TRILLIUM TECHNOLOGIES

The Promise of Artificial Intelligence for Bushfire Defence

RESULTS OF THE BUSHFIRE DATA QUEST 2020

Challenge Partner

Technology Partner

Research Partner

Research Partner

CONTENTS

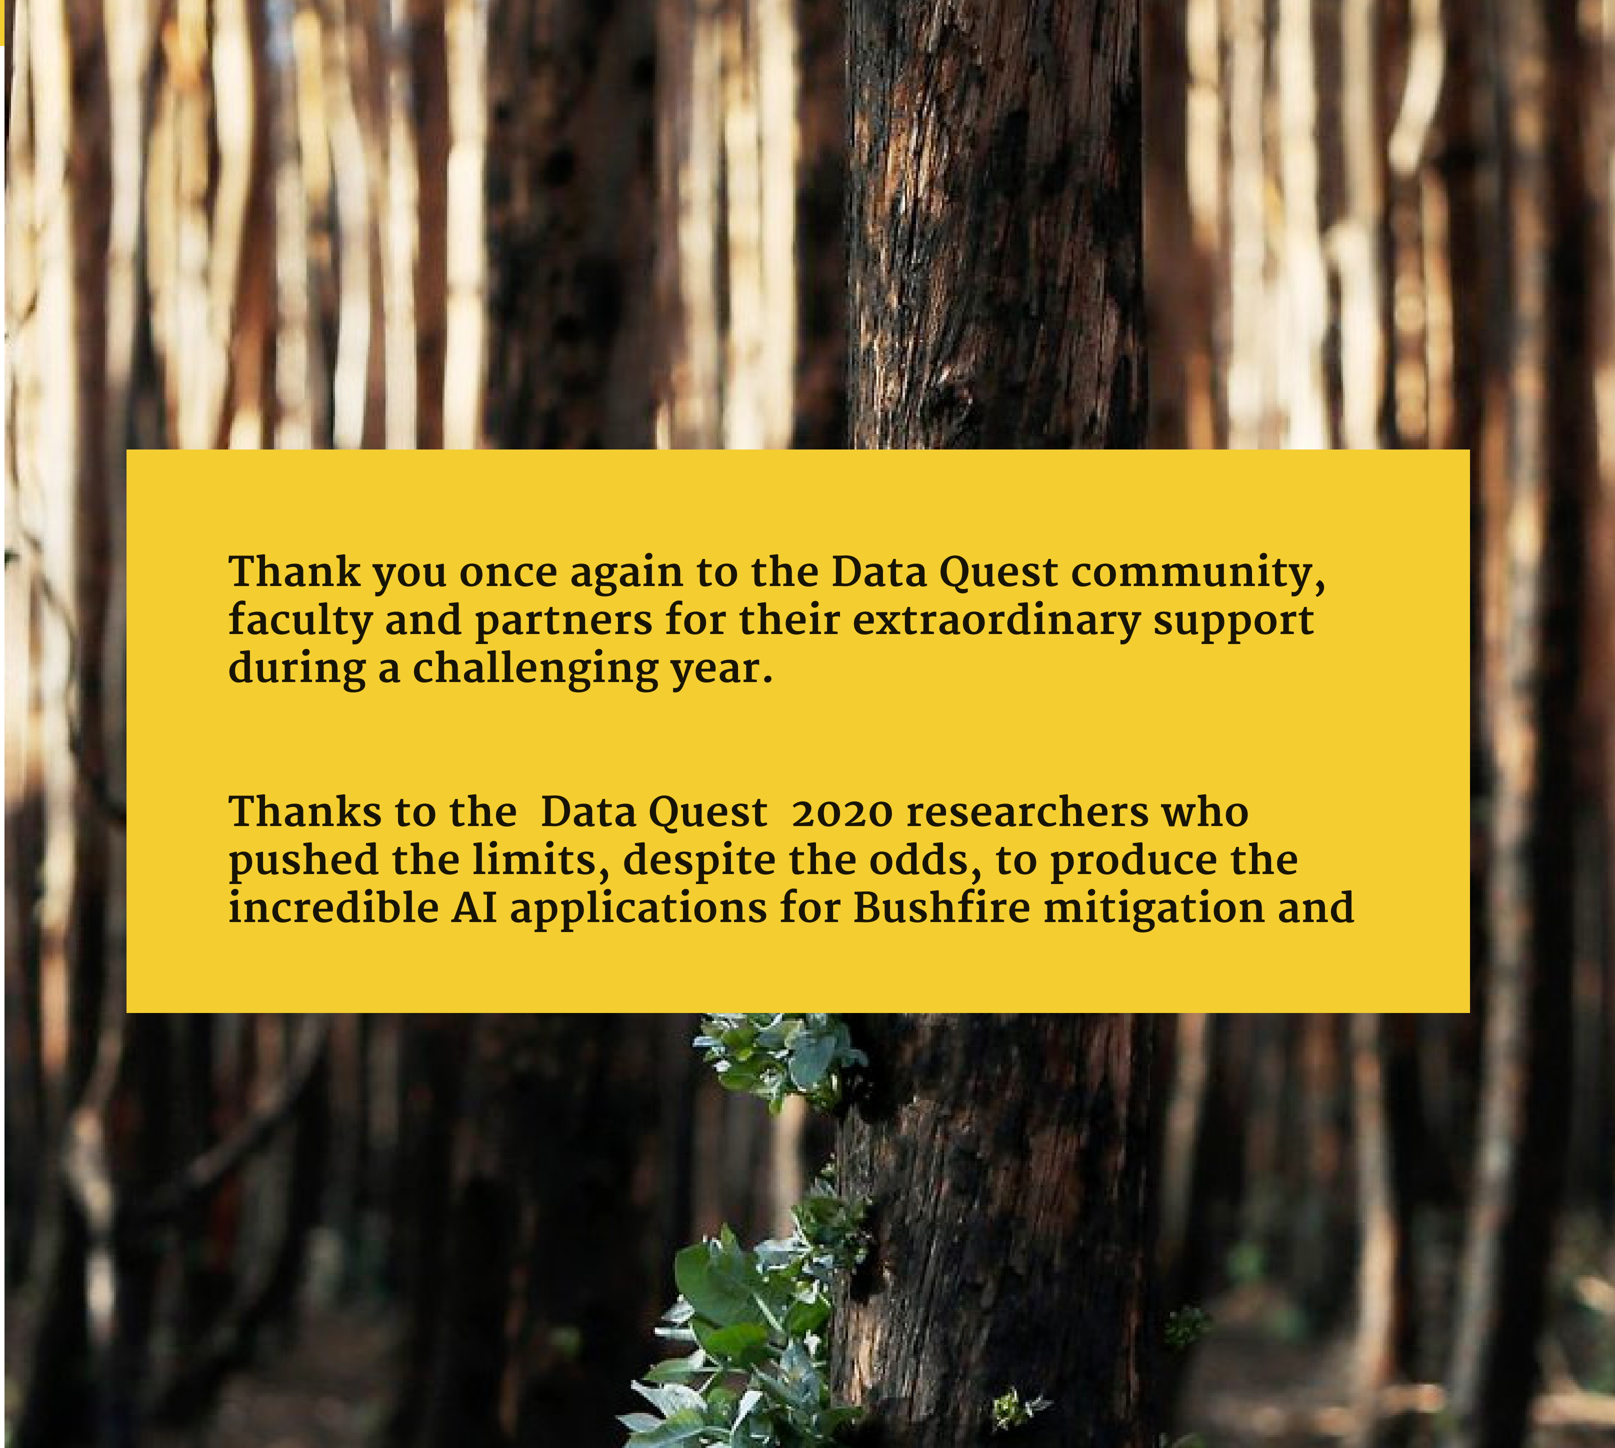

Thank you once again to the Data Quest community, faculty and partners for their extraordinary support during a challenging year.

Thanks to the Data Quest 2020 researchers who pushed the limits, despite the odds, to produce the incredible AI applications for Bushfire mitigation and

Minderoo Foundation

Challenge Partner

dug

Technology Partner

NSW Government

Research Partner

Australian National University

InSpace ANU Institute for Space

Research Partner

Australian Government | Australian Space Agency

Macquarie University Sydney · Australia

Australian Government Geoscience Australia

Northwest Nazarene University

CSIRO

UNSW Canberra

space UNSW Canberra

FIREBALL

planet.

Australian Government Bureau of Meteorology

THE PROMISE OF ARTIFICIAL INTELLIGENCE FOR BUSHFIRE DEFENCE

Artificial Intelligence (AI) is emerging as a particularly useful tool for detection and prediction of bushfires.

Once trained, AI pipelines are able to maintain vigilance over enormous data volumes and make accurate predictions on future events based on observations of the past. This is why the fusion of geospatial data from space- and ground-based sensors using AI holds so much promise for bushfire resilience and mitigation.

During the 2020 Data Quest the potential of space data and AI for bushfire defense was explored across the following challenge areas: (1) **Fuel Assessment** - in other words, the dryness and characteristics of the bush itself, which informs both the likelihood of a destructive fire and its progress; (2) **Early Detection** - crucial for rapid response and triage; (3) **Fire Behaviour** - how a fire evolves is critical for co-ordinated response planning.

In this document we will explore the potential for AI to tackle the following questions:

Can we use AI and satellite data to improve the resolution of Australia's fuel moisture map?

AI can produce moisture maps nearly 50 times more detailed and update them dynamically every five days.

Can we use AI and geostationary satellites (which sit over Australia and deliver an 'always on' view) to provide early warning of fire ignition?

Initial tests suggest that fire-warnings could come from geo-stationary satellites in near-real time. This demonstration has only been tested on historical data as of now, but the workflow is promising.

Once we have an ignition, can we help fire-fighters and communities understand the evolution of the bushfire?

The physics and co-factors of bushfire progress are complex, however, AI can potentially match the performance of models on supercomputers - which are stretched thin during fire season. This opens up the possibility of more localised prediction capabilities and faster predictions in the hands of first responders, giving decision-makers the possibility of a more tactical and effective response.

Can we determine the ferocity of a fire to help triage strategic decisions?

Fire clouds - pyrocumulonimbus – are visible from space and detectable with AI. These clouds are nature's way of saying, "things are getting serious", as the clouds create their own weather system around a fire. Early detection of this phase shift from space and the ability to warn fire-fighters and communities on the ground could be a lifesaver.

This work demonstrates the potential for AI and space data to help tackle one of humanity's most pernicious and urgent challenges.

In this Proceedings, we will cover these topics in more detail and expand on the implications for the 'bushfire-system' - the complex interaction of fire-ecology with the social and technical response architecture that bravely defends us each fire season.
This work is dedicated to all the bushfire and wildfire professionals who put their lives on the line, with a special thanks to everyone that provided the invaluable insight to ensure this work was fit for purpose.
We owe an enormous debt of thanks to our remarkable Data Quest researchers, who worked so hard to push the limits of what is currently possible.

We encourage feedback of all sorts.

Cormac Purcell

Program Director
Bushfires DataQuest 2020
Trillium Tech Pty Ltd

A NEW ERA OF MEGAFIRES

Wildfires have been front-and-centre in the public consciousness because of two unprecedented and catastrophic fire seasons.

The first in eastern Australia during 2019 / 2020, and the second in California, Oregon and Washington State in the late summer of 2020 - and at the time of writing still plaguing Colorado.

The emergency services and information infrastructure in both Australia and the USA have been hard-pressed to cope with the size and ferocity of the "megafires", leading to tragic loss of life, large numbers of citizens becoming internal refugees, and massive destruction of flora, fauna and property. In apocalyptic scenes, the populations of whole Australian towns were dramatically evacuated from beaches by the armed forces. In the US, fast-moving firestorms tore through towns, separating families and destroying homes and livelihoods. Similarly damaging fire seasons (especially in terms of pollution and health costs) were experienced in Siberia and Indonesia in recent years, even if these were less widely reported in the news.

Continued global heating is a major driver of the 2020 bushfires, which were also exacerbated by the El Niño/ La Niña weather oscillation. Climate models predict that fire seasons like those experienced in 2020 will be the new normal in the future.

The vital role of wildlands in ecosystem services.

Often overlooked in comparison to the human cost is the damage bushfires bring to Earth's vulnerable wildlands (what Antipodeans call 'The Bush'). Earth's wildlands provide habitat for around 6.5 million species according to the United Nations Environment Program. Wildlands contribute to energy development, recreational and spiritual opportunities for us humans and provide irreplaceable ecosystem services to the Earth system, such as clean water, nutrient cycling, pollination, and habitat for key actors in the food chain.

Effective management of prescribed fires is an essential step toward healthy and sustainable wildlands. Large expanses of bushland in Australia and elsewhere have evolved with fire and depend on periodic wildfires for regeneration.

A quantitative understanding of the relationships between fuel, fire behaviour, and the effects on human development and ecosystems can help stewards of wildlands develop nimble solutions to Australian bushfire problems.

