

But in these two scenarios, the camera that was used was not performing well enough to be able to detect the fire colour. Once it was substituted in the other two scenarios for a better camera, the detection ability of the program changed for the better.

In scenario three and four, positive fire detections were made. Using the data from the combined fire detections we are able to calculate the mean first fire detection frame.

Table 20: Total fire detection statistics of scenarios 3 and 4

	Scenario three	Scenario four
Total fire frames	385	370
Mean number of detected fire frames	70	101.375
Standard deviation	26.960	22.557

The results are not too encouraging at this stage, Scenario Three only has an accuracy of 18.18% while scenario four is 27.398% accurate. But I believe that this more due to synchronisation issues with the videos rather than inaccuracy in the algorithm. If we look at the probable fire detections made by the NIR and Colour detection tracks a different story is told.

Table 21: Colour and NIR detections statistics for scenario 3.

	Scenario three	
	NIR	Colour
Total fire frames	385	385
Mean of probable fire frames	344.75	81.75
Standard deviation	27.801	46.032

Based on only the NIR information the accuracy of detection is almost 90% but the colour accuracy is much lower at 21.23%. This is probably due to scenario three being run during the night time, which favours NIR detection.

Table 22: Colour and NIR detections statistics for scenario 4.

	Scenario four	
	NIR	Colour
Total fire frames	370	370
Mean of probable fire frames	163.25	246.75
Standard deviation	46.0317	18.910

When the fire is detected during the day time it is the colour detection method that out performs the NIR. Colour has an overall accuracy of 66.69% while the NIR falls to 44.12%.

It's hard to say exactly how accurate the system would be from these results. But we can see that the NIR camera has the potential to outperform the colour detection camera under low light or night time conditions. During the day however the situation is reversed, the colour detection method is able to outperform the NIR method.

Now let's look at the detection times, once again we'll first look at the combined detections

Table 23:First time fire detection statistics of scenarios 3 and 4

	Scenario three	Scenario four
First fire frame	365	380
Mean first detection frame	411.375	467.390
Standard deviation	28.2536	62.7988

For a system running at 25 frames per second there would be an elapsed time of 1.855 seconds in scenario three between the start of the fire and the mean first fire detection frame. In scenario four this elapsed time blows out to 3.295 seconds.

For this program 95% percent of the threshold values Tc and Tr will have their first fire detection frames within two standard deviations of the mean. For scenario three we can expect these detections to occur between frame 365 (since our positive fire detection should not occur before the fire starts) and frame 468 or a maximum elapsed time of 4.115 seconds between the first start and first detection.

In scenario four 95% of the detections will occur between frame 380 (Fire start) and frame 587.9728 or a maximum elapsed time of 8.3189 seconds.

To get a better idea of what the detection times are, lets now look how the NIR and colour methods performed for each test.

Table 24: First time fire detection statistics of scenario 3

	Scenario Three	
	NIR	Colour
First fire frame	365	391
Mean first detection frame	371.5	417.25
Standard deviation	13	42.287

Now we can see that the gap is much closer. There is only a 6.5 frame difference for NIR and 26.25 difference for colour, much better than the 46 frame difference in the combined scenario. We can also see that the standard variation is much larger in the colour detection method, meaning that it is much more sensitive to threshold changes.

Table 25: First time fire detection statistics of scenario 4.

	Scenario Four	
	NIR	Colour
First fire frame	370	380
Mean first detection frame	341.25	391
Standard deviation	36.105	9.2736

The NIR detection method has a mean fire detection in advanced of the actual fire start because it is picking up heat from a blow torch as it is lit in the frame. The colour method did not pick this up as the flame was blue. We can see that the colour detection method only needed 11 frames to detect the fire and that its standard deviation is very small, meaning we can expect good amount of accuracy most of the time.

Even though there is only a limited amount of data here we can reasonably assume that in most 'normal' fire situations, the system can be expected to detect the fire within 3-5 seconds or worst case no more than 10 seconds before the alarm is sounded.

6.2.2 Scenarios 5-8

The last four scenarios were designed to see how the system would respond to fire like phenomena, the last two scenarios were designed specifically to induce false positive conditions.

Scenarios five and six were designed to test out the NIR system. We can see that if the camera only used the NIR detection method, the system would have a number of false positives in scenario five. Scenario six was save by the post algorithm filtering processes designed to remove noise and small objects from the frame.

Table 26: Total fire detection statistics of scenarios 5 and 6

	Scenario five	Scenario six
Total fire frames	0	0
Mean number of NIR probable fire frames	34.25	0
Standard deviation	20.073	0

Scenario five had a false positive rate of 9.1% which upon first inspection does not seem bad. But given the nature of that test, which only had movement, it should be lower than this. But because the system used colour detection method as well, there was no false alarms actually made.

Scenarios seven and eight were designed to really test the program. Scenario seven was designed to test out the ability of the colour algorithm to differentiate between fire coloured objects and actual fire. The result was for the most part it was. Only 2-3 detections were made on average that were large enough to be considered a probable fire. And the NIR system did not make any detections so no combined detections were made.

Scenario eight was designed to specifically induce a failure condition. There are multiple fire coloured objects, some are moving and there is a fire lit during the video. But since it uses gas as a fuel, the flame is blue.

Table 27: Total fire detection statistics of scenario 8

	Scenario eight	
	NIR	Colour
Total fire frames	529	529
Mean probable fire frames	0	1
Standard deviation	0	1.414

The system fails to detect the fire in this scenario, which is not surprising. There is a large amount of confusing data that the system is currently not capable of dealing with. The lack of any probable fire detections by the NIR detection program is disappointing but not unexpected given the size of the fire.

The two threshold values, Tr and Tc, had a major effect on the programs detection ability but these are not the only variables that have an effect on detection ability. The values used for the morphological filtering operations also have a large effect on the detection ability of the program.

The filtering value used in the colour camera detection algorithm in particular is a big culprit. I believe that since this value was set so high initially and was not altered during testing it had a big effect on the number of detections made by the colour algorithm.

7. Conclusion

The overall aim of this project was to construct a fire detection system that would be suitable for use on a combine harvester. I believe that while the system developed by this project is not there yet, it is well on its way towards becoming a viable fire detection system.

We know that harvester fires are a problem for the industry, we hundreds happening every year. We know that three-quarters of the fires start in the engine bay and that the most common fire fuel is the chaff created by the fire process. We know that consumer cameras can be modified to detect NIR light and that this NIR light can be used to detect hot objects.

Looking at the specific objects set out in the introductory chapter, we can see that the system is capable of detecting fire and hotspots. It can differentiate between a static hotspot created by the exhaust and a dynamic hotspot created by a fire. So far the system has demonstrated reasonable accuracy and detection speed on equipment that is very economical.

I believe that the quality of the equipment used resulted in less than optimal results during the testing phase. We can see that in the scenarios where higher quality cameras were available the detection accuracy of the system increased. Despite these results I believe that this fire detection system shows promise and with further development could yield a very accurate and very effective fire detection system.

Ultimately it is still a very difficult task to detect under any conditions with machine vision, especially in the difficult conditions faced by this system. I am extremely happy that I was able to produce a system that was able to detect fire at all, let alone one that can detect fire with any accuracy.

Future work

Looking beyond this project towards a commercially viable system, the following tasks should be looked at.