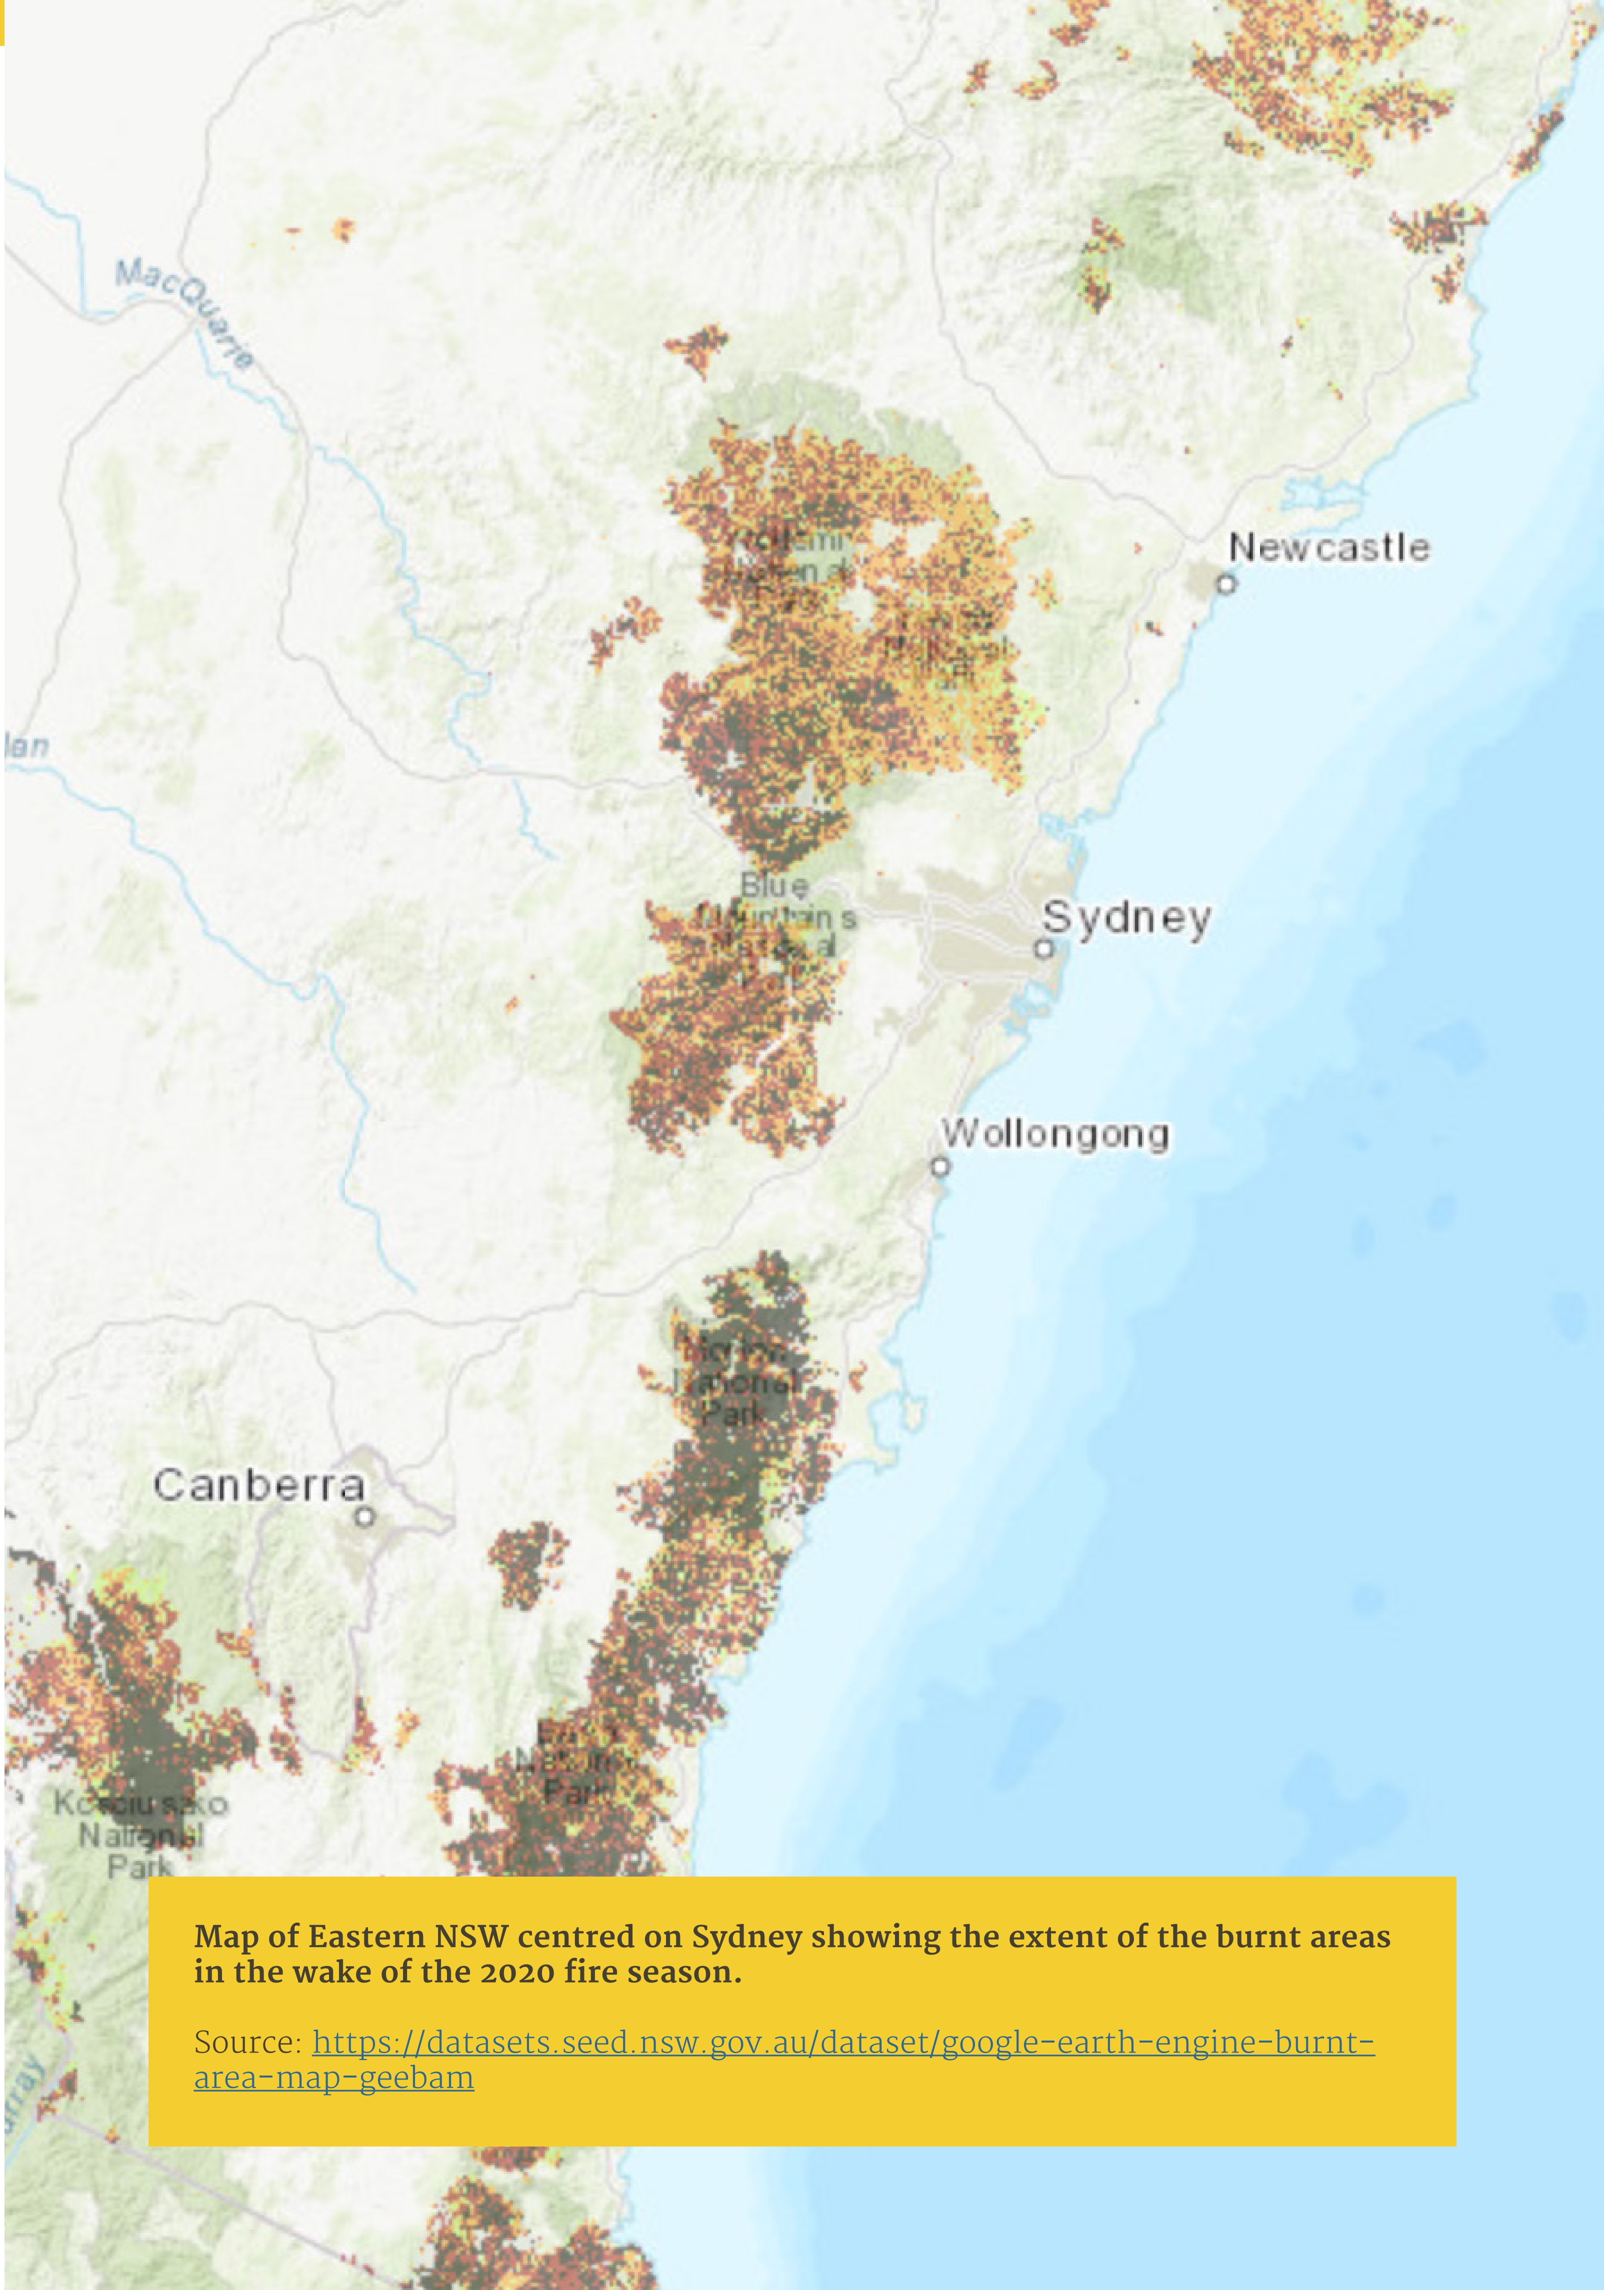

Map of Eastern NSW centred on Sydney showing the extent of the burnt areas in the wake of the 2020 fire season.

Source: https://datasets.seed.nsw.gov.au/dataset/google-earth-engine-burnt-area-map-geebam

NEW LEVELS OF INTENSITY REQUIRE NEW TOOLS

Fire consumes millions of hectares of Australian land each year, with the cost from the 2020 season alone estimated to be AU$100 billion. Whole towns have been scoured from the map and livelihoods have been destroyed, especially in regional areas.

High-intensity bushfires also contribute to post fire erosion, soil loss, flooding events and loss of timber resources. This results in negative impacts on wildlife habitat, ecosystem resilience, infrastructure, and recreational opportunities, in addition to knock-on effects for businesses tied closely to the land.

However, new capabilities in remote-sensing and artificial intelligence can help prepare for coming fire seasons, improve fire-fighting capability and boost the resilience of fire-affected communities, and ecosystems.

AI: ACTIONABLE INTELLIGENCE

There are very few places on Earth that are not photographed or otherwise 'measured' at least once per day.

Satellites image the whole globe, revisiting locations every few days or hours to record the changing landscape. Their cameras see deep into the infrared and ultra-violet, capturing information on chemistry, temperature and other physical conditions. Newer instruments can resolve details on centimetre scales, locating individual trees and buildings. At the same time, ground-based sensors like weather stations and lightning detectors record surface conditions unavailable from space. All of this data is stored - sometimes by commercial companies, but often by government agencies under open-access policies. However, there remain significant barriers for non-specialists to access and use this information as part of a project workflow.

For bushfire management, the real promise of this data is in translating it to actionable intelligence, directly answering queries like "**Which areas are at risk of ignition if struck by lightning?**" or "**Show me a map of the water level in streams wider than 1 meter.**" or "**The wind has just changed - where will the active fire spread to now**"?

To answer these questions requires the synthesis of disparate data-streams - imagery, time-series, spectral information, point-clouds and ad-hoc measurements. Machine learning (ML, a branch of artificial intelligence - AI) excels at this type of data-fusion.

Building a machine learning workflow is a specialist task requiring experience and knowledge to get right. However, the professional frameworks for building ML models (e.g., PyTorch, TensorFlow, Keras) are adaptable and have large libraries that enable fast prototyping. Once an appropriate ML algorithm has been identified (an art in itself), the first major task is to write a 'data-loader' that can translate input data into the required format.

Training the ML algorithm can take significant time and computing power, but once complete the model has the potential to make accurate predictions in seconds, on a portable device.

If deployed correctly, such ML algorithms could change how bushfires are prepared for, tackled and ultimately recovered from.

This proceedings document gives examples of four projects that leverage machine learning to address three key challenges managing bushfires in Australia and worldwide:

High-resolution and frequent Live Fuel Moisture maps from orbit.

Directly mapping fire risk and predicting fire-spread using satellite data.

Reliable early detection and localisation of new fires.

Investigating signatures of extreme fire behaviour.

Each of these projects is at an early stage of development (Technology Readiness Level 2 - 3) with proof-of-concept work done over an intense research sprint plus book-end periods. **When finalised and validated, these capabilities would enhance the ability to assess fire risk in the landscape on fine scales, detect fires earlier, predict the path of fire in response to dynamic weather conditions and potentially warn about extreme fires.**

TECHNOLOGY READINESS LEVEL

We assess the maturity of machine learning algorithms and their potential for deployment by using a scoring methodology based on the Technology Readiness Level (TRL) system.

TRL is a systems-engineering protocol created by NASA and DARPA in the USA for reliably developing, integrating and scaling complex interdependent technologies and cross-disciplinary teams.

The power of the TRL system is its ability to create a common language and set of standards between collaborators. Critical reviews and validation experiments are run at each level, firmly guiding the development of an idea, from the blackboards of researchers, via the test-bench, to an operational tool.

The TRL framework is useful for assessing near-term solutions, but also medium and longer term opportunities, and how they can fit together into an integrated system.

The Data Quest accelerated research sprint is designed to take proof-of-concept ideas to TRL 2 - 3 by showing how a challenge is solvable using machine learning and by documenting missing or highly desirable data.

Technology Readiness Levels for Machine Learning Systems

Alexander Lavin & Gregory Renard

https://arxiv.org/pdf/2006.12497.pdf

TRL **9** **Actual System Proven in Operation**	TRL4ML **9** **Deployment** Monitoring the current version, improving next
TRL **8** **System Complete and Qualified**	TRL4ML **8** **Flight-ready** The end of system development
TRL **7** **System Prototype Demonstration**	TRL4ML **7** **Integrations** ML infrastructure, product platform, data pipes, security protocols
TRL **6** **Technology Demonstrated in Relevant Environment**	TRL4ML **6** **Application Development** Robustification of ML modules, specifically towards one of more use-cases
TRL **5** **Technology Validated in Relevant Environment**	TRL4ML **5** **Machine Learning Capability** The R&D to product handoff
TRL**4** **Technology Concept Formulated**	TRL4ML **4** **Proof of Concept Development** Demonstration in a real scenario
TRL **3** **Experimental Proof of Concept**	TRL4ML **3** **System Development** Sound software engineering
TRL **2** **Technology Concept Formulated**	TRL4ML **2** **Proof of Principle Development** Active R&D is initiated
TRL **1** **Basic priniciples observed**	TRL4ML **1** **Goal-Oriented Research** Moving from basic principles to practical use

BUILDING A COMMUNITY: **FROM OUR PARTNERS**

The Bushfire Data Quest is as much about building a community of partners as it is about research. Each organisation in our network holds pieces of the puzzle and it is only by working together that we can realise our mutual benefit.

“

Minderoo Foundation Wildfire and Disaster Resilience Program is excited to be the Challenge Partner of the Bushfire Data Quest 2020. With this year’s fire season fast approaching and communities still feeling the disastrous effects of Australia’s bushfire crisis last summer make our collective work - including piloting new ways of preparing for and responding to disasters more vital than ever.

Minderoo Foundation’s experience working on the ground in fire affected communities this year has highlighted that existing technology and systems leave us ill-equipped to deal with large-scale disasters like fire, flood and other unforeseen challenges. We know further threats and shocks will come in the future, and this is why we need new approaches to lift national resilience. The Bushfire Data Quest 2020 builds on the excellent work many are already doing in this space, tackling change through the system as a whole, and supporting out of the box thinking and solutions. I’m pleased that we can support the great minds that have come together to solve parts of the problem and trial new methods to overcome the grand challenges we face in this area.

”

MINDEROO FOUNDATION

Adrian Turner
CEO, Minderoo Foundation Fire Fund

“

Using AI and geospatial data, the Bushfire Data Quest has the potential to unlock new ways to protect Australian, communities and wildlife from devastating bushfires.

We are proud to support this collaborative and innovative effort which has the potential to enhance worldwide bushfire management and prevention.

We look to support initiatives with the potential to solve profound problems and save lives. We are honoured to be part of the extremely valuable work of the Bushfire Data Quest. With only a few months until the next bushfire season, this research sprint could have an important role in protecting our nation.

”

Australian National University

InSpace
ANU INSTITUTE FOR SPACE

"

As the clock ticked over to welcome in the year 2020, we watched in horror as bushfires ravaged our country. More than 20% of our native forests burned, an estimated one billion Australian animals perished, lives were lost, and homes destroyed. All that we held so dear as Australians was severely damaged.