- Further testing and refinement using better equipment. Changing the cameras for ones with better quality sensors may give much better accuracy and detection speeds. The detection values used in the program currently are far from perfect and need to be refined.
- Testing the system on a harvester under normal operating conditions. One of the major steps that I regret missing in this project was the lack of testing I was able to do with harvesters.
- Adding addition methods of fire detection to increase the accuracy of the system. This could take two different forms. The first is integrating ‘dumb’ smoke or heat detectors into the system to provide a backup to the camera system. The second way is to add additional processing to the program either improving on a current method of detection or adding a new one that I have over looked.
- Moving the software onto an embedded system is the next major step in the design process of this system. An embedded system has the much smaller power and room requirements making it practical to mount on the harvester. The program needs to be optimised for the embedded system to get the fastest processing speed.
- Scaling the system to use multiple cameras. As it is the system only uses the two cameras. The final system will use multiple twin mounted cameras around the harvester. The program needs to be rewritten to take advantage of parallel processing and needs to be able to allocate processors to different camera ‘streams’

- Integrating the drone system into the fire detection system. I originally thought the drone would play a much bigger part in the overall system the it ultimately did. Still I believe that it would be a useful addition to the system but this needs to be tested.
- Harvester destructive testing. One thing I found during my research that all of information regarding harvester fires was done on a statistical basis. Nobody had tried to start fires on a harvester to see how these fires began, spread and acted in the harvester environment.

References

Adimec, 2011. *CCD vs CMOS image sensors in Defense Cameras.* [Online]
Available at: http://info.adimec.com/blogposts/bid/58840/CCD-vs-CMOS-Image-Sensors-in-Defense-Cameras
[Accessed 6 July 2016].

Adimec, 2015. *CCD vs. CMOS, sensitivity in low light improvements with industrial CMOS image sensors and cameras..* [Online]
Available at: http://info.adimec.com/blogposts/ccd-vs.-cmos-sensitivity-in-low-light-improvements-with-industrial-cmos-image-sensors-and-cameras
[Accessed 10 August 2016].

airstier UG & Co.KG, 2016. *yeair!.* [Online]
Available at: www.yeair.de
[Accessed 04 10 2016].

Australian Standard, 2016. *AS 5062-2016.* s.l.:SAI global.

Australian/New Zealand Standards, 1997. *AS/NZS 2153.1:1997.* s.l.:SAI Global.

Australian/New Zealand Standards, 1997. *AS/NZS 2153.7:1997.* s.l.:SAI Global.

B. Uğur Töreyina, Y. D. U. G. A. E. Ç., 2006. Computer vision based method for real-time fire and flame detection. *Pattern Recognition Letters,* 27(1), pp. 45-58.

Brandt J., M. H. R. F. F. M. a. O. R., 2013. *Test method for fire supression systems in bus and coach engine compartments,* s.l.: Society of Automotive Engineers.

Brogue, R., 2013. Sensors for fire detection. *Sensor Review,* 33(2), pp. 99-103.

Buchmann, I., 2016. *BU-401a: Fast and Ultra-fast Chargers.* [Online]
Available at: http://batteryuniversity.com/learn/article/ultra_fast_chargers
[Accessed 13 September 2016].

Buchmann, I., 2016. *BU-409: Charging Lithium-ion.* [Online]
Available at: http://batteryuniversity.com/learn/article/charging_lithium_ion_batteries
[Accessed 13 September 2016].

Byoung Chul Ko, K.-H. C. J.-Y. N., 2009. Fire detection based on vision sensor and support vector machines. *Fire Safety Journal,* Issue 44, pp. 322-329.

Byoung Chul Ko, K.-H. C. J.-Y. N., 2010. Early fire detection algorithm based on irregular patterns of flames and hierarchical Bayesian Networks. *Fire Safety Journal,* Issue 45, pp. 262-270.

Civil Aviation Safety Authority, 2016. *Commerial umanned flight - remotely piloted aircraft under 2kg.* [Online]
Available at: https://www.casa.gov.au/standard-page/commercial-unmanned-flight-remotely-piloted-aircraft-under-2kg
[Accessed 13 August 2016].

Civil Aviation Safety Authority, 2016. *Part 101 Amendments - Cutting red tape for remotely piloted aircraft.* [Online]
Available at: https://www.casa.gov.au/aircraft/standard-page/part-101-amendments-cutting-

red-tape-remotely-piloted-aircraft
[Accessed 24 May 2016].

Deere & Company, 2016. *Combine camera observation system.* [Online]
Available at:
http://salesmanual.deere.com/sales/salesmanual/en_NA/combines_headers/attachments/combines/monitoring_systems/combine_observation_system.html
[Accessed 2016].

Dempster, P., 2014. *CCD and CMOS Sensors.* [Online]
Available at: http://www.techbriefs.com/component/content/article/ntb/features/feature-articles/20063
[Accessed 6 July 2016].

Drysdale.D, 1996. *An Introduction to Fire Dynamics.* s.l.:Wiley-Interscience.

Glinz, A., 2013. *DJI Spreading Wings S800 hexacopter.* [Online]
Available at: https://commons.wikimedia.org/w/index.php?curid=29727889
[Accessed 2016].

Grains Research and Development Corperation, 2013. *Reducing Harvester Fire Risk: The Back pocket guide.* Canberra: Grains Research and Development Corperation.

Gtuav, 2014. *A Yamaha R-MAX.* [Online]
Available at: https://commons.wikimedia.org/w/index.php?curid=36610016
[Accessed 2016].

Hans Wasthuber, T., 2009. *Combine Harvester Diagram.* [Online]
Available at: https://commons.wikimedia.org/wiki/File:Maehdrescher_schema_nummeriert.svg
[Accessed 2016].

J.M. Shutske, M. V., 2002. *Combine Fire Prevention and Control Summit.* Chicago, ASAE.

J.M. Shutske, W. F. J., 1994. Grain Combine Fires: A Loss Reduction Approach. *Applied Engineering in Agriculture,* 10(2), pp. 175-182.

J.M. Shutske, W. F. L. G. S. P., 1990. Agricultural Machinery Fire Losses: A Preventative Approach. *Applied Engineering in Agriculture,* 6(5), pp. 575-581.

Joe, 2010. *InView being prepared for a night mission..* [Online]
Available at: https://upload.wikimedia.org/wikipedia/commons/9/9a/Inview_UAV_medium.jpg
[Accessed 2016].

Lawford, H., 2010. *New Holland combine at work near Stoneleigh, UK.* [Online]
Available at:
https://commons.wikimedia.org/wiki/Combine_harvester#/media/File:New_Holland_combine_at_work_near_Stoneleigh_1.jpg
[Accessed 2016].

Law, J., 2012. *Buring Issue: GRDC figures show that there are 1000 combine harvester fires in Australia each year..* [Online]
Available at: http://www.weeklytimesnow.com.au/machine/use-commonsense-to-prevent-combine-harvester-fire-during-harvest/news-story/f033c730547be732cedd997f4f6dbca7
[Accessed 2016].

Margques, O., 2011. Video fundamentals. In: *Practical Image and Video Processing using MATLAB.* New Jersey: John Wiley & Sons Inc, pp. 510-511.

Marques, O., 2011. Section 16 Color Image Processing. In: *Practical Image and Video Processing using MATLAB.* New Jersey: John Wiley & Sons, pp. 387-425.

Miu, P., 2016. *Combine Harvesters: Theory, Modeling, and Design.* 1st ed. Boca Raton, FL: CRC Press.

OpenCV, 2016. *How to Use Background Subtraction Methods.* [Online]
Available at: http://docs.opencv.org/trunk/d1/dc5/tutorial_background_subtraction.html
[Accessed 2016].

P. Kaewtrakulpong, R. B., 2001. *An Improved Adaptive Background Mixture Model for Realtime Tracking with Shadow Detection,* Middlesex: Kluwer Academic Publishers.

Piccardi, M., 2004. *Background subtraction techniques: a review.* [Online]
Available at: http://www.cs.nccu.edu.tw/~whliao/es2008/BackgroundSubtractionReview-Piccardi.pdf
[Accessed 24 September 2016].

Queensland Fire and Rescue Service, 2016. *Number of fires attended by QFES personnel where equipment involved in ignition,* Brisbane: Queensland Fire and Rescue Service.

Quick, D. G., 2010. *Harvester Fires,* s.l.: Grains Research and Development Corperation (GRDC).

Stauffer, C. &. G. W. E. L., 1999. *Adaptive background mixture models for real-time tracking.* s.l., s.n.

Strangefarmer.com, 2016. *Farm Dangers - John Deere Combine on fire in the field.* [Online]
Available at: http://www.strangefarmer.com/images/content/168109.JPG
[Accessed 01 10 2016].

Tamersoy, B., 2009. *Background Subtraction.* [Online]
Available at:
https://www.cs.utexas.edu/~grauman/courses/fall2009/slides/lecture9_background.pdf
[Accessed 04 08 2016].

Tamersoy, B., 2009. *Background Subtraction.* [Online]
Available at:
https://www.cs.utexas.edu/~grauman/courses/fall2009/slides/lecture9_background.pdf
[Accessed 25 09 2016].

Turgay Celik, H. D., 2009. Fire Detection in video sequence using generic color model. *Fire Safety Journal,* Issue 44, pp. 147-158.

University of California, 2013. *Background Subtraction.* [Online]
Available at: https://www.ics.uci.edu/~dramanan/teaching/cs117_spring13/lec/bg.pdf
[Accessed 24 September 2016].

Waltham, N., 2013. CCD and CMOS sensors. In: M. C. Huber, ed. *Observing Photons in Space.* New York: Springer, pp. 423-442.

Worksafe Queensland, 2015. *Be prepared for harvester risks.* [Online]
Available at: https://www.worksafe.qld.gov.au/news/newsletter/esafe-newsletters/esafe-

rural/september-2013/breaking-news/be-prepared-for-harvester-risks
[Accessed 20 03 2016].

Xuejun Chen, F. D., 2015. A dual-band flame detector based on video. *Optik,* 127(1), pp. 478-483.

Y. Le Maoult, T. S. J. J. O. a. J. P. A., 2007. Fire Detection: A New Approach based on a low cost CCD Camera in the Near Infrared. *Process Safety and Environmental Protection,* 85(3B), pp. 193-206.

Appendices

Appendix A: Project Specification

Appendix B: Queensland Fire and Rescue service data.

Appendix C: MATLAB code for live video

Appendix D: MATLAB code for recorded video

Appendix D: Software testing results

Appendix A: Project Specification

ENG4111/4112 Research Project

Project Specification

For: Steven Mosetter

Title: Using UAV's and machine vision in the early detection of combine harvester fires

Major: Mechanical Engineering

Supervisor/s: Dr. Tobias Low

Enrolment: ENG4111 - ONC S1, 2016
ENG4112 - ONC S2, 2016

Project Aim: To develop an integrated system that can detect fires on combine harvesters and warn the operator, giving them time to extinguish it before the total loss of the harvester and any other collateral damage can result.

Programme: Issue A, 23th March 2016

1. Research background information relating to harvester design, harvest conditions and what causes harvester fires. Identify sources of ignition, along with sources of fuel.
2. Design and construct testing prototype using the two CCD Cameras, modifying one so that it can detect Near Infrared Radiation (NIR).
3. Using testing prototype, observe fire causing conditions under various states including
    a. Fire under various states of combustion ie, pre-ignition, smoldering, starting to burn, fully alight, etc
    b. Fire from different types of fuel, including diesel, oil, chaff, hay, dust etc
4. Using data from initial observations attempt to create an automated fire detection program.
5. Review and optimize program by exposing it to increasingly more complex situations.
6. Test prototype on actual combine harvester.
    a. Look at the feasibility of different mounting options on and around the Harvester.
    b. Investigate the feasibility of using a remote platform to monitor the Harvester.
    c. Investigate the ability of the system to observe the harvester under normal operating conditions and under adverse conditions if possible.
    d. Determine the overall usefulness of the system, can it warn the operator in time to change the outcome of a fire event.

If time and resources permit

7. If testing with the harvester has determined that a remote platform could be useful, investigate the types of remote platforms that are currently available and which one best fits the requirements.
8. Look at the basic design of the platform.
9. Investigate the ways feedback can be provided to the operator.

Appendix B: Queensland Fire and Emergency service data.

Operational Performance Analysis Office
Queensland Fire and Emergency Services

This document has been produced and verified by the Operational Performance Analysis Unit on 19/04/2016

Number of fires attended by QFES personnel where equipment involved in ignition was identified as agricultural equipment

Month	Mobile property fire	Structure fire	Special structure fire	Outside storage fire, not rubbish	Stock fodder fire	Cultivated grain or crop fire	Scrub/bush/grass fire	Grand Total
Jul-15	1						1	2
Aug-15	3		2				3	8
Sep-15	2	1					2	5
Oct-15	2		2		1	1	5	11
Nov-15	2	1						3
Dec-15	1	1					1	3
Jan-16	3						2	5
Feb-16	3						2	5
Mar-16	1			1			1	3
Totals	**18**	**3**	**4**	**1**	**1**	**1**	**17**	**45**

Notes:
Data extracted from Fire database on 19/04/2016. Figures are based on live data and so are subject to change if query is repeated at a later date as records may be updated.
The records pertaining to incidents attended by fire fighters from the Queensland Fire and Emergency Services (QFES) are kept in accordance with Australasian Fire and Emergency Service Authorities Council (AFAC) Australian Incident Reporting Standard (AIRS).
In order to isolate incidents relevant to the query (agricultural machinary fires), AIRS field E8 - Equipment Involved In Ignition has been used. However, the code for this field that specifies "Agricultural Equipment and/or vehicle, including harvesters, headers, bailers and seeders" (966) only commenced use in QFES systems in July 2015, so this is the earliest data that can be provided.
To give an indication of the impact of the fires, AIRS field A23 - Type of Incident has been used to classify the incidents. The Type of Incident field reflects the most serious situation that occurred as part of the incident.

For further information and additional requests please contact the Statistics Team
Email: QFESStats@qfes.qld.gov.au

This document has been produced and verified by the Operational Performance Analysis Unit on 19/04/2016

Notes:

Data set includes incidents where:

AIRS field A23 - Type of Incident is between '100' and '199' (Fires and Explosions)
AIRS field A6 - Alarm Date is between '2015-07-01' and '2016-03-31'
AIRS field E8 - Equipment Involved In Ignition is '966'
AIRS field A5 - Exposure Number is '0' (indicating that fire incident is not as a result of another fire)
The incident record is marked as completed

Types of Incidents included in this data set are defined as follows:

Mobile property fire	*A23 between '150' and '159'*
Structure fire	*A23 between '110' and '129'*
Special structure fire	*A23 between '130' and '139'*
Outside storage fire, not rubbish	*A23 = '141'*
Stock fodder fire	*A23 = 143'*
Cultivated grain or crop fire	*A23 = '164'*
Scrub/bush/grass fire	*A23 = '162' or A23 = '163'*

Appendix C: MATLAB Program for live video

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Steven Mosetter
% Engineering Project 2016
% FIRE DETECTION PROGRAM
% IR CAM LIVE VIDEO
% VERSION 5
% COLOUR CAM LIVE VIDEO
% VERSION 3
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
clc
clear
%% Video%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
IR_feed = imaq.VideoDevice('winvideo',2,'YUY2_720x576');
IR_feed.ReturnedColorSpace = 'grayscale';
%COLOUR_feed = imaq.VideoDevice('winvideo',1);
CC_feed = imaq.VideoDevice('winvideo',3,'UYVY_720x576');
CC_feed.ReturnedColorSpace = 'rgb';
CC_feed.ReturnedDataType = 'uint8';

%% Infrared (IR) Camera processing values %%%%%%%%%
    IR_Threshold_percentage = 75;
%Sets the filtering threshold (Tr) for the IR camera. The higher
this number is the more luminescent the object must be. Set between
0 and 100
    IR_bit_size= 255;
%Set this value to 1 for double array(GREYSCALE OR INTENSITY), 255
for uint8 and 655353 for uint8
    IR_Threshold= (IR_Threshold_percentage/100)*IR_bit_size;
    Disk_size = 10;
%Size of the disk used to filter out non fire bloom shaped objects.
    disk = strel('disk',Disk_size); %Disk filter object
    Noise_Cancellation_Threshold_IR = 100;
%Any object with less than this number of pixels will be removed
from the picture. Used to filter out noise in the image.