Data Quest aims to push the frontiers of research and develop new tools to help solve some of the biggest challenges that humanity faces. This year's Bushfire Data Quest 2020 is indeed addressing a challenge close to our hearts.

The team at DUG is excited to be part of a collaboration of industry, technology, science and research that will push the boundaries of conventional thinking in the quest to discover real solutions to improve our planet. We are honoured to provide the technology platform that will be the launching pad of truly powerful thinking.

"

dug

"

In a country with a high bushfire risk, satellite images and data play an increasingly vital role in protecting our communities and environment. Not only do satellites help detect bushfires, they can allow us to predict their movement and assess the damage they cause.

The NSW Government is supporting the development of local capability through our $5 million Space Industry Development Strategy, generating investment and jobs in this important sector and contributing to safety and well-being in our fire-prone landscapes. We're proud to support FDL's Bushfire Data Quest and look forward to seeing the unique solutions developed to manage and prevent bushfires.

"

NSW GOVERNMENT

The Hon. Stuart Ayres MP
Minister for Jobs, Investment, Tourism and Western Sydney

"

Bushfire Data Quests collaborative, data-led approach will not only consolidate our historic knowledge in this crucially important environmental area, but also allow for the development of strategies to plan for, and mitigate, future bushfire events.

"

NSW GOVERNMENT

Professor Hugh Durrant-Whyte
NSW Chief Scientist & Engineer

BUILDING A COMMUNITY: **FROM OUR PARTNERS**

"

In response to the devastating Australian bushfires in January the Minister for Industry, Science and Technology, the Hon Karen Andrews MP, tasked the Australian Space Agency to consider the role of space-based Earth Observation to support planning, response and recovery efforts related to bushfires. We now have the opportunity to further innovate how we use satellite data through Earth Observation imagery. In partnership with the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Geoscience Australia (GA) and the Bureau of Meteorology, the Australian Space Agency has established a taskforce which has engaged with emergency management agencies, state and territory governments, and the research community. The taskforce is working together to understand this issue and consider opportunities for the future, including how we can effectively apply satellite data to mitigate the risk of bushfires occurring and further enhance our response during bushfires.

"

Australian Government

Australian Space Agency

"

UNSW Canberra Space gets out of bed each day to develop and demonstrate the art of the possible for combining artificial intelligence and space technologies, to help meet challenges and opportunities on the ground. Intelligent space systems offer a pathway to rapidly turn remotely sensed data into actionable information, piped directly to the user. The Bushfire Data Quest is an excellent specific example of the broader opportunity offered by FDL, to develop both the science and the talent pool that are needed for that future, and UNSW Canberra Space is excited to be closely involved.

"

UNSW CANBERRA

space UNSW CANBERR A

“

The recent bushfire seasons have been startlingly severe. The impacts are felt throughout the entire community. At Fireball we use sensor fusion and machine learning to detect bushfires within minutes of ignition and to model bushfire behaviour. It is exciting for Fireball to be a part of Data Quest and to contribute to the development of new technologies that may help us better understand and respond to bushfire risks.

”

FIREBALL

“

The Bushfire Data Quest represents a fantastic opportunity to advance our thinking, and to help realise the true potential of near real-time Earth Observation applications to emergency management.

”

Australian Government

Geoscience Australia

“

Macquarie University is delighted to support Data Quest in their first Australian Data Quest: Detecting and Responding to Bushfires. Data Quest innovative approach exemplifies the tradition of cross-disciplinary, industry-engaged research upon which Macquarie University prides itself and we are excited to see the ideas and solutions generated at this ground breaking event.

”

MACQUARIE University

“

Northwest Nazarene University is pleased to partner with the Bushfire Data Quest. This is an excellent opportunity for the faculty and students of NNUs College of Natural & Applied Sciences to share the benefits of our ongoing applied research in disciplines as diverse as Fire Ecology, Astronautics, Biomedical, Agriculture and Archaeology. We look forward to continuing our partnership in the Bushfire Data Quest, with the goal of assisting in the development of solutions specific to the problems faced with management of bushfires in Australia and New Zealand.

”

NORTHWEST NAZARENE UNIVERSITY

AI FOR MEGAFIRES

Artificial intelligence has matured to the point where it can compliment, and improve on, the physics-based bushfire modeling methods that have been employed for the past 50 years.

The enhanced capabilities of these new AI tools (sometimes called ecoinformatics) will position fire managers to better leverage the dramatic increase in information at their disposal in the next decade.

The use of AI, along with data from newer satellites with higher resolution sensors, will enable countries such as Australia to lead the world in fire science and prediction. Traditional modelling methods, such as those used in the US, have been based on older empirical measurements and data from legacy observation platforms, such as Landsat. AI-based tools will allow Australia and other countries to rapidly integrate the latest data into more reliable predictions - without having to expend resources at the level that has been necessary in the past.

The new methods, tools and datasets that are developed to manage bushfires in Australia can also be leveraged to address the wildfire problems faced by countries around the world. This will result in better planning, suppression and mitigation responses to bushfire globally and its effect on ecosystems worldwide.

The revolution in the use of AI is being driven by technological advances that are still accelerating today.

ADVANCES

Advances in compute

Advances in computing hardware, algorithms and tools are making analytic capabilities possible in 2020 that were relegated to science fiction just a few years ago. There has been an increasing effort to direct those analytic capabilities toward bushfire management, enabling land managers to better predict future fire behaviour and effects. These improved analytic approaches also promise to guide fuel treatment strategies, mitigating the severity and extent of future bushfires.

Advances in data availability

As computing capabilities have increased, so also has the amount of data to which we can apply artificial intelligence to derive actionable knowledge. A significant source of rich data are Earth-observing satellites. In March 2020, the one-hundred-millionth Landsat scene was downloaded from the USGS web portal, which first started offering Landsat imagery for download at no-cost starting Oct 2008. Landsat 'scenes' have sizes of up to 1.5 gigabytes each and almost 150 petabytes (1.5 x 1017 bytes) of imagery have been downloaded to date. This is not to mention data that's been collected and downloaded from other satellite systems such as Sentinel, Himawari, and Planet Labs.

Advances in resolution and fidelity

The increased resolution of the latest remote sensing platforms has already been shown to lead to more accurate fire-related data products compared to older systems (Hamilton 2017a, Hamilton, 2019). The accuracy of the analytics will continue to improve as we launch the next generation of satellites into orbit, offering even higher spatial, temporal and spectral resolution. The resulting flood of data will make the need for AI tools even more acute - this is because AI is uniquely capable of extracting actionable insight from large and heterogeneous datasets. As fire seasons become more intense, timely knowledge from these data and tools will be a step-change in our ability to combat bushfires.

Advances in machine learning methods for bushfire mitigation

A recent review of potential machine learning applications for bushfires (Jain et al, 2020), identified potential applications for several types of ML models. For supervised learning models, applications included fire susceptibility, fire spread/burn area prediction, fire occurrence, severity, smoke prediction and climate change effects. For unsupervised learning models, applications included fire detection, fire mapping, burned area prediction, fire weather prediction, landscape controls on fire, fire susceptibility, and fire spread/burn area prediction. Agent-based learning models were identified with applications in optimising fire simulators, fire spread and growth, fuel treatment, planning and policy, and wildfire response. The 2020 Data Quest investigated several of these research opportunities and ML applications, but there are many more opportunities that can be investigated.

Fuels

AI applied to higher resolution imagery (spatial and spectral) will enable the generation of high-fidelity fuel mapping layers. The increased information content will provide bushfire analytic tools with geospatial layers representing the fuel state with higher accuracy and more frequent cadence than currently available. This improved fuel data is the key driver of fire fisk maps and fire behaviour models.

Behaviour

ML-models can make rapid predictions of what an existing fire will do, adapting quickly to changes in the weather conditions. This is critical information for ground forces fighting campaign fires as early warnings of changes will likely save lives. Before each fire season, predictive analytics will inform fire managers on the best fuel treatments to apply by modelling their impact on the behavior of a potential future bushfire. On longer timescales, behaviour models will help determine what affects climate change is likely to have on ecosystems during future fire seasons.

Post-fire effects measure the change that a bushfire has on an ecosystem, focusing on the days, weeks and years following the fire. Improved bushfire analytics provide an intriguing opportunity to provide actionable insights to fire managers, enabling them to develop and deploy post-fire recovery strategies on a burned area in a timely manner. Improved knowledge of the immediate effects will efficiently enable managers to prescribe mitigation strategies that will provide for quicker recovery of the burned area, resulting in more resilient ecosystems.

Detection

Timely localisation of new bushfire ignitions is a critical component of the toolset that can be built to improve management of our wildlands. The ability to maximise both the temporal and spatial resolution of detections can greatly reduce the amount of time from initial ignition to the start of suppression efforts. These ecoinformatic tools have the capability to improve the initial response times of the fire service crews, enabling the initiation of suppression efforts even quicker than is currently possible. These improvements can be accomplished by leveraging the data streams of current and future generations satellites, augmented with a network of land based cameras.

Where can we go from here?

One of the primary accomplishments of the 2020 Data Quest was to show what could be accomplished by bringing the brightest AI engineers together with exceptional fire ecologists, determining what is ‘solvable’ using currently available tools, but also where we need to invest to close data-gaps.

The findings are discussed in the following section.

DATA FOR AI AND BUSHFIRE MANAGEMENT

AI is only as good as the data it has to work on.

A promising range of Earth-observation, geospatial and environmental data exists in Australia and New Zealand, that is suitable for use with AI-enhanced bushfire response systems. This includes multi-spectral imagery, synthetic aperture radar, LIDAR measurements of elevation and tree-cover, digital elevation models, vegetation type maps, weather station measurements (temperature, humidity, pressure), historic fire boundaries, historic fire progression maps, lightning ground-strike detections and fire ignition points.

All of these datasets inform models of fire risk in the landscape, or predictions of fire behaviour. Some datasets can be used to detect potential ignition sources and fires just starting to burn. However, to complete the vision of AI enabled bushfire resilience and mitigation, more investment is needed.

"AI ready data" is different from scientific data.

Data veracity, availability and completeness are the holy trinity of Artificial Intelligence. Incomplete data and outdated documentation complicate effective alignment of datasets and can add weeks of painstaking effort to a scientific machine learning project, with significant increases to research overhead and cost.

Input data needs to be assessed for missing or confounding values, such as blank or misaligned swaths, which can erroneously draw the attention of ML algorithms during the training phase. Similarly, noise and redundancy (causing data or class imbalance) also confound model performance. Supervised learning methods require well-labelled data to train models, while self-supervised methods need large volumes of examples. However, often labeling is non-existent or incomplete. This is further compounded by lack of enough data for a model to learn any patterns within the corpus to enable self-supervised methods to be applied. Synthetic data can help here, however, if ground truth is unavailable, any ML outcome is problematic to validate.

Oftentimes project data may require significant pre-processing to be ready for a machine learning pipeline. Data is also often widely distributed, difficult or costly to aggregate, requiring permissions and negotiations with data curators as part of the ETL (Extract / Transform / Load) process - a further time sink. In this way, cleaned, pre-aggregated and pre-processed data saves time and allows researchers to focus on the science. However, some emerging ML techniques require raw data to extract as much information as possible - critical for concepts such as autocorrection and super-resolution.