%% Colour Camera (CC) processing values
    CC_Threshold = 50;
%Sets the filtering threshold (Tc) for colour filtering. Sets the
minimum difference between the Cr and Cb channels needed for a pixel
to be read as a fire region. Max difference must be 1.0 for double,
255 for uint8 and 655353 for uint16
    Noise_Cancellation_Threshold_CC = 800;
%Any object with less than this number of pixels will be removed
from the picture. Used to filter out noise in the image.

%% IR Camera movement processing objects
    detector_MIR = vision.ForegroundDetector('NumTrainingFrames',
20, 'LearningRate',
0.0001,'MinimumBackgroundRatio',0.7,'NumGaussians',5); %
Motion detection system object, using foreground detection.
    blob_MIR = vision.BlobAnalysis('CentroidOutputPort', true,
'AreaOutputPort', true,'BoundingBoxOutputPort',
true,'MaximumCount',5,'MinimumBlobAreaSource', 'Property',
'MinimumBlobArea',500); %Blob detection. Used for centroid and area.
```

```matlab
%% Colour Camera movement processing objects
    detector_CC = vision.ForegroundDetector('NumTrainingFrames', 20,
'LearningRate', 0.005,'MinimumBackgroundRatio',.9,'NumGaussians',3);
%Motion detection system object, using foreground detection.
    blob_CC = vision.BlobAnalysis('CentroidOutputPort', true,
'AreaOutputPort', true,'BoundingBoxOutputPort',
true,'MaximumCount',5,'MinimumBlobAreaSource', 'Property',
'MinimumBlobArea',500); %Blob detection. Used for centroid and area.

%% Video Processing
videoPlayer_IR = vision.DeployableVideoPlayer;
%Video player object for useroutput, normally only one of these
would be used at once. Usually the colour feed.
videoPlayer_CC = vision.DeployableVideoPlayer;
width = 2; %Sets size of the detection dot in the image.
nframes = 1; %Starts counter at first frame.
Fire_size_CC=0;   % Size of the fire as determined by the NIR
Emissions
Fire_size_IR=0;   % Size of the fire as determined by the Colour
processing
disp('The detection program will now begin. To terminate the
program, press control-c.')

while nframes<nframes+1

%% Default values
        %These values are reset every frame
        dist_NIR = 10000;
% Distance between the center of the NIR Emission centroid and NIR
movement centroid
        dist_CC= 10000;
% Distance between the center of the Colour Emission centroid and
Colour movement centroid
        dist_C=10000;
% Distance between the mean centers of Dist_NIR and Dist_CC
        Useroutput_1=step(IR_feed);
%Video feed from NIR Video/Camera
        Useroutput_2=step(CC_feed);
%Video feed from Colour Video/Camera

%% NIR_emission_processing
% This algorithm block looks for the fire based on NIR emissions.
        taggedfire_IR = read(IR_feed,nframes);
%Take frame from the NIR video feed
        Z=taggedfire_IR>IR_Threshold;
% Filter the image using the fire_luminance, leaves only very bright
objects. Refer to equation 13.
        No_reflect = imopen(Z,disk);
% Remove non disk shaped objects from video frame
        No_noise = bwareaopen(No_reflect,
Noise_Cancellation_Threshold_IR);
%Remove noise and artefacts from the picture
stats_IR = regionprops(No_noise, {'Centroid', 'Area'});
% Find the centroid and area of the remaining shapes in the picture
if ~isempty([stats_IR.Area]);
%If there are remaining objects in the picture
areaArray = [stats_IR.Area];
%Store the area of the objects in the picture
```

```
[size,idx] = max(areaArray);
%Find the object with the largest size, usually the flame due to the bloom.
b_IR = stats_IR(idx).Centroid;
%Centroid of the the largest object in frame.
c_IR = floor(fliplr(b_IR));
%Flip left to right and round towards negative infinity
row_IR = c_IR(1)-width:c_IR(1)+width;
%Places dot in the center row of the centroid area
col_IR = c_IR(2)-width:c_IR(2)+width;
%Places dot in the center column of the centroid area
Useroutput_1(row_IR,col_IR,1) = 0;
% Black dot on detected center of fire
Useroutput_1(row_IR,col_IR,2) = 0;
Useroutput_1(row_IR,col_IR,3) = 0;
Fire_size_IR= size;
%Sets fire size based on area of fire object in picture.
anomaly_IR = 1;
 elseif isempty([stats_IR.Area])
% If there are no remaining objects in the picture
Fire_size_IR = 0;% Recorders fire size as 0
anomaly_IR = 0;
end

%% MOVEMENT Processing NIR
% This algorithm block looks for movement in the NIR camera.
taggedFire_MIR= double(No_noise);
%Grabs image for the movement detection from the completed NIR emission block.
fgMask_MIR = step(detector_MIR, taggedFire_MIR );
%Background subtraction using detector
fgMnoise_MIR = bwareaopen(fgMask_MIR, 20);
%Filtering of small objects and noise from the image frame.
 [area,centroids, bboxes] = step(blob_MIR, fgMnoise_MIR); %Detecting the area and centroid of the remaing blob in the image.
if ~isempty(area)
c_MIR=floor(fliplr(centroids));
row_MIR = c_MIR(1)-width:c_MIR(1)+width;
%Places dot in the center row of the centroid area
col_MIR = c_MIR(2)-width:c_MIR(2)+width;
%Places dot in the center column of the centroid area
Useroutput_1(row_MIR,col_MIR,1) = 1; % White dot
Useroutput_1(row_MIR,col_MIR,2) = 1;
Useroutput_1(row_MIR,col_MIR,3) = 1;
anomaly_MIR = 1;
elseif isempty(area)
anomaly_MIR = 0;
end

%% NIR_Camera_Combination_detection
% This algorithm block determines the severity of the fire
% detection made by the two NIR blocks.
if all([anomaly_IR==1,anomaly_MIR==1])
dist_NIR=sqrt((c_IR(1)-c_MIR(1)).^2+(c_IR(2)-c_MIR(2)).^2);
%Calculates the distance between the NIR emission blob and the NIR movement blob
NIR_x=(c_IR(1)+c_MIR(1))/2; %Creates a center point based on the position of the two points.
NIR_y=(c_IR(2)+c_MIR(2))/2;
             end
```

```matlab
if all([dist_NIR<200,10000<Fire_size_IR,Fire_size_IR<40000]) %Fire
size criteria, possible fire. Fire_size describes the size of the
blob in the NIR emission frame.
NIR_detection_Possible=1;
NIR_detection_Probable=0;
elseif all([dist_NIR<200,Fire_size_IR>40000])
%Fire size criteria, probable fire. Fire_size describes the size of
the blob in the NIR emission frame.
NIR_detection_Probable=1;
NIR_detection_Possible=0;
else
NIR_detection_Possible=0;
NIR_detection_Probable=0;
end


%% COLOUR Processing_CC
% This algorithm block processes the colour video based on the
% flame colour
taggedFire_CC=step(CC_feed); %Reads the video frame from the video
object.
taggedFire_CC=rgb2ycbcr(taggedFire_CC); %This step is not normally
needed. The video from the cameras in the live program is already
YCbCr
Y=taggedFire_CC(:,:,1); %Splitting the image into its three
Cb=taggedFire_CC(:,:,2);
Cr=taggedFire_CC(:,:,3);
Y_ave=mean(mean(Y)); % Take the mean of the three channels of the
image
Cb_ave=mean(mean(Cb));
Cr_ave=mean(mean(Cr));
a=Y>Cb; % Equation 7
b=Cr>Cb;
c=Y>Y_ave; %Equation 9
d=Cb<Cb_ave;
e=Cr>Cr_ave;
f=(c+d+e)>=3;
%Adding the results of equation 9 into one logical matrix
g=(Cr-Cb)>=CC_Threshold; %Equation 10
h=(a+b+f+g)==4;
%Adds Equations 7,9 and 10 together into one logical matrix
h_2 = bwareaopen(h, Noise_Cancellation_Threshold_CC);
%Filtering of small objects and noise from the image frame.
stats_Cb = regionprops(h_2, {'Centroid', 'Area'});
if ~isempty([stats_Cb.Area]);
%If there are remaining objects in the picture
areaArray = [stats_Cb.Area];
%Store the area of the objects in the picture
[size,idx] = max(areaArray); %Find the object with the largest size
c_Cb = stats_Cb(idx).Centroid;
%Centroid of the largest object in frame.
c_Cb = floor(fliplr(c_Cb));
%Flip left to right and round towards negative infinity
row_Cb = c_Cb(1)-width:c_Cb(1)+width;
%Places dot in the center row of the centroid area
col_Cb = c_Cb(2)-width:c_Cb(2)+width;
%Places dot in the center column of the centroid area
 Useroutput_2(row_Cb,col_Cb,1) = 0;
%Places dot in center of detected fire
```