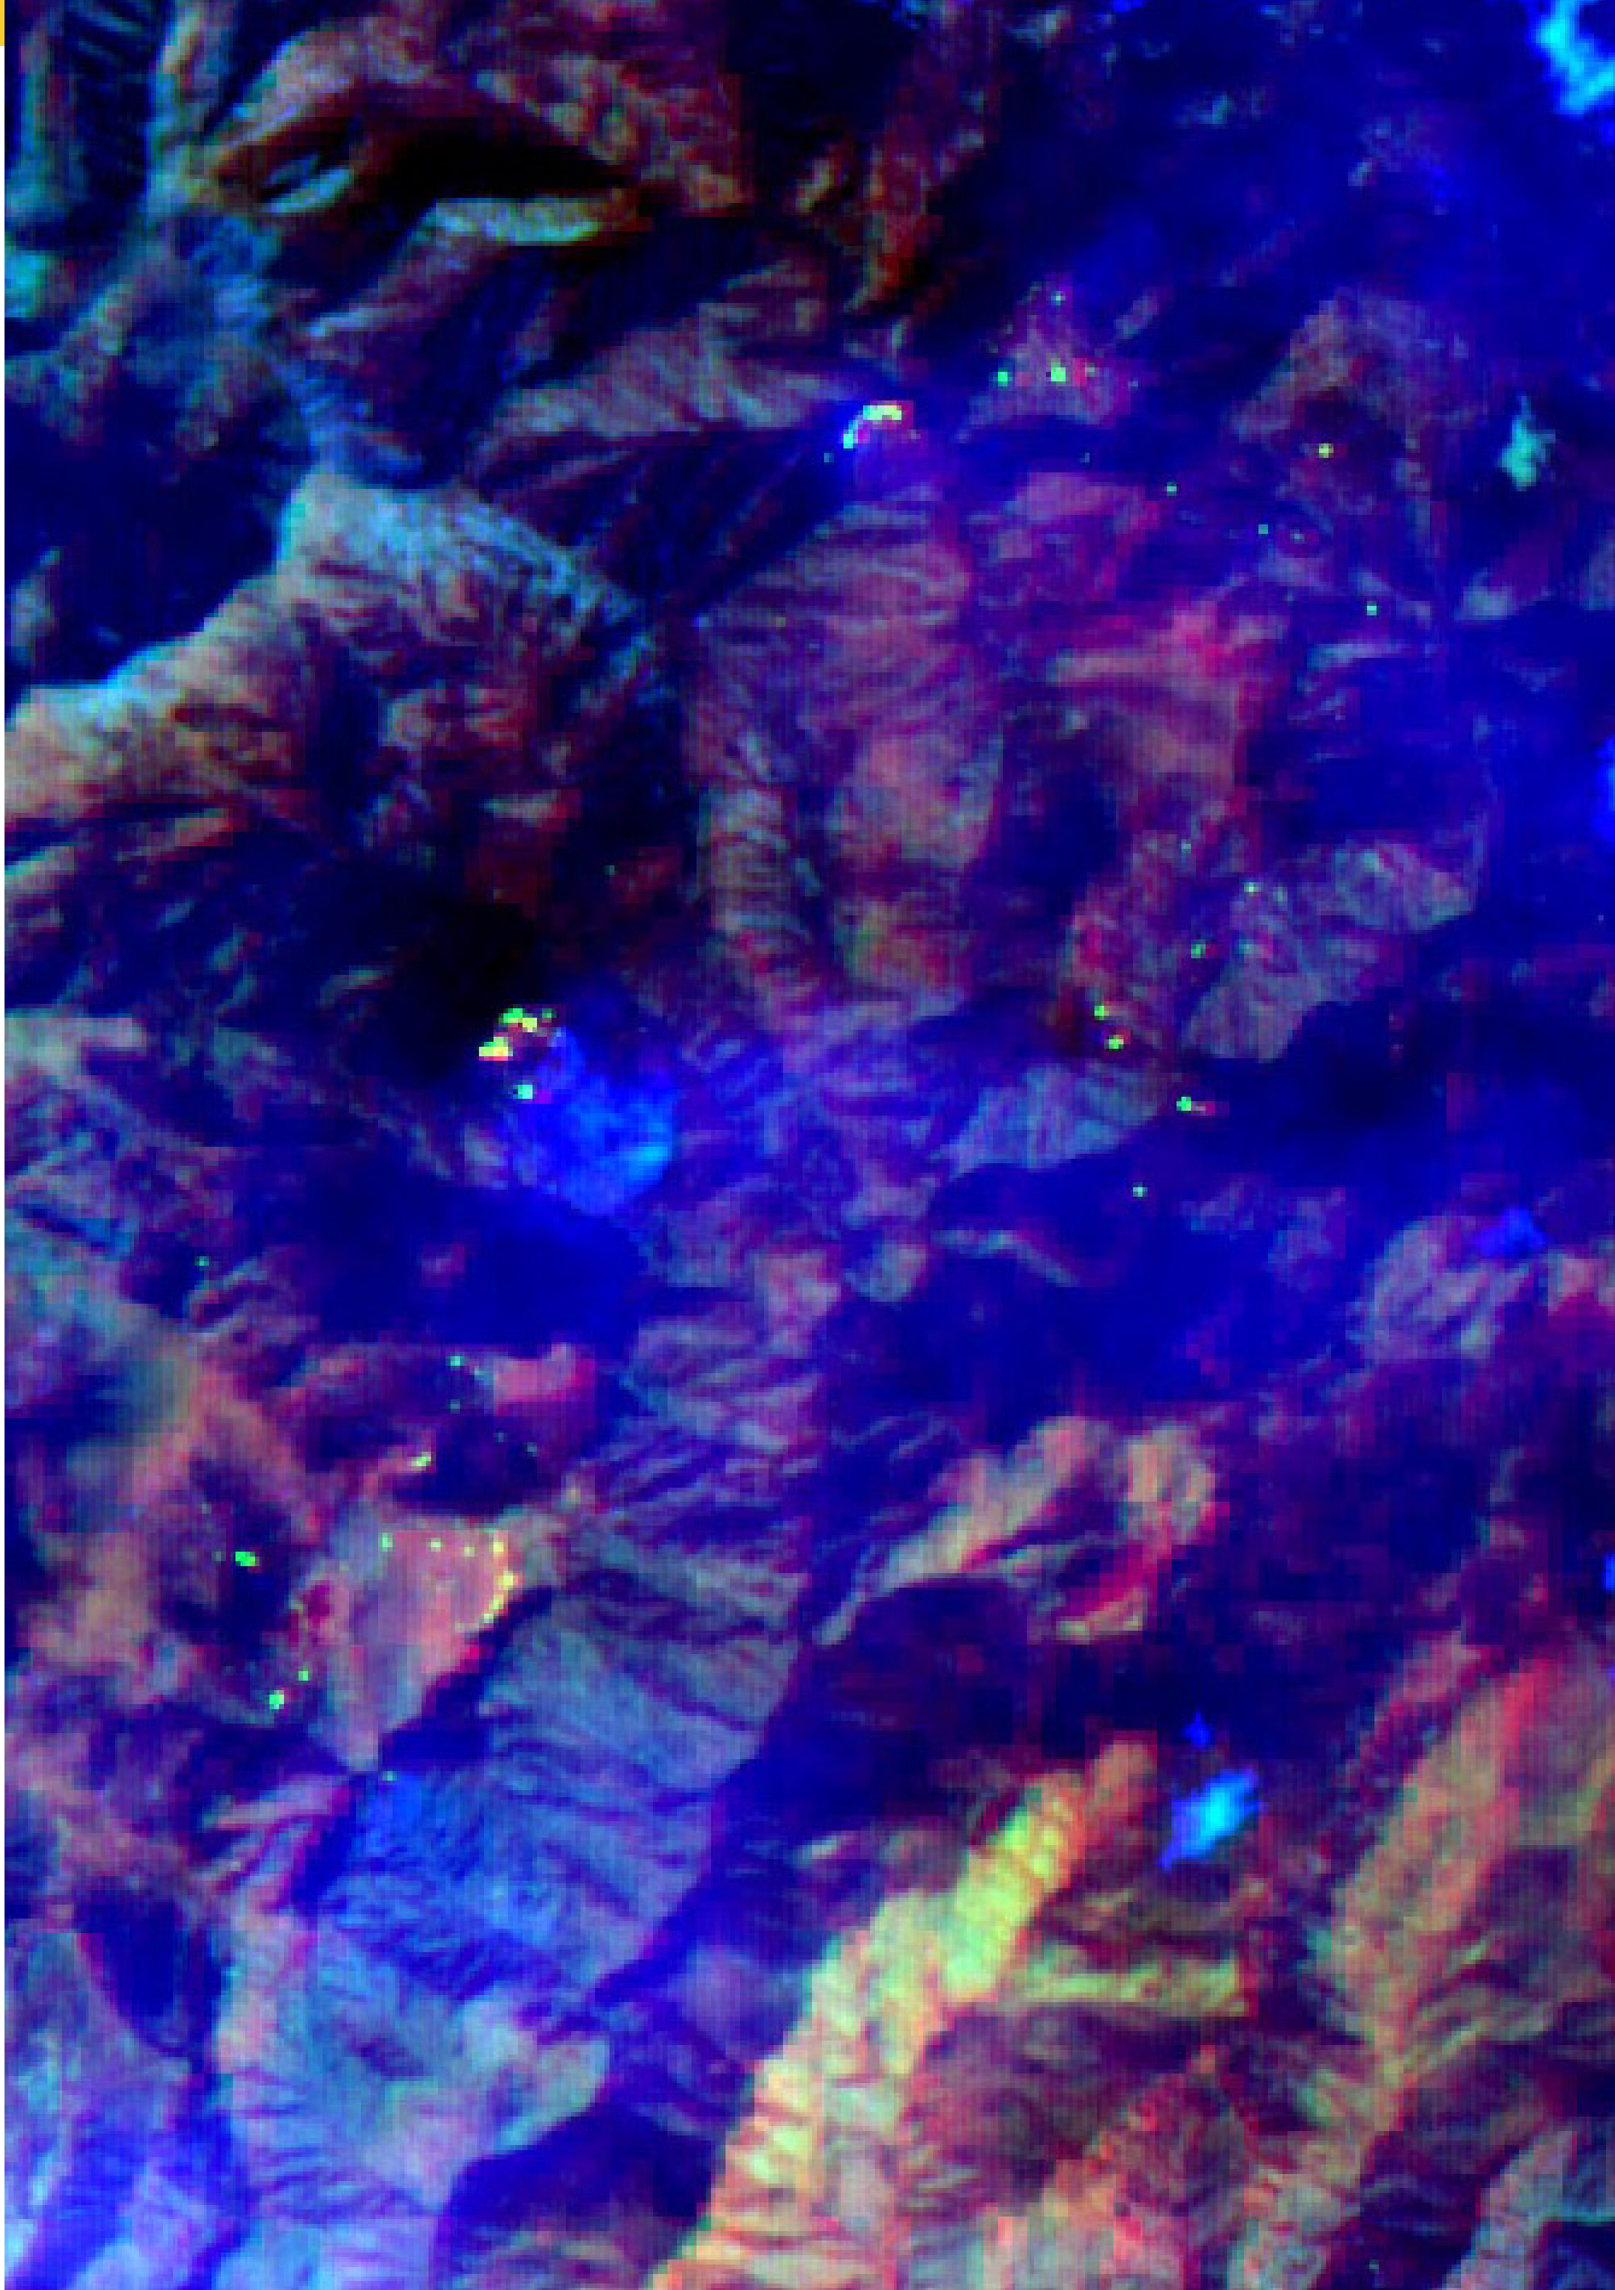

DATA FOR AI AND BUSHFIRE MANAGEMENT

Deployed ML-systems need maintenance in the face of data drift

Establishing performance metrics and an agreed range of output parameters are key to establishing veracity of predictions, and for AI ethics. Deployed systems also need to have continuous access to a complete and well-maintained flow of data. The data itself can often drift (distribution drift in properties of the data and concept drift in the use of the data). Understanding critical confounding factors, such as instrument drift, noise and other variables is also vital before model pipelines are built. Monitoring data quality during acquisition is also important, as consistency is key in model validation.

Moreover, AI systems are rarely set-up to adapt to these dynamic drifts and this remains a key impediment to the efficacy of machine learning in the real world and often the cause of embarrassing failures that get worse as time goes on.

AVAILABLE GEOSPATIAL DATA SOURCES FOR BUSHFIRE IN AUSTRALIA & NEW ZEALAND

The table below presents a list of data and data custodians. **Of particular note is the Japanese Space Agency Himawari-8 geostationary weather satellite that records kilometre-resolution images of Australia every 10 minutes and was used by two of the Data Quest research teams.** The European Space Agency Sentinel-1 (synthetic aperture radar) and Sentinel-2 (optical & infrared) satellites are also notable for imaging Australia every 5 days.

Himawari-8

PARAMENTERS	AVAILABILITY
Geosynchronous Earth orbit (GEO) at 140.7E longitude, **16 band multispectral imager** (visible, infrared), **500 m resolution**, full disk images **every 10 min**, imaging in selected areas at 2.5 min or 0.5 min frequency	Japan Meteorological Agency Australian Bureau of Meteorology (BOM)

Sentinel-1

PARAMENTERS	AVAILABILITY
Polar low Earth orbit (LEO), **sub-synchronous** (SSO), **Synthetic Aperture Radar** (SAR) images, C-Band, dual polarisation capability, four imaging modes, resolution (**down to 5 m**) and coverage (up to 400 km), **2-4 day** repeat coverage	European Union Copernicus Programme Sentinel Australasia

Sentinel-2

PARAMENTERS	AVAILABILITY
Polar LEO SSO, **13 band multispectral** (visible, near infrared, short wave far infrared), **10-60 m resolution**, 290 km swath width, **3-4 day** revisit coverage	European Union Copernicus Programme Sentinel Australasia

Digital Earth Australia

PARAMENTERS	AVAILABILITY
Digital Earth Australia program / Open Data Cube	- Geoscience Australia

Landsat 8

PARAMETERS
Polar LEO SSO, **11 bands** (visible, infrared), **30 m resolution**, **16 day revisit time**, Geoscience Australia Landsat Analysis Ready Data (ARD)

AVAILABILITY
US Geological Survey (USGS)

Geoscience Australia

Hotspots

PARAMETERS
Combined data from sensors **AVHRR, MODIS, VIIRS** and **Himawari** hotspots (but not the source data used to create them – only the point data and attributes).

AVAILABILITY
Digital Earth Australia

Elevation Data

PARAMETERS
Digital elevation model (DEM) data, **Shuttle Radar Topography Mission** (SRTM) 1 second elevation models

AVAILABILITY
NASA, USGS

NCI ANU

MODIS

PARAMETERS
Moderate Resolution Imaging Spectroradiometer (MODIS) on NASA Terra (EOS-AM) and Aqua (EOS-PM) satellites, **Polar LEO SSO**, **36 bands** (visible, infrared), **250 m to 1 km resolution**, **2-3 day** repeat cycle

AVAILABILITY
NASA

USGS

VIIRS

PARAMETERS
Visible Infrared Imaging Radiometer Suite (VIIRS) sensor on Suomi National Polar-orbiting Partnership (NPP) and National Oceanic and Atmospheric Administration (NOAA)-20 spacecraft, **Polar LEO SSO**, **22 bands** (visible, infrared), 3000 km swath width, **375 – 750 m resolution**

AVAILABILITY
NASA
NOAA

NSW Fire History map

PARAMETERS
NPWS Fire History – Wildfires and Prescribed Burns. FireHistory is a feature class that holds final fire boundaries for every year for which there is data.

AVAILABILITY
NSW Data Portal

Historic Australian Climate Data

PARAMETERS
SILO is a database of Australian climate data from 1889 to the present. It provides daily meteorological datasets for a range of climate variables in ready-to-use formats suitable for biophysical modelling, research and climate applications.

AVAILABILITY
SILO is hosted by the Queensland Department of Environment and Science

DATA GAPS

Experience gained during the Data Quest has highlighted missing data, or lack of capability, that are urgently needed to fulfil the promise of AI-driven bushfire management.

These 'data gaps' were discovered by the research teams while pursuing their projects and highlighted as being critical to future work. The table below records the gaps and suggested targets to aim towards. Some gaps are easy to fill by exposing currently available data, or by processing data into more digestible forms. Closing other gaps requires new measurements or sensors.

Some critical data is only available from commercial providers and should be purchased on behalf of organisations building bushfire management systems, or open alternatives developed.

The table below presents a list of data gaps identified during the Bushfire Data Quest 2020.

DATA NEEDED

Ignition Point Data

RECOMMENDATION

Temporal maps of ignition points would have greatly improved the fire spread mapping project. A large ignition point dataset was previously accessible via the NSW Department of Planning, Industry and Environment Data Portal, but was unavailable due to quality control issues.

DATA NEEDED

Lightning strike data

RECOMMENDATION

The majority of rural fires are started by lightning strikes, so temporal maps of ground-strike data would be very useful for many ML projects. However, the current best measurements of lightning strikes are made by commercial companies in Australia and are not freely accessible.

DATA NEEDED

Early Detection: High frequency and high resolution GEO satellite data.

RECOMMENDATION

The current source of images for early detection came from the Himawari-8 satellite in GEO. The data was valuable for this purpose, but there were challenges encountered. There were occasional gaps in coverage. The standard refresh rate was 10 minutes, with selected areas covered more frequently at 0.5 or 2.5 minutes. During fire season, targeting these selected areas toward high risk regions for higher frequency sampling would be very valuable if the data collection were combined with the appropriate analysis algorithms. The current sensor provided images at 1 km resolution.

Sensor designs for future GEO spacecraft should consider the need for bushfire detection applications to achieve higher frequency and higher resolutions at the appropriate wavelength bands.

DATA NEEDED

Fuel Moisture: Live Fuel Moisture Content ground truth

RECOMMENDATION

Currently, only 111 ground-truth LFMC measurements are available across two years (2015 and 2016) and in three locations - all near Canberra. Ground measurements of LFMC are required close in time to satellite passes and sampling of a range of vegetation types and terrain.

DATA NEEDED

Burn probability: High resolution wind and climate data

RECOMMENDATION

The fuel and terrain conditions can be mapped and modeled, but the wind and climate conditions are known at a much lower resolution. Various wind and climate sensors and ground stations do not provide sufficient resolution to guide the models of the burn probability.

DATA NEEDED

Extreme Fire Behaviour: Higher resolution and nighttime sensors

RECOMMENDATION

The most significant bands of Himawari-8 for analysing smoke plumes, and detecting the initiation of pyrocumulonimbus (pyroCB) extreme fire events, are the visible ones and near-IR. These bands require daylight to detect the smoke plumes and are late to detect any pyroCB events that are initiated at night. In addition, the near-IR bands are measured at only 2 km resolution.

THE DATA QUEST RESULTS

The **2020 Bushfire Data Quest** was created to accelerate interdisciplinary research that leveraged machine learning, remote sensing and bushfire science. Outcomes at TRL 2 - 3 were produced over a short period of time by creating small teams with a carefully balanced mix of expertise, and putting them to work in a 'pressure cooker' sprint environment. Teams were supported by a faculty of top-level experts.

Summaries of the results are presented in the following pages and the teams have written detailed technical memos, that are available in the Appendix.

How can we use historic fuel and moisture data to predict burn rate and intensity?

How can we use high r
and data fusion to dete
faster?

MAPPING FUEL MOISTURE CONTENT

FIRE RISK MAPS AND FIRE PROGRESSION

EARLY DETECTION

ETECTION

UEST 2020

gh resolution models
detect fires earlier/

FIRE BEHAVIOUR

DATA QUEST 2020

How can we use historical satellite imagery to improve predictive models of the behavior of wildfires and, in turn, better inform fire risk management and response?

N OF FIRE IGNITION

SIGNATURES OF EXTREME FIRE BEHAVIOUR

MAPPING FUEL MOISTURE CONTENT

RESEARCHERS


Maoying Qiao


Yang Chen


Vlad Tudor


Caitlin Adams

NEED:

High resolution and frequently updated Live Fuel Moisture Content map (dominating factor in fire risk) to feed risk maps and fire behaviour models.

CHALLENGE:

Make an accurate LFMC map on human scales (20m).

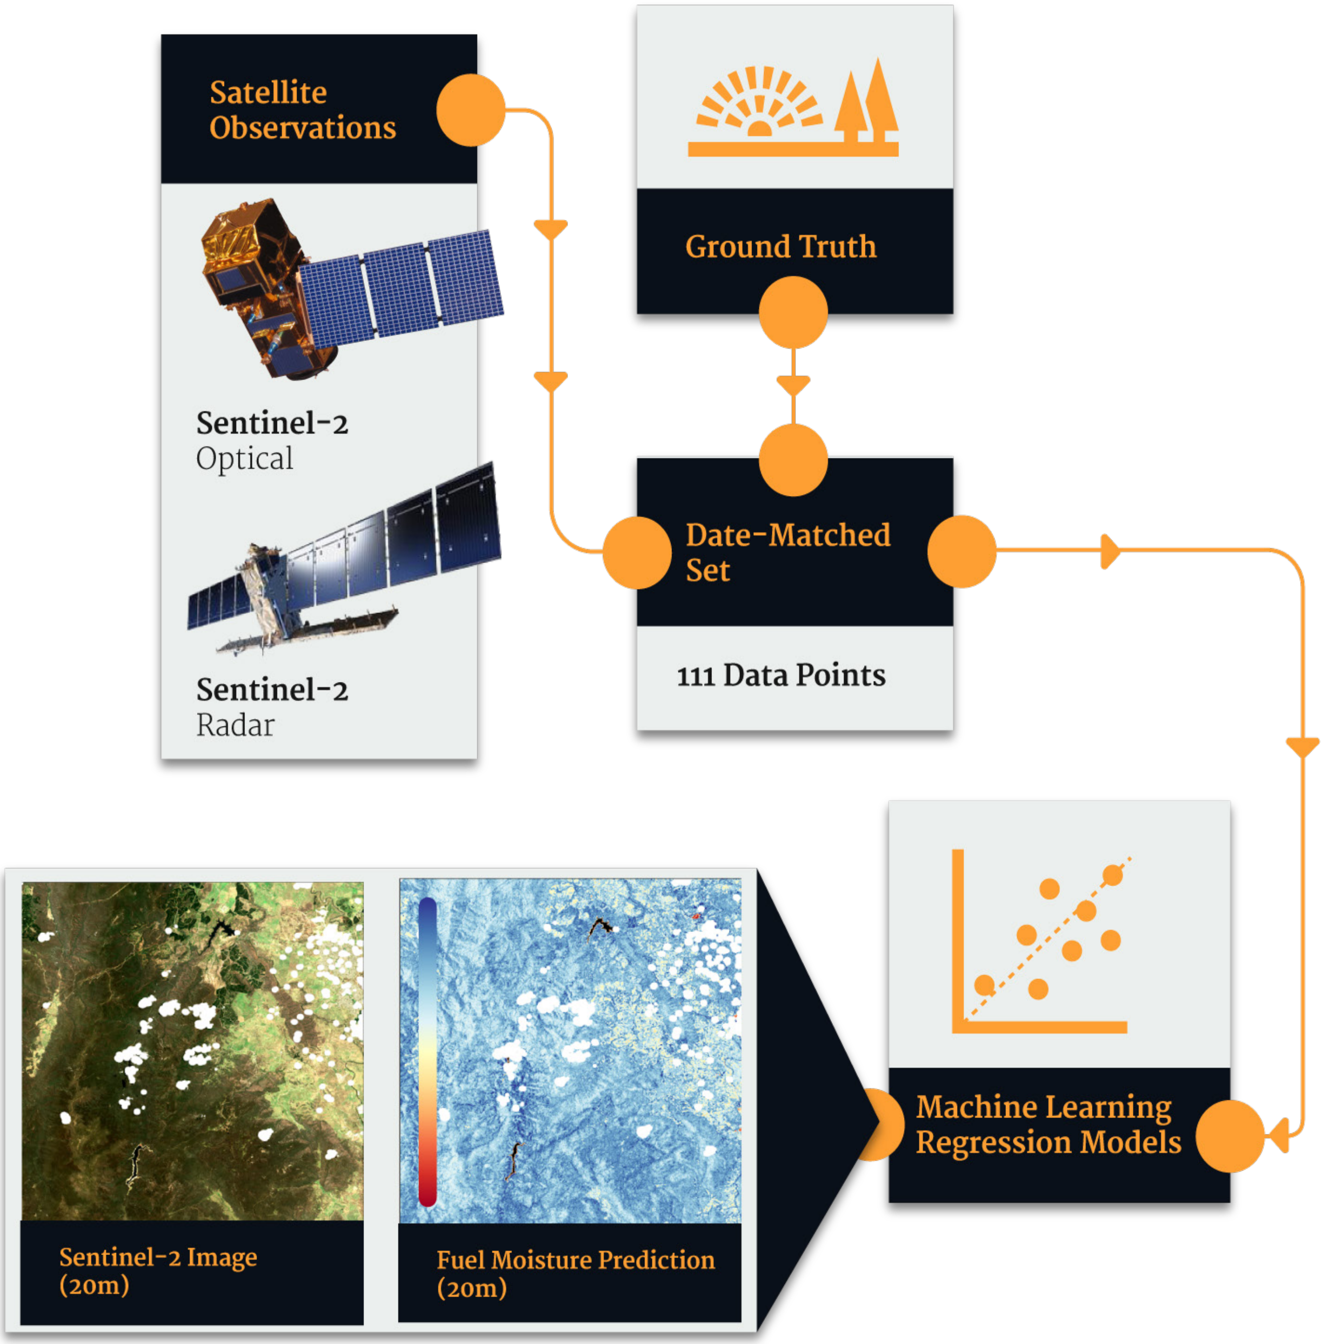

METHOD:

Data fusion of optical and SAR data from ESA Sentinel satellites, calibrated with ground measurements. Used suite of regression models.

RESULT:

Proof of concept pipeline to create a LFMC map.

NEXT STEPS:

Need many more ground measurements that sample range of moisture, vegetation type and landscapes.

EARLY DETECTION OF FIRE IGNITION

RESEARCHERS


Ilze Pretorius


Jack White


Kate Melnik


Alex Codoreanu

NEED:

Detect new fires earlier, with few false-positves and localise fire better.

CHALLENGE:

Make a better fire detection pipeline using satellite data.

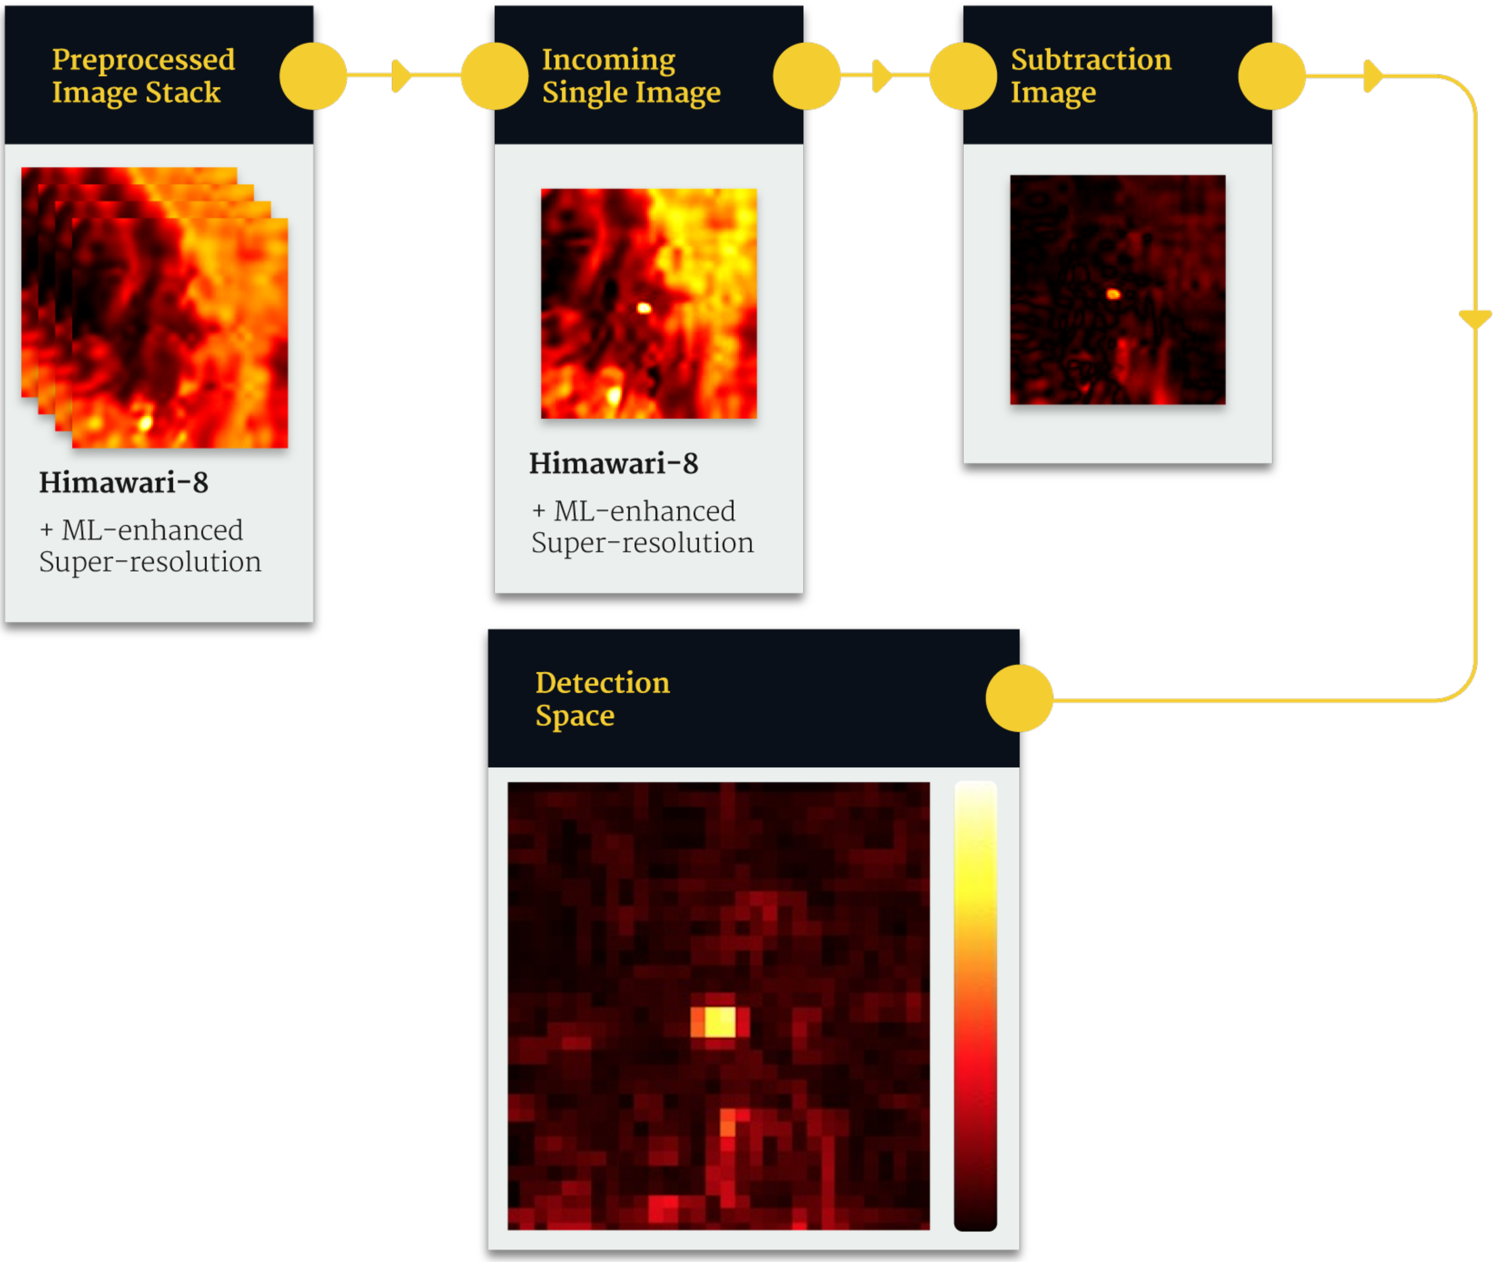

METHOD:

Apply astronomical image stacking and ML **super-resolution** algorithms to infrared images from the geostationary **Himawari-8** satellite.

RESULT:

Prototype pipeline for detecting fires within 10-20 minutes

NEXT STEPS:

Pipeline needs to incorporate reliable cloud masking and scale to cover larger areas.

FIRE RISK MAPS AND FIRE PROGRESSION

RESEARCHERS


Mahdi Kazemi


Ehsan Abbasnejad


Martyn Eliott


Yuri Shendryk


Sam Van Holsbeeck

NEED:

Accurate fire-risk maps and just-in-time prediction of fire spread.

CHALLENGE:

Create a map of fire risk from recent EO data and predict fire spread given ignition points.