```
Useroutput_2(row_Cb,col_Cb,2) = 255;
Useroutput_2(row_Cb,col_Cb,3) = 255;
Fire_size_CC= size;
anomaly_Cb = 1;
elseif isempty([stats_Cb.Area])
% If there are no remaining objects in the picture
Fire_size_CC=0;
anomaly_Cb = 0;
end

%% MOVEMENT Processing_CC
% This algorithm block looks for movement in the CC camera.
taggedFire_MCC=double(h_2);
%Grabs image for the movement detection from the completed CC
detection block
fgMask_CC = step(detector_CC, taggedFire_MCC );
 %Background subtraction using detector
fgMnoise_CC = bwareaopen(fgMask_CC, 20);
%Filtering of small objects and noise from the image frame.
[area,centroids, bboxes] = step(blob_CC, fgMnoise_CC);
%Detecting the area and centroid of the remaing blob in the image.
if ~isempty(area)
c_MCC=floor(fliplr(centroids));
row_MCC = c_MCC(1)-width:c_MCC(1)+width;
%Places dot in the center row of the centroid area
col_MCC = c_MCC(2)-width:c_MCC(2)+width;
%Places dot in the center column of the centroid area
Useroutput_2(row_MCC,col_MCC,1) = 0;
Useroutput_2(row_MCC,col_MCC,2) = 0;
Useroutput_2(row_MCC,col_MCC,3) = 255; %Places dot in center of
fire.
anomaly_MCC = 1;
elseif isempty(area)
anomaly_MCC = 0;
end

%% Colour Camera Combined Detection
%This algorithm block determines the severity of the fire
%detection made by the two CC blocks.
if all([anomaly_Cb==1,anomaly_MCC==1])
dist_CC=sqrt((c_Cb(1)-c_MCC(1)).^2+(c_Cb(2)-c_MCC(2)).^2); %If both
blocks make detections, this finds the distance between the two.
CC_x=(c_Cb(1)+c_MCC(1))/2; %Creates a center point between the two
colour detections
CC_y=(c_Cb(2)+c_MCC(2))/2;
             end
if all ([dist_CC<300,10000<Fire_size_CC,Fire_size_CC<20000]) %Fire
size criteria, possible fire. Fire_size describes the size of the
blob in the CC detection frame.
CC_detection_Possible=1;
CC_detection_Probable=0;
elseif all ([dist_CC<300,Fire_size_CC>20000]) %Fire size criteria,
probable fire. Fire_size describes the size of the blob in the CC
detection frame.
CC_detection_Possible=0;
CC_detection_Probable=1;
else
CC_detection_Possible=0;
CC_detection_Probable=0;
end
```

```
%% COMBINED DETECTION
% This final block determines if the two probable
%detections from the fire detection tracks constitute a
% true fired detection
                if
all([NIR_detection_Probable==1,CC_detection_Probable==1])
dist_C=sqrt((NIR_x-CC_x).^2+(NIR_y-CC_y).^2);
end
if dist_C<200
Useroutput_1 = insertText(Useroutput_1, [10 10], 'ALERT! PROBABLE
FIRE! TAKE ACTION NOW!', 'BoxOpacity', 1,'FontSize', 20);
Useroutput_2 = insertText(Useroutput_2, [10 10], 'ALERT! PROBABLE
FIRE! TAKE ACTION NOW!', 'BoxOpacity', 1,'FontSize', 20);
detection=1;
beep;
elseif dist_C>100
detection=0;
end
step(videoPlayer_IR,Useroutput_1) %Displays the processed video feed
step(videoPlayer_CC,Useroutput_2)
    nframes = nframes+1;
end
```

Appendix D: MATLAB Program for recorded video

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Steven Mosetter
% Engineering Project 2016
% FIRE DETECTION PROGRAM
% IR CAM RECORDED VIDEO
% VERSION 5
% COLOUR CAM RECORDED VIDEO
% VERSION 3
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
clc
clear
%% Infrared (IR) Camera processing values
IR_Threshold_percentage = 75;
%Sets the filtering threshold for the IR camera. The higher this
number is the more luminescent the object must be. Set between 0 and
100
IR_bit_size= 255;
%Set this value to 1 for double array(GREYSCALE OR INTENSITY), 255
for uint8 and 655353 for uint8
IR_Threshold= (IR_Threshold_percentage/100)*IR_bit_size;
Disk_size = 10; %Size of the disk used to filter out non fire bloom
shaped objects.
disk = strel('disk',Disk_size); %Disk filter object
Noise_Cancellation_Threshold_IR = 100;
%Any object with less than this number of pixels will be removed
from the picture. Used to filter out noise in the image.

%% Colour Camera (CC) processing values
CC_Threshold = 50;
%Sets the filtering threshold for colour filtering. Sets the minimum
difference between the Cr and Cb channels needed for a pixel to be
read as a fire region. Max difference must be 1.0 for double, 255
for uint8 and 655353 for uint16
Noise_Cancellation_Threshold_CC = 800; %Any object with less than
this number of pixels will be removed from the picture. Used to
filter out noise in the image.

%% IR Camera movement processing objects
    detector_MIR = vision.ForegroundDetector('NumTrainingFrames',
20, 'LearningRate',
0.0001,'MinimumBackgroundRatio',0.7,'NumGaussians',5); %Motion
detection system object, using foreground detection.
    blob_MIR = vision.BlobAnalysis('CentroidOutputPort', true,
'AreaOutputPort', true,'BoundingBoxOutputPort',
true,'MaximumCount',5,'MinimumBlobAreaSource', 'Property',
'MinimumBlobArea',500); %Blob detection. Used for centroid and area.
%% Colour Camera movement processing objects
    detector_CC = vision.ForegroundDetector('NumTrainingFrames', 20,
'LearningRate', 0.005,'MinimumBackgroundRatio',.9,'NumGaussians',3);
%Motion detection system object, using foreground detection.
    blob_CC = vision.BlobAnalysis('CentroidOutputPort', true,
'AreaOutputPort', true,'BoundingBoxOutputPort',
true,'MaximumCount',5,'MinimumBlobAreaSource', 'Property',
'MinimumBlobArea',500);
%Blob detection. Used for centroid and area.
colour video
```

```
%% Video replay
% Use this section for testing on recorded video
IR_feed=VideoReader('Test_4_NIR.mp4');
total_frames_NIR = IR_feed.NumberOfFrames;
%Get frame count of the NIR video
CC_feed=VideoReader('Test_4_Colour.mp4');
total_frames_CC = CC_feed.NumberOfFrames;
%Get frame count of the colour video

%% detection setup
detect=zeros(10,total_frames_NIR-1); %Preload the memory for
detection statistics, not used under normal live operating
conditions.
Frame_trigger_IR=-1;
Frame_trigger_MIR=-1;
Frame_trigger_IR_PO=-1;
Frame_trigger_IR_PR=-1;
Frame_trigger_CC_CB=-1;
Frame_trigger_MCC=-1;
Frame_trigger_CC_PO=-1;
Frame_trigger_CC_PR=-1;
Frame_trigger_CM=-1;

%% Video Processing
width = 2; %Sets size of the detection dot in the image.
%nframes = 1; %Starts counter at first frame.
Fire_size_CC=0;   % Size of the fire as determined by the NIR
Emissions
Fire_size_IR=0;   % Size of the fire as determined by the Colour
processing
disp('The detection program will now begin. To terminate the
program, press control-c.')

for nframes = 1:total_frames_CC
%% Default values
        %These values are reset every frame
        dist_NIR = 10000;
% Distance between the center of the NIR Emission centroid and NIR
movement centroid, reset
        dist_CC= 10000;
% Distance between the center of the Colour Emission centroid and
Colour movement centroid
        dist_C=10000;
% Distance between the mean centers of Dist_NIR and Dist_CC

%% NIR_emission_processing
% This algorithm block looks for the fire based on NIR emmissions.
taggedfire_IR = read(IR_feed,nframes);
%Take frame from the NIR video feed
taggedfire_IR=rgb2gray(taggedfire_IR); %Convert image to greyscale
Z=taggedfire_IR>IR_Threshold; % Filter the image using the
fire_luminance, leaves only very bright objects. Refer to equation
13.
No_reflect = imopen(Z,disk); % Remove non disk shaped objects from
video frame
No_noise = bwareaopen(No_reflect, Noise_Cancellation_Threshold_IR);
%Remove noise and artifacts from the picture
stats_IR = regionprops(No_noise, {'Centroid', 'Area'}); % Find the
centroid and area of the remaining shapes in the picture
```

```
if ~isempty([stats_IR.Area]);
%If there are remaining objects in the picture
areaArray = [stats_IR.Area];
%Store the area of the objects in the picture
 [size,center] = max(areaArray); %Find the object with the largest
size, usually the flame due to the bloom.
b_IR = stats_IR(center).Centroid; %Centroid of the the largest
object in frame.
c_IR = floor(fliplr(b_IR)); %Flip left to right and round towards
negative infinity
row_IR = c_IR(1)-width:c_IR(1)+width; %Places dot in the center row
of the centroid area
col_IR = c_IR(2)-width:c_IR(2)+width; %Places dot in the center
column of the centroid area
Fire_size_IR= size %Sets fire size based on area of fire object in
picture.
anomaly_IR = 1; %Sets IR anomaly counter to 1 for this frame.
elseif isempty([stats_IR.Area]) % If there are no remaining objects
in the picture
anomaly_IR = 0; %Sets IR anomaly counter to 0 for this frame.
Fire_size_IR = 0% Recordeds fire size as 0
end
if all ([anomaly_IR==1,Frame_trigger_IR==-1]) %Records first frame
detection of IR anomaly
Frame_trigger_IR=nframes;
                end
detect(1,nframes)=anomaly_IR; %Records detection for further
analysis


%% MOVEMENT Processing_NIR
% This algorithm block looks for movement in the NIR camera.
taggedFire_MIR= double(No_noise); %Grabs image for the movement
detection from the completed NIR emission block.
fgMask_MIR = step(detector_MIR, taggedFire_MIR ); %Background
subtraction using detector
fgMnoise_MIR = bwareaopen(fgMask_MIR, 20); %Filtering of small
objects and noise from the image frame.
 [area,centroids, bboxes] = step(blob_MIR, fgMnoise_MIR); %Detecting
the area and centroid of the remaing blob in the image.
if ~isempty(area) %If the area is not empty
c_MIR=floor(fliplr(centroids));
row_MIR = c_MIR(1)-width:c_MIR(1)+width; %Places dot in the center
row of the centroid area
col_MIR = c_MIR(2)-width:c_MIR(2)+width; %Places dot in the center
column of the centroid area
anomaly_MIR = 1; %Sets NIR movement anomly to 1 for this frame.
elseif isempty(area) %If the area is empty (=0)
anomaly_MIR = 0; %Sets NIR movement anomly to 0 for this frame.
             end
if all ([anomaly_MIR==1,Frame_trigger_MIR==-1]) %Detects the first
movement detection made during the video
Frame_trigger_MIR=nframes;
            end
detect(2,nframes)=anomaly_MIR; %Records detection for further
analysis
```

```matlab
%% NIR_Camera_Combination_detection
% This algorithm block determines the severity of the fire
% detection made by the two NIR blocks.
if all([anomaly_IR==1,anomaly_MIR==1])
dist_NIR=sqrt((c_IR(1)-c_MIR(1)).^2+(c_IR(2)-c_MIR(2)).^2);
%Calculates the distance between the NIR emission blob and the NIR
movement blob
NIR_x=(c_IR(1)+c_MIR(1))/2; %Creates a center point based on the
position of the two points.
NIR_y=(c_IR(2)+c_MIR(2))/2;
end
if all([dist_NIR<200,10000<Fire_size_IR,Fire_size_IR<40000]) %Fire
size criteria, possible fire. Fire_size describes the size of the
blob in the NIR emssion frame.
NIR_detection_Possible=1;
NIR_detection_Probable=0;
elseif all([dist_NIR<200,Fire_size_IR>40000]) %Fire size criteria,
probable fire. Fire_size describes the size of the blob in the NIR
emssion frame.
NIR_detection_Probable=1;
NIR_detection_Possible=0;
else
NIR_detection_Possible=0;
NIR_detection_Probable=0;
end
if all ([NIR_detection_Possible==1,Frame_trigger_IR_PO==-1])
%Records the first possible fire frame detected
Frame_trigger_IR_PO=nframes;
elseif all ([NIR_detection_Probable==1,Frame_trigger_IR_PR==-1])
%Records the first probable fire frame detected
Frame_trigger_IR_PR=nframes;
end
detect(3,nframes)=NIR_detection_Possible; %Records detection for
further analysis
detect(4,nframes)=NIR_detection_Probable;

%% COLOUR Processing_CC
% This algorithm block processes the colour video based on the
% flame colour
taggedFire_CC=read(CC_feed,nframes); %Reads the video frame from the
video object.
taggedFire_CC=rgb2ycbcr(taggedFire_CC); %This step is not normally
needed. The video from the cameras in the live program is already
YCbCr
Y=taggedFire_CC(:,:,1); %Splitting the image into its three
Cb=taggedFire_CC(:,:,2);
Cr=taggedFire_CC(:,:,3);
Y_ave=mean(mean(Y)); % Take the mean of the three channels of the
image
Cb_ave=mean(mean(Cb));
Cr_ave=mean(mean(Cr));
a=Y>Cb; % Equation 7
b=Cr>Cb;
c=Y>Y_ave; %Equation 9
d=Cb<Cb_ave;
e=Cr>Cr_ave;
f=(c+d+e)>=3; %Adding the results of equation 9 into one logical
matrix
g=(Cr-Cb)>=CC_Threshold; %Equation 10
h=(a+b+f+g)==4; %Adds Equations 7,9 and 10 together into one logical
matrix
```

```
stats_Cb = regionprops(h_2, {'Centroid', 'Area'});
if~isempty([stats_Cb.Area]); %If there are remaining objects in the
picture
areaArray = [stats_Cb.Area]; %Store the area of the objects in the
picture
 [size,center] = max(areaArray); %Find the object with the largest
size
c_Cb = stats_Cb(center).Centroid; %Centroid of the the largest
object in frame.
c_Cb = floor(fliplr(c_Cb)); %Flip left to right and round towards
negative infinity
row_Cb = c_Cb(1)-width:c_Cb(1)+width; %Places dot in the center row
of the centroid area
col_Cb = c_Cb(2)-width:c_Cb(2)+width; %Places dot in the center
column of the centroid area
anomaly_Cb = 1; %Creates Colour anomaly. Cb is left over from when
the colour detection was split into two blocks Cb and Cr
Fire_size_CC= size;
elseif isempty([stats_Cb.Area]) % If there are no remaining objects
in the picture
anomaly_Cb = 0;
Fire_size_CC=0;
end
if all ([anomaly_Cb>=1,Frame_trigger_CC_CB==-1]) %Records first
frame detection of CC anomaly
Frame_trigger_CC_CB=nframes;
end
detect(5,nframes)=anomaly_Cb; %Records detection for further
analysis

%% MOVEMENT Processing_CC
% This algorithm block looks for movement in the CC camera.
taggedFire_MCC=double(h_2); %Grabs image for the movement detection
from the completed CC detection block
fgMask_CC = step(detector_CC, taggedFire_MCC ); %Background
subtraction using detector
fgMnoise_CC = bwareaopen(fgMask_CC, 20); %Filtering of small objects
and noise from the image frame.
 [area,centroids, bboxes] = step(blob_CC, fgMnoise_CC);  %Detecting
the area and centroid of the remaing blob in the image.
if ~isempty(area) %If a block is remains
c_MCC=floor(fliplr(centroids));
row_MCC = c_MCC(1)-width:c_MCC(1)+width; %Places dot in the center
row of the centroid area
col_MCC = c_MCC(2)-width:c_MCC(2)+width; %Places dot in the center
column of the centroid area
anomaly_MCC = 1;
elseif isempty(area)
anomaly_MCC = 0;
end
if all ([anomaly_MCC==1,Frame_trigger_MCC==-1]) %Records first frame
detection of MCC movement anomaly
Frame_trigger_MCC=nframes;
end
detect(6,nframes)=anomaly_MCC; %Records detection for further
analysis
```

```
%% Colour Camera Combined Detection
%This algorithm block determines the severity of the fire
%detection made by the two CC blocks.
if all([anomaly_Cb==1,anomaly_MCC==1])
dist_CC=sqrt((c_Cb(1)-c_MCC(1)).^2+(c_Cb(2)-c_MCC(2)).^2); %If both
blocks make detections, this finds the distance between the two.
CC_x=(c_Cb(1)+c_MCC(1))/2; %Creats a center point between the two
colour detections
CC_y=(c_Cb(2)+c_MCC(2))/2;
end
if all ([dist_CC<300,10000<Fire_size_CC,Fire_size_CC<20000]) %Fire
size criteria, possible fire. Fire_size describes the size of the
blob in the CC detection frame.
CC_detection_Possible=1;
CC_detection_Probable=0;
elseif all ([dist_CC<300,Fire_size_CC>20000]) %Fire size criteria,
probable fire. Fire_size describes the size of the blob in the CC
detection frame.
CC_detection_Possible=0;
CC_detection_Probable=1;
else
CC_detection_Possible=0;
CC_detection_Probable=0;
end
if all ([CC_detection_Possible==1,Frame_trigger_CC_PO==-1])
Frame_trigger_CC_PO=nframes;
elseif all ([CC_detection_Probable==1,Frame_trigger_CC_PR==-1])
Frame_trigger_CC_PR=nframes;
end
detect(7,nframes)=CC_detection_Possible; %Records detection for
further analysis
detect(8,nframes)=CC_detection_Probable;

%% COMBINED DETECTION
% This final block detemines if the two probable detections from the
fire detection tracks constitute a true fired detection
if all([NIR_detection_Probable==1,CC_detection_Probable==1])
dist_C=sqrt((NIR_x-CC_x).^2+(NIR_y-CC_y).^2); %Takes the distance
between the NIR and CC probable fire detections.
end
if dist_C<300
detection=1;
elseif dist_C>300
detection=0;
end
if all ([detection==1,Frame_trigger_CM==-1])
Frame_trigger_CM=nframes;
end
detect(9,nframes)=detection; %Records detection for further analysis
nframes = nframes+1; %Frame counter adds one at the end of the loop.
percent_done=(nframes/total_frames_CC)*100 %Displays percentage of
test done.
if percent_done == 100
beep;
end
```

```
%% Video statistics
% This section displays the number of frame detections and when the
first
% one was.
totals=sum(detect,2);
DETECTIONS_NIR_EMISSIONS=totals(1)
DETECTIONS_NIR_MOVEMENT=totals(2)
DETECTIONS_NIR_POSSIBLE_FIRE=totals(3)
DETECTIONS_NIR_PROBABLE_FIRE=totals(4)
DETECTIONS_CC_Emissions=totals(5)
DETECTIONS_CC_Movement=totals(6)
DETECTIONS_CC_POSSIBLE_FIRE=totals(7)
DETECTIONS_CC_PROBABLE_FIRE=totals(8)
DETECTIONS_COMBINED=totals(9)

FIRST_DETCTION_IR=Frame_trigger_IR
FIRST_DETECTION_IR_MOVE=Frame_trigger_MIR
FIRST_DETCTION_IR_POSSIBLE_FIRE=Frame_trigger_IR_PO
FIRST_DETCTION_IR_PROBABLE_FIRE=Frame_trigger_IR_PR
FIRST_DETECTION_CC=Frame_trigger_CC_CB
FIRST_DETECTION_CC_MOVEMENT=Frame_trigger_MCC
FIRST_DETCTION_CC_POSSIBLE_FIRE=Frame_trigger_CC_PO
FIRST_DETECTION_CC_PROBABLE_FIRE=Frame_trigger_CC_PR
First_DECECTION_COMBINED=Frame_trigger_CM
```