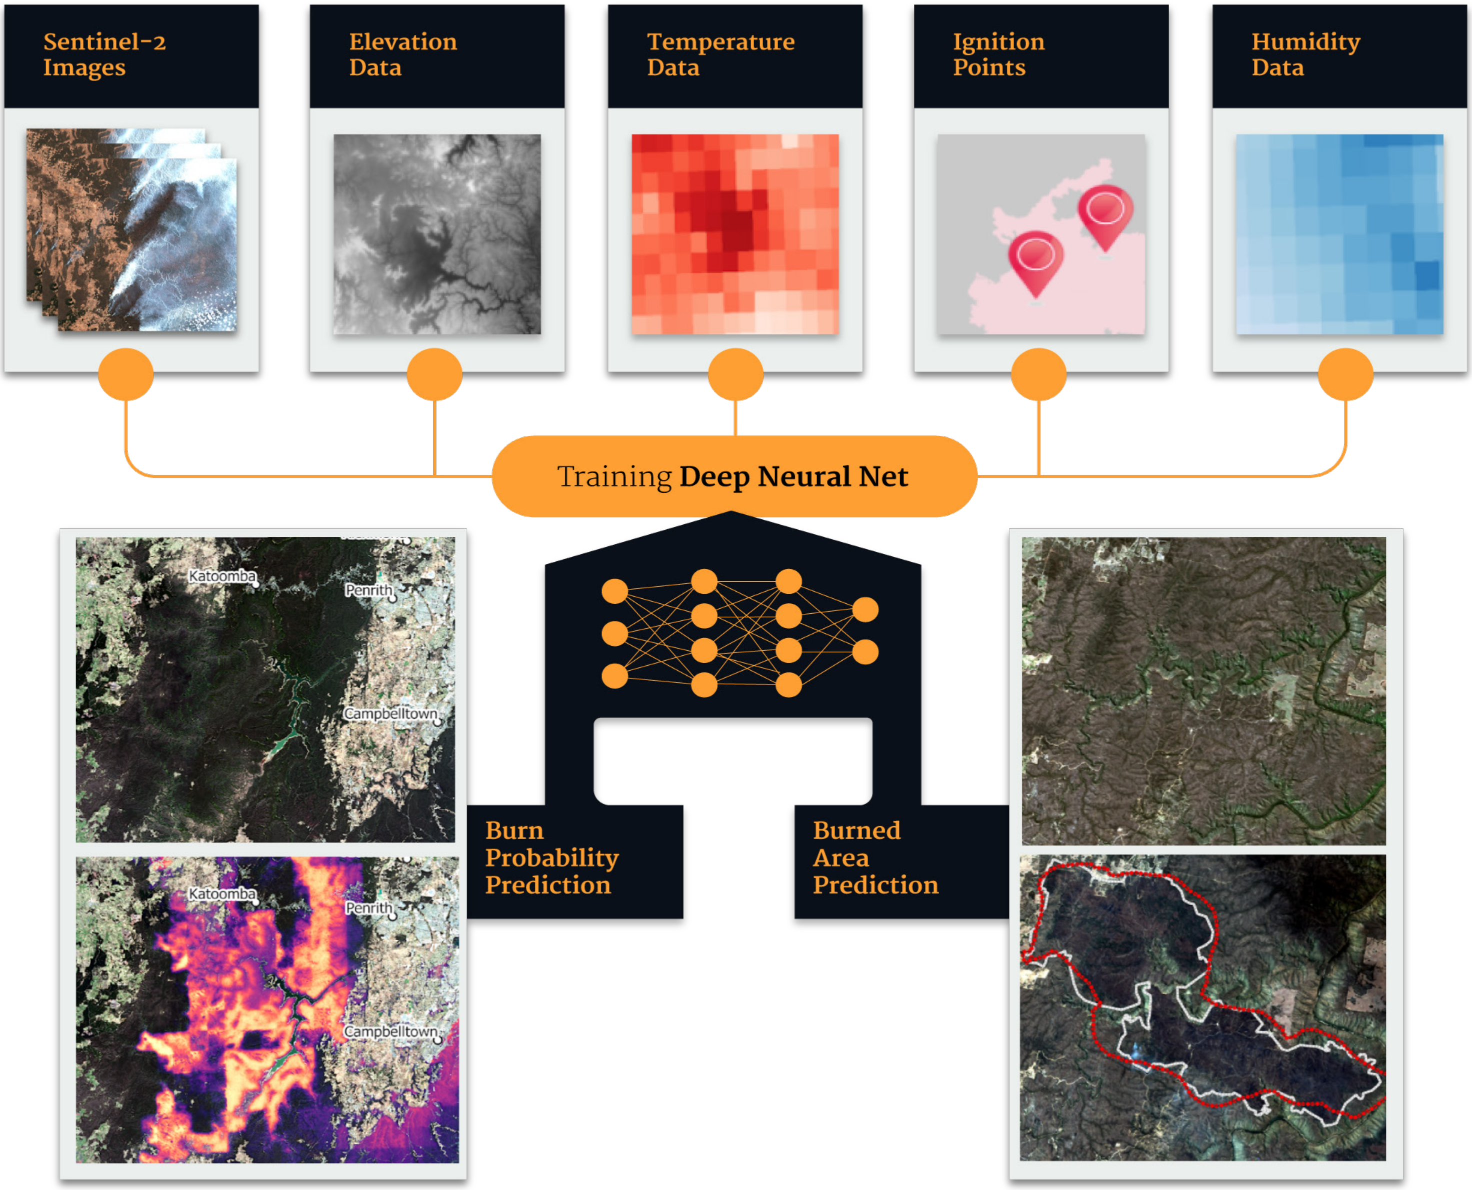

METHOD:
Neural networks that fuse Sentinel-2 images with terrain and weather information.

RESULT:
Two segmentation models that predict fire risk, fire scar boundary and temporal progression of fire.

NEXT STEPS:
Add high resolution grid of weather data and scale to larger areas.

SIGNATURES OF EXTREME FIRE

RESEARCHERS

Todd Ellis

Thomas McCavana

Stephane Mangeon

Anna Matala

NEED:

Early warning of when a fire will exhibit extreme behaviour, like generating a pyrocumulonumbus (pyro-Cb) cloud.

CHALLENGE:

Investigate if remote-sensing data can be used to detect extreme fire behaviour earlier.

True Color

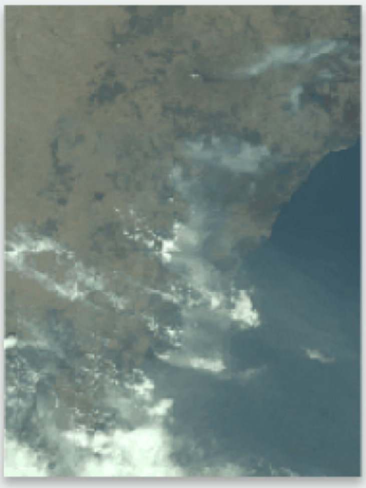

Plume Mask

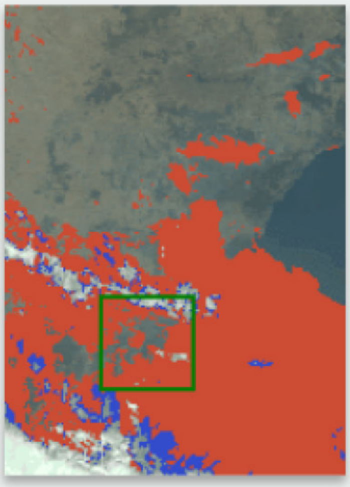

Band 7

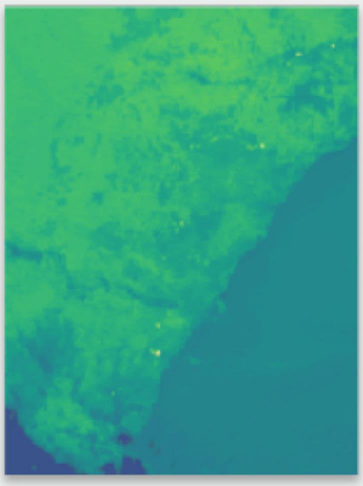

PCA Scatterplot (Masked Pixels

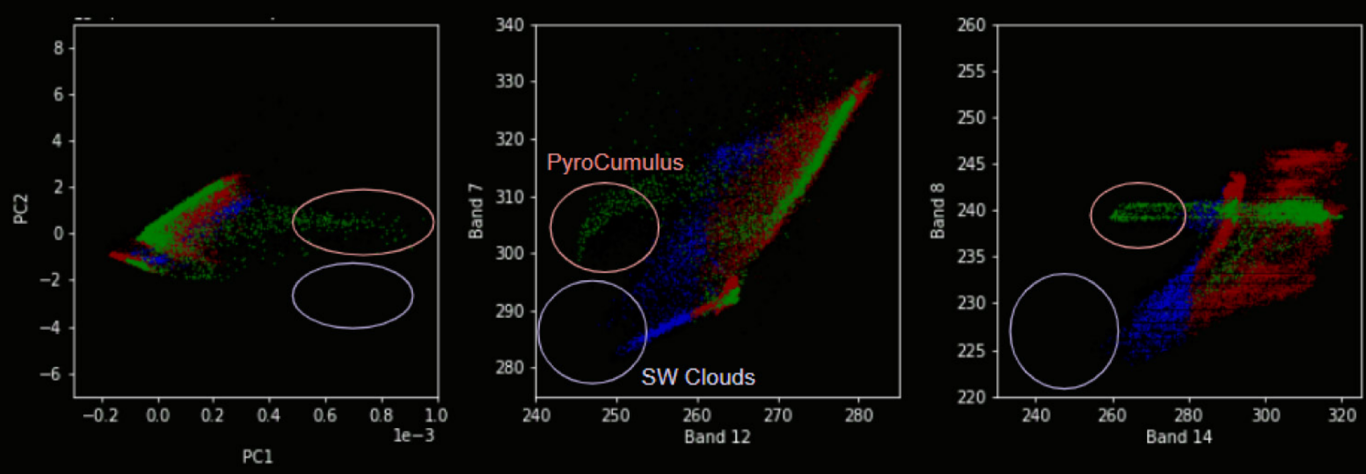

METHOD:

Use dimensionality reduction techniques (principal component analysis, PCA) and clustering to detect and isolate signals in multi-spectral data from the Himawari-8 satellite

RESULT:

Candidate signals from extreme fires and pyro-Cb events were detected in PCA and colour ratio plots.

NEXT STEPS:

Expand the analysis to include higher resolution data and analyse spatial patterns.

DATA QUEST FACULTY

SCIENCE FACULTY

DR TARA STRAND
SCION CROWN RESEARCH INSTITUTE, NZ

DR MARTA YEBRA
AUSTRALIAN NATIONAL UNIVERSITY

A/PROF DALE HAMILTON
NORTHWEST NAZARENE UNIVERSITY

DR RUTH LUSCOMBE
FIREBALL

BRAD CARTER
FIREBALL

DR MARTIN COPE
CSIRO

ZHENG-SHU ZHOU
CSIRO

LEON MAJEWSKI
BOM

SIMON OLIVER
GEOSCIENCE AUSTRALIA

ML FACULTY

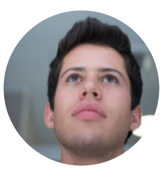

PROF YARIN GAL
AITC CO-CHAIR
UNIVERSITY OF OXFORD

DR CHEDY RAISSI
AITC CO-CHAIR
UBISOFT

DATA ANALYSIS SUPPORT

BREALYN BOERNER
NNU STUDENT

JEFFREY FAIRBANKS
NNU STUDENT

KAMDEN BROTHERS
NNU STUDENT

KYLE DUNCAN
NNU STUDENT

DATA QUEST FACILITATION & SUPPORT

CORMAC PURCELL
DATA QUEST LEAD,
TRILLIUM TECHNOLOGIES

SARAH MCGEEHAN
DATA QUEST CO-LEAD,
TRILLIUM TECHNOLOGIES

JAMES PARR
FDL FOUNDER, CEO
TRILLIUM TECHNOLOGIES

EMELINE PAAT-DAHLSTROM
DATA QUEST TEAM

ERIC DAHLSTROM
DATA QUEST TEAM

LEO SILVERBERG
DESIGNER,
TRILLIUM TECHNOLOGIES

RICHARD STRANGE
DATA QUEST DATA
WRANGLER

JODIE HUGHES
DATA QUEST PRODUCER
TRILLIUM TECHNOLOGIES

RUSSELL BOYCE
UNSW SPACE &
DATA QUEST STEERING
GROUP

MARK CHEUNG
LOCKHEED MARTIN &
DATA QUEST STEERING
GROUP

SUDANTHA BALAGE
UNSW SPACE &
DATA QUEST STEERING
GROUP

LEE SPITLER
MACQUARIE UNIVERSITY &
DATA QUEST STEERING
GROUP

JENN ZHU
CSIRO & DATA QUEST
STEERING GROUP

ADITYA CHOPRA
ANU & DATA QUEST
STEERING GROUP

DR ZHITAO XIONG
NSSN & DATA QUEST
SUPPORT

TECH MEMOS

FUEL ASSESSMENT

DATA QUEST 2020

MAPPING FUEL MOISTURE CONTENT

BUSHFIRE DATA QUEST 2020

Technical Memorandum

High Resolution Prediction of Live Fuel Moisture Content

Caitlin Adams (FrontierSI, Melbourne)

Yang Chen (DATA61, CSIRO, Melbourne)

Maoying Qiao (DATA61, CSIRO, Sydney)

Vlad Tudor (Code for Australia, Sydney)

Science and Data Leads

Dr. Marta Yebra (Australian National University)

Abstract / Executive Summary

In a country as dry as Australia, it is well understood that the absence of moisture in the landscape can lead to devastating bushfires. Reliable and frequent measurements of how wet or dry our vegetation is can aid fire preparation and management, particularly in helping plan controlled burns. The amount of moisture in vegetation across the country is currently estimated by the Australian Flammability Monitoring System, with one moisture value every 500 m^2 (an area close to the size of four cricket fields). We set out to see if we could reduce this to a moisture value every 20 m^2, by applying machine learning to data from the Copernicus Sentinel Satellites. While we were able to produce maps of vegetation moisture at this resolution, the accuracy and applicability of our method were hampered by the lack of ground measurements to inform our model. There is promise in this work, and we recommend the collection of more ground-based measurements of vegetation moisture, both by experts and citizen science programs, which would improve the accuracy and applicability of our approach.

Introduction

The moisture content of vegetation is an important factor in understanding fire behaviour, both in the context of managing prescribed burns and predicting how wildfires might spread. Fires interact with vegetation in complex ways, often igniting in dead vegetation on the ground but spreading to the living canopy. Thanks to remote sensing (RS) techniques leveraging satellite data, we are able to monitor live fuel moisture content (LFMC) from space on a regular basis. The Australian Flammability Monitoring System (AFMS, http://wenfo.org/afms/) currently provides timely spatial information on LFMC across the continent. However, this map is only available at 500m resolution, which prevents detailed insight into potential fire behaviour. Features in the terrain as small as 50m can influence the rate of spread and intensity of a fire, meaning that a ten-fold increase in resolution would have a significant impact on predictions. Our team has applied machine learning to 20m spatial resolution satellite data from the Copernicus Sentinel satellites[1] to provide high resolution estimates of LFMC, aiding bushfire management.

Identified Needs and Opportunities

We identify the following needs that may be addressed using machine learning techniques to integrate passive and active RS images:

1) Create a finer-resolution map of LFMC in Australia at shorter cadence using Sentinel satellite observations.
2) Obtain more accurate estimates of LFMC by integrating optical with synthetic aperture radar data.

The algorithm to create the current AFMS product is described by Yebra et al. 2018 and uses a computationally expensive radiative transfer model to derive the LFMC map from MODIS[2] reflectance data. The radiative transfer models are applied using a 16-day composite of MODIS data (MCD43A4 Version 6) at 500 m spatial resolution. Such a coarse resolution is acceptable for setting the national fire danger rating, but is too coarse to model fire behavior at smaller scales. A finer-resolution LFMC estimate can be implemented using Sentinel 2 observations at 10-20m spatial resolution every 3 to 5 days. Sentinel-1 C-band (5.4 GHz) synthetic aperture radar (SAR)

[1] The Sentinel satellites are part of the European Space Agency (ESA) Copernicus network of Earth observation satellites. https://www.esa.int/Applications/Observing_the_Earth/Copernicus
[2] Moderate Resolution Imaging Spectroradiometer on the Terra satellite. https://terra.nasa.gov/about/terra-instruments/modis

data has the potential to make the LFMC measurement more reliable because its longer wavelength can penetrate through canopy to detect understory vegetation and it is also less sensitive to weather conditions. Here we use both Sentinel-1 and -2 composite data through a computer vision approach to estimate LFMC for a test region in the Australian Capital Territory. We also make recommendations on increasing the number and diversity of ground-truth FMC measurements for various fuel types (e.g., grass fuel, shrubs, and forest overstorey and understorey fuel). Enriching the quality and quantity of the ground-truth data will improve the prediction capacity of the machine learning models in a more heterogeneous environment.

Data description

Our project used three datasets: Sentinel-2 optical imagery, Sentinel-1 radar imagery, and Globe-LFMC on-ground FMC measurements. The data from the two Sentinel missions was used to generate the feature variables for our approach, and the Globe-LFMC data was used as the target variable.

Sentinel-2

The Sentinel-2 mission is part of the European Space Agency's Copernicus programme, and consists of two polar-orbiting satellites that measure the Earth's reflectance over 13 bands in the visible, near-infrared, and short-wave infrared portions of the electromagnetic spectrum. The pair of satellites have a combined revisit time for Australia of around 3-5 days. The first satellite, Sentinel-2A, was launched on the 23rd of June, 2015, and the second satellite, Sentinel-2B, was launched on the 7th of March, 2017.

There are a number of processing steps involved in making data from Sentinel-2 analysis-ready. Importantly, the satellites measure light that has passed through the atmosphere, so the effects of this passage must be removed before useful insight can be gained about the Earth's surface. This data product, known as bottom-of-atmosphere reflectance (sometimes referred to as Level-2 data), can be directly used in analysis, and is the recommended product. There are a number of steps involved in producing analysis ready data, which can lead to minor differences in the products produced by different agencies. To this point, Level-2 data is available for Australia from Geoscience Australia (through the National Computational Infrastructure, Amazon Web Services, and the Digital Earth Australia Sandbox) and the European Space agency (through the Copernicus Hub).

Service	Access	Description
National Computational Infrastructure (NCI)	http://dap.nci.org.au/thredds/remoteCatalogService?catalog=http://dapds00.nci.org.au/thredds/catalog/if87/catalog.xml	Direct access to Geoscience Australia's Sentinel-2 ARD stored on the NCI through a THREDDS catalogue.
Amazon Web Services (AWS)	http://dea-public-data.s3-website-ap-southeast-2.amazonaws.com/?prefix=L2/	Direct access to Geoscience Australia's Sentinel-2 ARD data stored on AWS through S3 client
Digital Earth Australia Sandbox	https://app.sandbox.dea.ga.gov.au/	A JupyterHub platform that can be used to access and analyse Geoscience Australia's Sentinel-2 ARD through the Open Data Cube python API

Copernicus Hub	https://copernicus.nci.org.au/sara.client/#/explore	An online search interface that allows a user to query and access the European Space Agency's Sentinel-2 ARD data (stored in Australia on the NCI). It can also be queried with a Python package.

In our work, we chose to access the Sentinel-2 Level-2 data through the Digital Earth Australia Sandbox, an implementation of the Open Data Cube housing Geoscience Australia's Earth observation data. The Open Data Cube has a Python API that allows for the extraction and processing of raster data stored in its database. We chose to use the Sandbox as it provided easy access to processed Sentinel-2 data for our area and time of interest. We discuss our use of the Sandbox to acquire Sentinel-2 data in more detail below.

Sentinel-2 captures 13 spectral bands, ranging from visible to short-wave infrared. We used the blue, green, red, vegetation red edge 1, near-infrared and shortwave infrared bands (bands 2, 3, 4, 5, 8 and 11 respectively) as features for our models. From these, we calculated additional band index features, which can indicate physical properties. We calculated the normalised difference vegetation index (NDVI), the visible atmospherically resistant index (VARI) and the modified normalised difference water index (MNDWI). The first two are designed to highlight vegetation in the landscape, the final highlights water. All three are valuable in attempting to predict fuel moisture, making them informative features.

Sentinel-1

Sentinel-1 is another mission in the Copernicus programme, again consisting of two polar-orbiting satellites. Unlike Sentinel-2, the Sentinel-1 satellites transmit microwaves, then record how much of the outgoing signal was reflected back from the surface. Due to operating at longer wavelengths than visible light, these satellites pick up different features from the ground and can see through clouds. Because they actively send radiation, they can also operate at night. Sentinel-1 operates in the C-Band (central frequency of 5.405 GHz) and can transmit either vertically- or horizontally-polarised light; it then measures the phase and intensity of the reflected light in both polarisation bands. Between the two satellites, it has a revisit time of around 6 days.

Sentinel-1 analysis ready data for Australia is not as readily available as Sentinel-2. Level-1 data from Sentinel-1 is provided through the Copernicus Hub. There are two products: Single Look Complex (SLC) and Ground Range Detected (GRD). SLC preserves phase information, whereas GRD only keeps intensity information. Phase information is typically used when measuring changes in height (through interferometry) and is not needed for our application. Sentinel-1 also has four observing modes, corresponding to different ground-coverage and resolution; we used data collected in the Interferometric Wide swath mode, which is the most commonly used configuration.

We used Sentinel-1 GRD data that had been processed to a Level-2 backscatter product specifically for the Data Quest by Zheng-Shu Zhou from CSIRO. This involves multiple steps, including removing noise sources and applying calibration. The resulting output is a digital number (DN; unitless), which can be converted to the typical radar backscatter parameter (γ_0) through the equation $\gamma_0 = 10 \times \log_{10}(DN)$. This is done for both polarisation bands, where Sentinel-1 transmits vertically-polarised light, and receives either vertically-polarised light (VV band), or horizontally-polarised light (VH band). We use both of these as features in our model.

Globe-LFMC

Live fuel moisture content (LFMC) measures the amount of water contained in live vegetation as a percentage of the vegetation's dry mass. The Globe-LFMC database was established to provide a collection of the measurements from around the world, which could then be used to better calibrate algorithms attempting to predict LFMC from remote sensing measurements, such as in this work.

The database is stored as a csv file and can be downloaded, and opened, with Excel or Python. It includes 161,717 LFMC measurements from 1,383 sampling sites in 11 countries. The most useful columns for this analysis are *Country*, *Latitude*, *Longitude*, *Sampling date*, *Sampling year*, *Land Cover* and *LFMC value*. After loading, the data can be filtered to show measurements for Australia only.

Data Gaps

The single biggest factor that restricts the generalising power of the model is the availability of ground-truth LFMC measurements. Currently, only 111 LFMC measurements are available post-2015 (the launch date of Sentinel-2), spread across two years (2015 and 2016) and three locations - all near Canberra.

In addition, the availability of Sentinel-2 data was frequently compromised due to cloud presence. As such, long time gaps (> 20 days) between ground measurements and Sentinel-2 observations were common. Fuel moisture can vary on a timescale of hours, especially during hot weather or after rainfall, meaning such long observation gaps pose a severe threat to the capability of the model to learn the associations between remote sensing and ground truth data.

If many more ground measurements of LFMC were available, spread over land and time, two problems would be solved: 1) the lack of ground truth observations close in time to satellite passes and 2) the poor sampling of a range of vegetation types and terrain, which leads to a model that is not *generalizable*.

We note that since the Globe-LFMC measurements ceased, the Sentinel-2B satellite has been launched, doubling the number of passes over Australia and providing more opportunities to match data.

Data Product	Link
Sentinel-2 Optical and NIR Imagery (Level-2, bottom of atmosphere corrected)	https://docs.dea.ga.gov.au/setup/sandbox.html
Sentinel-1 SAR (Level-2, Analysis Ready Data)	Custom-processed Level-2 data product [2]
Globe-LFMC	https://www.nature.com/articles/s41597-019-0164-9

[2] Level-2 Sentinel-1 SAR data was processed by CSIRO especially for the Data Quest. In the future it will be available from CloudStor.

Data Exploration

The nature of our dataset means that it's valuable to understand how our observed variables relate to the fuel moisture content. In particular, examining how the different feature variables correlate with fuel moisture content can give us insight into which features we select for our machine learning models. Variables with higher correlations should have good predictive power.

Here we create scatter plots between each of the feature variables from Sentinel-2 and LFMC to quickly assess any existing correlations. We observe some correlation between our variables and the observed fuel moisture content. For example, higher normalised difference vegetation index (NDVI) values correspond to higher fuel moisture; this is consistent with higher NDVI values corresponding to healthy vegetation, and lower values corresponding to dry or dead vegetation. The scatter plots also reveal that there is significant scatter in the data, which is to be expected given our small sample.

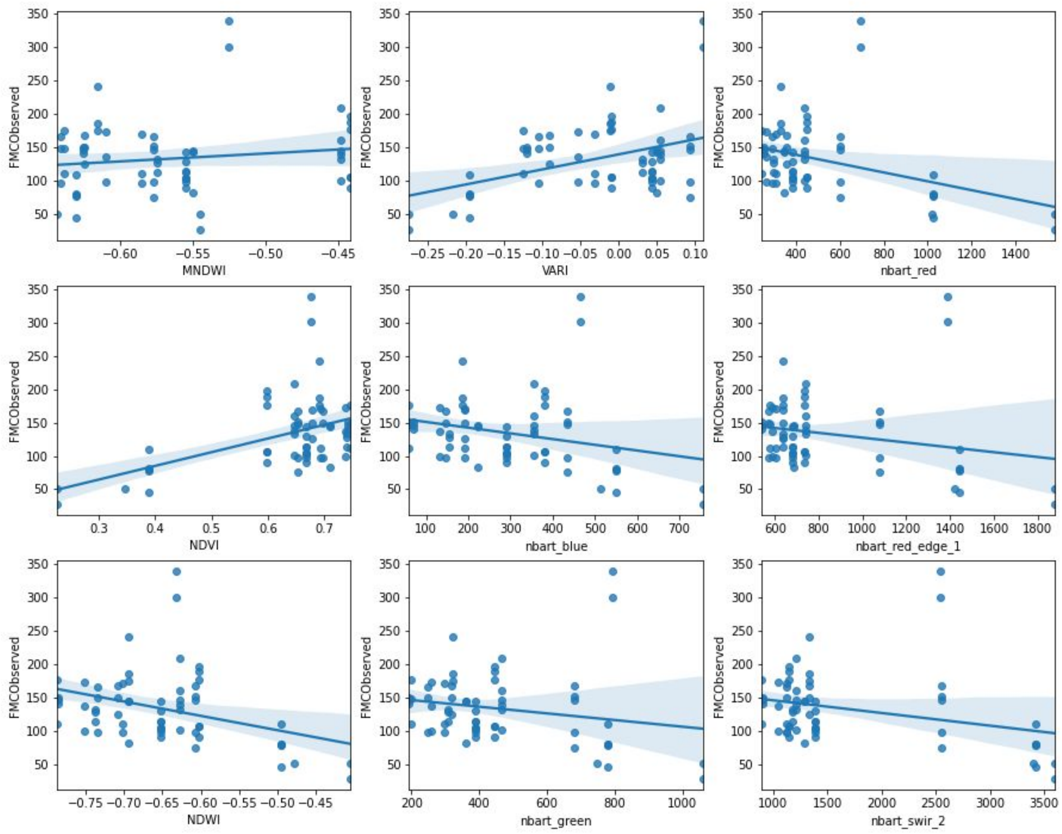

Figure 1: Plots of feature variables from Sentinel-2 optical data against the ground-based measurements of fuel moisture content (FMCObserved). The feature variables are: *modified normalised water index* (MNDWI), *visual atmospheric resistance index* (VARI), *optical red band* (nbart_red), *normalised difference vegetation index* (NDVI), *optical blue band* (nbart_blue), *red-edge band* (nbart_red_edge_1), *normalised difference water index* (NDWI), *optical green band* (nbart_green), *shortwave infrared* (nbart_swir_2).

Methodology

Workflow overview

Our method involves individually processing each of our datasets then preparing them for ingestion into our regression modelling pipeline. For each Sentinel dataset, we identify data that falls within our areas of interest (the three sampling sites around Canberra) for the time period overlapping with available LFMC measurements (2015 to 2016). This is then matched against the available LFMC measurements to build a dataset that lists the Sentinel-1 and -2 observations closest in time and space to each LFMC measurement. This is formatted as a table containing the satellite observations (feature variables) and the matched LFMC measurements (target variable).

We then fit a number of models to the data and compare outcomes. We evaluate the models based on their root-mean-square error (RMSE) when fitted to testing data set aside from our main dataset.

Sentinel-2 Pipeline

We accessed Sentinel-2 data through the Digital Earth Australia Sandbox, a Python-based platform for Earth observation analysis. The DEA Sandbox was chosen as the most user-friendly and convenient method to access analysis-ready Sentinel-2 data for Australia. In particular, the sandbox provides a full environment to process the Sentinel data, with all modules installed and ready to use. The LFMC data used in our analysis was collected from three sites, specified by latitude and longitude. We created shapefiles defining 100m by 100m areas around each site, matching the approximate collection area of the ground-truth measurements and providing a reasonable sample of pixels around the area of interest from Sentinel-2.

We then loaded Sentinel-2 data between August 2015 and November 2016 (the dates covered by the LFMC samples). The Sandbox allowed us to retrieve specific Sentinel-2 bands, so we selected the blue, green, red, vegetation red edge 1, near-infrared and shortwave infrared bands (bands 2, 3, 4, 5, 8 and 11 respectively). We also required that the query only return observations where more than 80% of our area of interest was cloud free. This was to ensure we got a representative measurement of the area of interest.

After loading the data, we also calculated three band indices as additional features. These were the normalised difference vegetation index (NDVI), the Visible Atmospherically Resistant Index (VARI) and the modified normalised difference water index (MNDWI). The first two are designed to highlight vegetation in the landscape, the final highlights water. All three are valuable in attempting to predict fuel moisture.

Since the LFMC data is a single observation for each site and sampling date, we reduced the loaded Sentinel-2 data and calculated band indices to representative values by taking the spatial median across each loaded observation. This resulted in a single measurement of the loaded bands and indices for each observation, which we could use as features for our machine learning model. These features were saved to a csv, along with the date of observation, which is later used to match these observations to the LFMC values.

Sentinel-1 Pipeline

The Sentinel-1 data we used was provided in the form of GeoTIFF files for each observation, with one file for the VH band and one for the VV band. We loaded these using the GDAL package, and extracted the pixel corresponding to our sampling site locations by calculating the location of the site relative to the image's coordinate bounds. The file names also contain the date and time-stamp of each observation, so we used text processing code to extract these.

Sentinel-1 data is provided as a unitless digital number for each band. We converted it to the more commonly used radar backscatter coefficient, γ_0, by taking the base 10 logarithm of the digital number and multiplying it by 10 (see Sentinel-1). We then wrote the backscatter values along with the date of observation to a comma-separated-variable (CSV) file to be matched by timestamp with the LFMC values.

Data Preparation Pipeline

To be useful for our machine learning workflow, we need to format the data into a table, with each feature variable from Sentinel-1 and -2, and the LFMC target variable as columns. Rows correspond to matched observations.

To begin, we load the LFMC data and keep only observations from the Australian Capital Territory between January 2015 and December 2016. We then load the processed Sentinel-1 and -2 datasets from the previous, and ensure that there is only one observation per row.

For each LFMC observation, we then identify the Sentinel-1 and -2 observations that are closest in date and position, and record the matched values. We note the date differences between the LFMC and each Sentinel observation so that we can filter the matches to those that were within reasonable time of each other.

Regression Modelling Pipeline

We identified our task as a regression problem and fit the simplest linear regression model as our baseline. Then, we applied L1 and L2 regularizers to handle the overfitting issue. In addition, we applied nonlinear regression models to study our tasks, including random forest regression, K-nearest-neighbour regression, Bayesian ridge regression, multi-layer perceptron, linear support vector regression, and Tweedie regression with Gamma distribution. Table 1 summarises the properties of each regression algorithm, noting any significant advantages or pitfalls.

Table 1: Properties of regression algorithms used to model the LFMC from SAR and optical images.

Algorithm	**Properties**
Linear Regression	Simplest baseline model, assumes linear relationships between features and target variable.
Random Forest	A collection of individual decision trees that randomly use different features to predict the target variable. As an ensemble method, the decision trees "vote" on the predicted target variable, making the prediction more robust to errors than from an individual decision tree. The disadvantage is that these models can overfit the training data, and not generalise well to new examples.

KNN Regression	This method predicts the target value of new points by locally interpolating the values of the point's nearest neighbours. An advantage of KNN is the predicted values will be in the range of the target values from the training dataset. Thus, it may avoid predicting meaningless values. This is a simple algorithm, and easy to interpret, but may suffer when working with small training sets.
Bayesian Ridge Regression	A probabilistic model for regression problems. This can have a number of advantages, particularly around incorporating and predicting uncertainty. It can also work well for small datasets. This method can be time-consuming to run if the dataset is larger, and may not perform better than standard linear regression.
Multi-layer Perceptron	A simple neural network, with fully connected layers. Multi-layer perceptron models often perform well on complex datasets, but may take longer to train depending on the complexity of their architecture.
Linear SVM Regression	This model attempts to fit a line and decision boundary that captures the training data. As such, it may perform poorly on data with a large amount of scatter.
Tweedie Regression with Gamma Distribution	This model returns non-negative predictions, which is valuable here as our target variable must be non-negative given the definition of fuel moisture content.

We applied a machine learning pipeline to compare different models fairly. We measured the average performance of each model via 5-fold cross validation. Within each fold's experiment, we applied grid search methods to pick the best hyperparameters of a model in terms of model fitness.

Results

We present the results of fitting the seven algorithms to the data in Table 2 below. Simple linear regression models resulted in over-fitting, which we attempted to control using L1 and L2 regularizers, leading to nine models for comparison.

The matching window in time between the Sentinel satellites passing overhead and the date of the ground-truth measurements is critical for this analysis. Ideally this should be as short as possible so that the ground-truth measurements and satellite images are measuring the same moisture levels. However, the Sentinel-1 SAR observations for this time were relatively sparse meaning that a matching window of 26 days was necessary to provide a sufficiently large number of points to. We present the R-squared and RMSE values for time-filtered data and all data as a comparison.

Comparison of fitting all data versus time-mathed data

If we don't restrict the time-matching window between S1 and S2 data, there are 111 data points available. We also choose a matching window (through trial-and-error) of less than 26 days, leaving 74 data points available for fitting.