Appendix D: Software Testing results

Legend

A -1 means that no detection was made.

NIR Probable fire: Distance between centroids of NIR emission anomaly and NIR movement anomaly must be less than 200 pixels. Detected fire anomaly size must be bigger than 40000 pixels.

NIR Possible fire: Distance between centroids of NIR emission anomaly and NIR movement anomaly must be less than 200 pixels. Detected fire anomaly size must be bigger than 10000 pixels, smaller than 40000 pixels.

CC Probable fire: Distance between centroids of CC emission anomaly and CC movement anomaly must be less than 300 pixels. Detected fire anomaly size must be bigger than 20000 pixels.

CC Possible fire: Distance between centroids of CC emission anomaly and CC movement anomaly must be less than 300 pixels. Detected fire anomaly size must be bigger than 10000 pixels, smaller than 20000 pixels.

Combined detection: Must have both a CC and NIR probable detection in the frame. The distance between these probable fire detections must be less than 300 pixels.

Scenario 1

VIDEO 1	INFRARED THRESHOLD (TR) (%)				COLOUR THRESHOLD (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	150	150	147	107	147	147	147	147
NIR MOVEMENT	38	52	64	90	64	64	64	64
NIR POSSIBLE	0	0	0	0	0	0	0	0
NIR PROBABLE	0	0	0	0	0	0	0	0
CC EMISSIONS	0	0	0	0	0	0	0	0
CC MOVEMENT	0	0	0	0	0	0	0	0
CC POSSIBLE	0	0	0	0	0	0	0	0
CC PROBABLE	0	0	0	0	0	0	0	0
COMBINED	0	0	0	0	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	1	1	1	1	1	1	1	1
NIR MOVEMENT	112	99	74	43	74	74	74	74
NIR POSSIBLE	-1	-1	-1	-1	-1	-1	-1	-1
NIR PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
CC EMISSIONS	-1	-1	-1	-1	-1	-1	-1	-1
CC MOVEMENT	-1	-1	-1	-1	-1	-1	-1	-1
CC POSSIBLE	-1	-1	-1	-1	-1	-1	-1	-1
CC PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
COMBINED	-1	-1	-1	-1	-1	-1	-1	-1

Scenario 2

VIDEO 2 INFRARED THRESHOLD (TR) (%) COLOUR THRESHOLD (TC)

TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	750	750	750	750	750	750	750	750
NIR MOVEMENT	724	724	723	719	723	723	723	723
NIR POSSIBLE	236	235	247	240	575	575	575	575
NIR PROBABLE	468	455	443	400	463	463	463	463
CC EMISSIONS	6	6	6	6	252	48	6	0
CC MOVEMENT	6	6	6	6	44	35	6	0
CC POSSIBLE	0	0	0	0	0	0	0	0
CC PROBABLE	0	0	0	0	0	0	0	0
COMBINED	0	0	0	0	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	1	1	1	1	1	1	1	1
NIR MOVEMENT	2	2	5	26	5	5	5	5
NIR POSSIBLE	2	2	35	49	35	35	35	35
NIR PROBABLE	177	269	273	277	273	273	273	273
CC EMISSIONS	35	35	35	35	1	34	35	-1
CC MOVEMENT	35	35	35	35	31	34	35	-1
CC POSSIBLE	-1	-1	-1	-1	-1	-1	-1	-1
CC PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
COMBINED	-1	-1	-1	-1	-1	-1	-1	-1

Scenario 3

VIDEO 3	INFRARED THRESHOLD (TR) (%)				COLOUR THRESHOLD (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	398	397	395	391	395	395	395	395
NIR MOVEMENT	393	391	391	391	391	391	391	391
NIR POSSIBLE	4	10	28	49	28	28	28	28
NIR PROBABLE	378	355	332	314	332	332	332	332
CC EMISSIONS	232	232	232	232	543	338	232	105
CC MOVEMENT	144	144	144	144	325	179	144	99
CC POSSIBLE	12	12	12	12	27	20	12	17
CC PROBABLE	88	88	88	88	116	108	88	15
COMBINED	67	81	77	60	96	92	77	10
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	312	312	312	312	312	312	312	312
NIR MOVEMENT	44	312	312	312	312	312	312	312
NIR POSSIBLE	364	444	406	406	406	406	406	406
NIR PROBABLE	391	365	365	365	365	365	365	365
CC EMISSIONS	391	391	391	391	1	383	391	406
CC MOVEMENT	395	391	391	391	2	383	391	406
CC POSSIBLE	395	395	395	395	104	397	395	446
CC PROBABLE	405	405	405	405	391	391	405	447
COMBINED	405	405	405	405	393	393	405	480

Scenario 4

VIDEO 4	INFRARED THRESHOLD (TR) (%)				COLOUR THRESHOLD (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	750	728	625	449	625	625	625	625
NIR MOVEMENT	748	745	578	433	578	578	578	578
NIR POSSIBLE	249	190	202	241	202	202	202	202
NIR PROBABLE	219	181	144	109	144	144	144	144
CC EMISSIONS	370	370	370	370	371	370	370	366
CC MOVEMENT	328	328	328	328	283	311	328	332
CC POSSIBLE	39	39	39	39	10	16	39	55
CC PROBABLE	253	253	253	253	247	266	253	221
COMBINED	145	119	100	69	86	95	100	97
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	1	1	1	28	1	1	1	1
NIR MOVEMENT	3	3	3	104	3	3	3	3
NIR POSSIBLE	3	96	306	323	306	306	306	306
NIR PROBABLE	310	310	374	371	374	374	374	374
CC EMISSIONS	381	381	381	381	379	381	381	383
CC MOVEMENT	381	381	381	381	379	381	381	383
CC POSSIBLE	387	387	387	387	381	383	387	394
CC PROBABLE	391	391	391	391	383	386	391	404
COMBINED	391	504	512	512	383	386	512	505

Scenario 5

VIDEO 5	INFRARED THRESHOLD (TR)				COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	375	375	375	375	375	375	375	375
NIR MOVEMENT	212	327	259	299	259	259	259	259
NIR POSSIBLE	0	0	17	107	17	17	17	17
NIR PROBABLE	31	30	62	14	62	62	62	62
CC EMISSIONS	0	0	0	0	211	23	0	0
CC MOVEMENT	0	0	0	0	153	23	0	0
CC POSSIBLE	0	0	0	0	0	0	0	0
CC PROBABLE	0	0	0	0	0	0	0	0
COMBINED	0	0	0	0	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	1	1	1	1	1	1	1	1
NIR MOVEMENT	3	2	2	2	2	2	2	2
NIR POSSIBLE	-1	-1	49	2	49	49	29	49
NIR PROBABLE	3	7	9	55	9	9	9	9
CC EMISSIONS	-1	-1	-1	-1	2	123	-1	-1
CC MOVEMENT	-1	-1	-1	-1	2	123	-1	-1
CC POSSIBLE	-1	-1	-1	-1	-1	-1	-1	-1
CC PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
COMBINED	-1	-1	-1	-1	-1	-1	-1	-1

Scenario 6

VIDEO 6	INFRARED THRESHOLD (TR)				COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	500	500	500	500	500	500	500	500
NIR MOVEMENT	296	309	296	296	296	296	296	296
NIR POSSIBLE	21	7	2	1	2	2	2	2
NIR PROBABLE	0	0	0	0	0	0	0	0
CC EMISSIONS	0	0	0	0	2	0	0	0
CC MOVEMENT	0	0	0	0	0	0	0	0
CC POSSIBLE	0	0	0	0	0	0	0	0
CC PROBABLE	0	0	0	0	0	0	0	0
COMBINED	0	0	0	0	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	1	1	1	1	1	1	1	1
NIR MOVEMENT	204	18	204	204	204	204	204	204
NIR POSSIBLE	229	322	371	227	371	371	371	371
NIR PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
CC EMISSIONS	-1	-1	-1	-1	239	-1	-1	-1
CC MOVEMENT	-1	-1	-1	-1	-1	-1	-1	-1
CC POSSIBLE	-1	-1	-1	-1	-1	-1	-1	-1
CC PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
COMBINED	-1	-1	-1	-1	-1	-1	-1	-1

Scenario 7

VIDEO 7	INFRARED THRESHOLD (TR)				COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	637	507	352	206	352	352	352	352
NIR MOVEMENT	544	411	251	173	251	251	251	251
NIR POSSIBLE	32	19	3	0	3	3	3	3
NIR PROBABLE	0	0	0	0	0	0	0	0
CC EMISSIONS	432	432	432	432	723	522	432	181
CC MOVEMENT	338	338	338	338	543	402	338	129
CC POSSIBLE	1	1	1	1	55	48	1	0
CC PROBABLE	0	0	0	0	5	6	0	0
COMBINED	0	0	0	0	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	7	145	149	181	149	149	149	149
NIR MOVEMENT	11	159	227	311	227	227	227	227
NIR POSSIBLE	315	541	717	-1	717	717	717	717
NIR PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
CC EMISSIONS	1	1	1	1	1	1	1	4
CC MOVEMENT	5	5	5	5	2	4	5	4
CC POSSIBLE	117	117	117	117	4	15	117	-1
CC PROBABLE	-1	-1	-1	-1	102	86	-1	-1
COMBINED	-1	-1	-1	-1	-1	-1	-1	-1

Scenario 8

VIDEO 8	INFRARED THRESHOLD (TR)				COLOUR THRESHOLD VARIANCE (TC)			
TOTAL DETECTIONS (FRAMES)	60	70	80	90	20	40	60	80
NIR EMISSION	284	197	107	63	107	107	107	107
NIR MOVEMENT	118	63	31	14	31	31	31	31
NIR POSSIBLE	0	0	0	0	0	0	0	0
NIR PROBABLE	0	0	0	0	0	0	0	0
CC EMISSIONS	750	750	750	750	750	750	750	121
CC MOVEMENT	201	201	201	201	246	216	201	198
CC POSSIBLE	35	35	35	35	101	77	35	2
CC PROBABLE	0	0	0	0	3	1	0	0
COMBINED	0	0	0	0	0	0	0	0
FIRST DETECTION FRAME	60	70	80	90	20	40	60	80
NIR EMISSION	206	206	208	208	208	208	208	208
NIR MOVEMENT	210	213	213	213	213	213	213	213
NIR POSSIBLE	-1	-1	-1	-1	-1	-1	-1	-1
NIR PROBABLE	-1	-1	-1	-1	-1	-1	-1	-1
CC EMISSIONS	1	1	1	1	1	1	1	131
CC MOVEMENT	128	128	128	128	4	2	128	131
CC POSSIBLE	548	548	548	548	154	167	548	573
CC PROBABLE	-1	-1	-1	-1	573	574	-1	-1
COMBINED	-1	-1	-1	-1	-1	-1	-1	-1