Table 2: Performance comparison when using all data points and data points with maximum 26

days between S1 and S2. For our chosen metrics, higher R^2 and lower RMSE indicate better performance (indicated by the directional arrows after each metric name). The metrics for the best performing models are highlighted in bold text and underlined for similar performing models.

	All data points		Date gap ≤ 26 days (74 data points)	
	R^2 ↑	**RMSE ↓**	**R^2 ↑**	**RMSE ↓**
linear regression	0.06	56.51	0.32	45.25
linear regression with L1	0.06	54.38	0.30	47.47
linear regression with L2	0.07	55.41	<u>0.38</u>	**40.12**
Tweedie Regressor with Gamma distribution	0.03	54.43	0.19	47.80
Bayesian Ridge Regressor	0.06	52.72	0.31	46.87
K-Neighbors Regressor	0.01	56.45	0.32	45.87
MultiLayer Perceptron	0.06	59.61	**0.43**	41.64
LinearSVR	0.07	55.24	<u>0.38</u>	<u>40.76</u>
Random Forest Regressor	0.07	58.12	0.24	47.32
MEDIAN:	***0.06 ± 0.02***	***55.41 ± 2.08***	***0.32 ± 0.07***	***45.87 ± 3.09***

From the results in Table 2 we can see that the RMSE values are high and R-squared values low when fitting using all data points (e.g. $\langle RMSE\rangle_{All}$ = 55.41 versus $\langle RMSE\rangle_{Matched}$ = 45.87 and $\langle R^2\rangle_{All}$ = 0.06 versus $\langle R^2\rangle_{Matched}$ = 0.32). Thus we do not recommend performing regression without date-gap matching. When using the date-gap matching, we found that the multi-layer perceptron model had the highest R^2, and had a reasonable RMSE value, with linear regression with L1 and LinearSVR having similar performance.

Comparison between S1 and S2 data (time-matched only)

We next restrict our analysis to only Sentinel-1 and -2 data that can be matched in a 26 day time window. To determine which of the S1 and S2 data dominate, we fit models to both SAR and optical data, and to each dataset individually and in combination.

There are 12 features in the combination: two from S1 (*VV, VH*) and ten from S2 (*nbart_blue, nbart_green, nbart_red, nbart_nir_1, nbart_red_edge_1, nbart_swir_2, MNDWI,NDVI, NDWI, and VARI*). We present the results in Table 3.

Table 3: Performance comparison when using data points with maximum date gap of 26 days between S1 and S2 data, and different combinations of features from each dataset. For our chosen metrics, higher R^2 and lower RMSE indicate better performance (indicated by the directional arrows after each metric name). The metrics for the best performing models are in bold and underline, as before.

	S1+ S2		**S1**		**S2**	
	R^2 ↑	**rmse ↓**	**R^2 ↑**	**rmse ↓**	**R^2 ↑**	**rmse ↓**
linear regression	0.32	45.25	0.22	51.72	<u>0.37</u>	<u>41.47</u>
linear regression with L1	0.30	47.47	0.07	52.78	0.31	46.54
linear regression with L2	<u>0.38</u>	**40.12**	0.11	51.94	<u>0.37</u>	<u>41.45</u>
Tweedie Regressor with Gamma distribution	0.19	47.80	0.14	52.63	0.30	45.58
Bayesian Ridge Regressor	0.31	46.87	0.17	50.59	0.31	46.87
K-Neighbors Regressor	0.32	45.87	0.26	45.99	0.29	46.76
MultiLayer Perceptron	**0.43**	<u>41.64</u>	0.16	49.30	**0.39**	**40.66**
Linear SVR	<u>0.38</u>	<u>40.76</u>	0.13	49.10	<u>0.38</u>	<u>40.87</u>
Random Forest Regressor	0.24	47.32	**0.36**	**43.36**	0.33	45.42
MEDIAN:	***0.32 ± 0.07***	***45.87 ± 3.09***	***0.16 ± 0.09***	***50.59 ± 3.21***	***0.33 ± 0.04***	***45.52 ± 2.75***

From the results in Table 3, we see that:

1. In general, models fit to the S2 data alone perform similar to the combination of S1 plus S2 data. Most models fit to the S1 data alone perform poorly, which could be due to the higher number of features for S2 (10) compared to S1 (2), as well as the types of physical information these features capture.
2. The best performing model is the multi-layer perceptron, which achieved the highest R-squared and a reasonable RMSE on the combined S1 and S2 data.

3. The Random Forest Regressor outperformed all other algorithms on the S1 data alone, approaching the performance of models fit to the combined dataset. This suggests that the radar images have predictive value, but that care is needed in their use.

Selecting Features via Regularisation

The L1 and L2 regularisers apply penalties to the each of the features to constrain the model and avoid over-fitting. This allows the user to apply a threshold to the regularisation coefficients, dropping the least-important features. Similarly, the decision trees in the Random Forest method can assess the importance of each feature by position in the tree. In Table 4 below we present the results of models run on data with the least important features omitted.

Using features selected by L1:

The features selected by the L1 regulariser are VV, VH (S1 data) and nbart_blue, nbart_green, nbart_red, nbart_nir_1, nbart_red_edge_1, nbart_swir_2, I (S2 data). The dropped features are MNDWI, NDVI, NDWI, and VAR.

Table 4: Performance comparison when using data points with maximum date gap 26 and features selected by L1 regularizer. For our chosen metrics, higher R^2 and lower RMSE indicate better performance (indicated by the directional arrows after each metric name). The metrics for the best performing model are in bold.

	R^2 ↑	RMSE ↓
linear regression	<u>0.38</u>	<u>39.84</u>
linear regression with L1	0.30	46.40
linear regression with L2	**0.39**	**39.92**
Tweedie Regressor with Gamma distribution	0.30	45.57
Bayesian Ridge Regressor	0.31	46.87
K-Neighbors Regressor	0.32	45.87
MultiLayer Perceptron	0.34	43.75
LinearSVR	<u>0.38</u>	<u>40.84</u>
Random Forest Regressor	<u>0.37</u>	<u>41.68</u>

From the results in Table 4, we see that when using the reduced feature set, linear regression with L2 regularisation is the best performing model. However, simple linear regression, LinearSVR and Random Forest display similar performance. This may indicate that linear-based models perform better for a smaller set of more correlated features. Even so, these models do not outperform the multi-layer perceptron when using all 12 features. If processing time became an issue when using the multi-layer perceptron in future, this experiment might support using a simpler linear model with a reduced set of features, providing reasonable performance with a lower processing time.

Qualitative results

It is useful to examine the distribution of LFMC values and how the values are spread over the landscape. In Table 5 we compare the values and extent of LFMC for the Globe-LFMC target variable, the current state-of-the art AFMS and selected models from this work.

Table 5: Histograms and spatial distributions of LFMC values for Globe LFMC, the AFMS and selected models from this work.

	histogram of FMC values	predicted maps
ground-truth (Globe-LFMC)	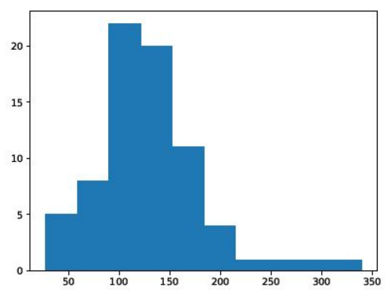	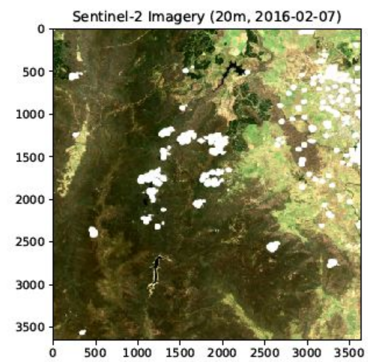
AFMS	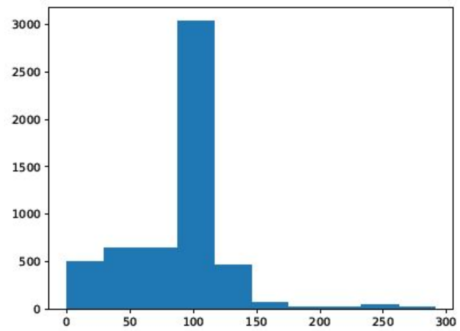	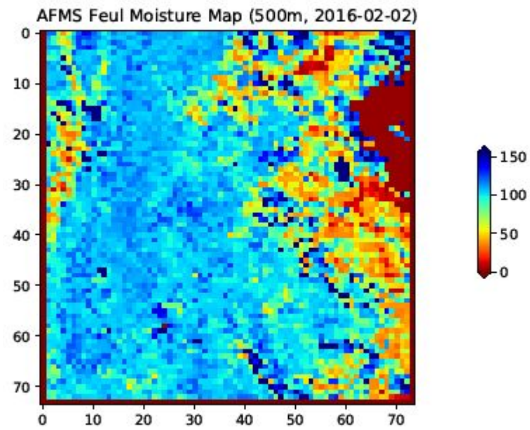

linear regression	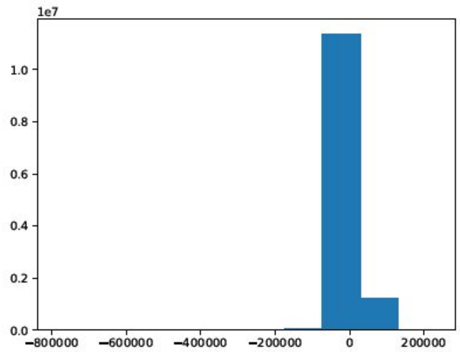	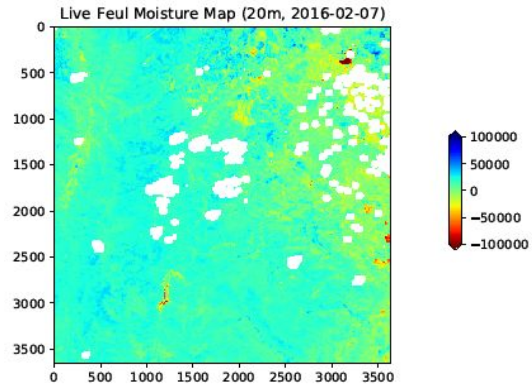
Tweedie Regressor with Gamma distribution	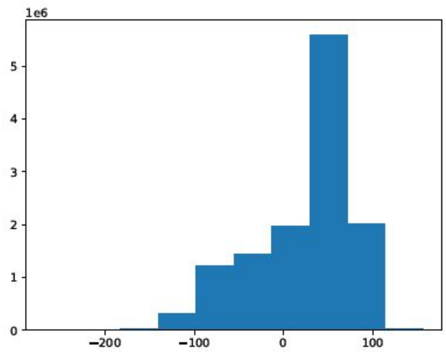	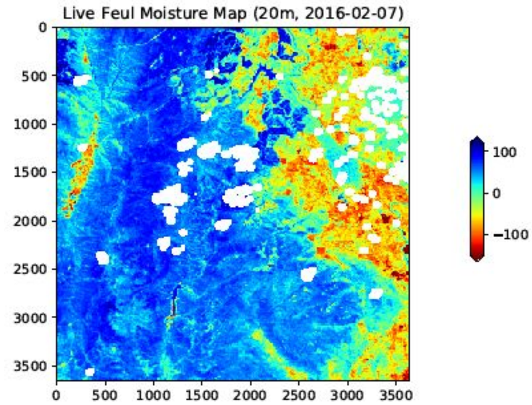
Bayesian Ridge Regressor	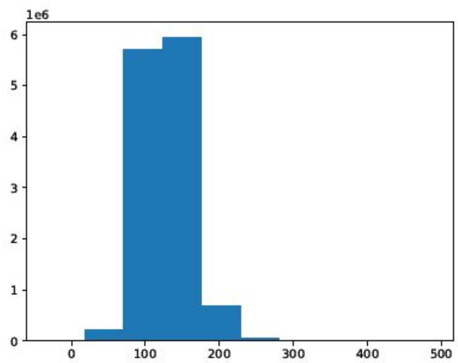	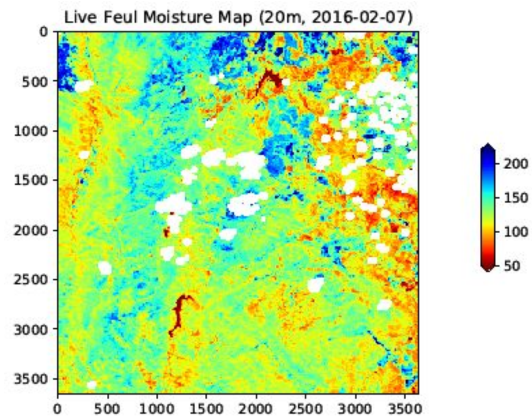

K-Neighbors Regressor	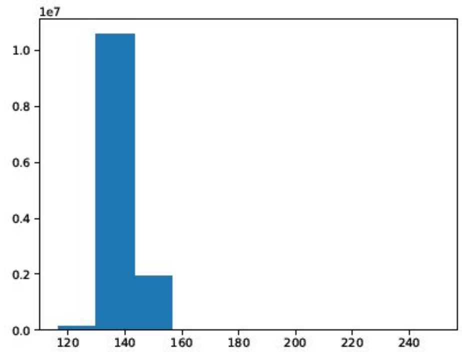	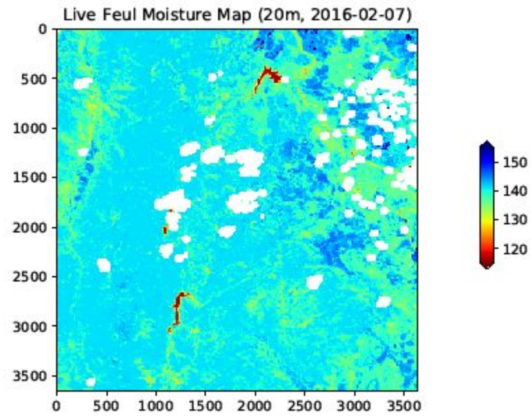
MultiLayer Perceptron	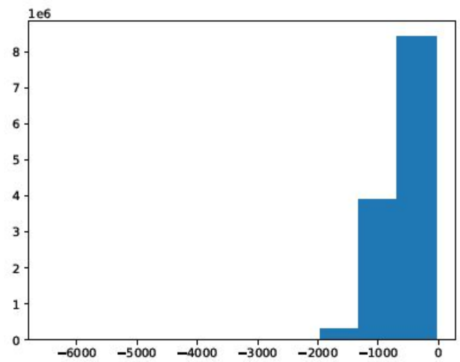	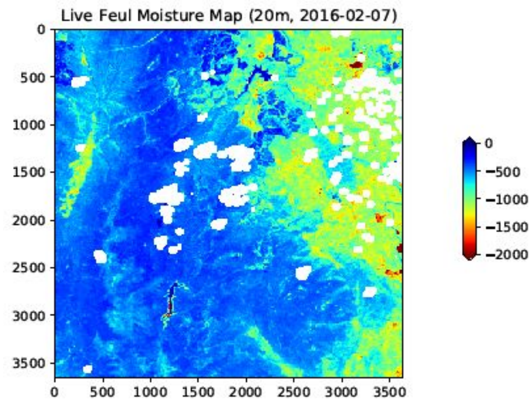
Random Forest Regressor	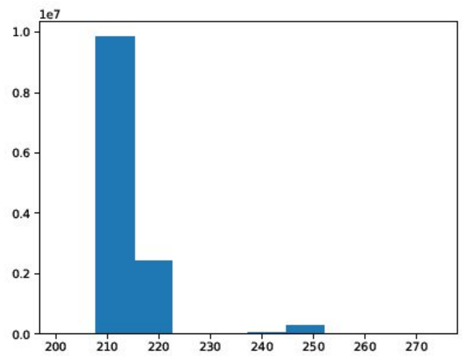	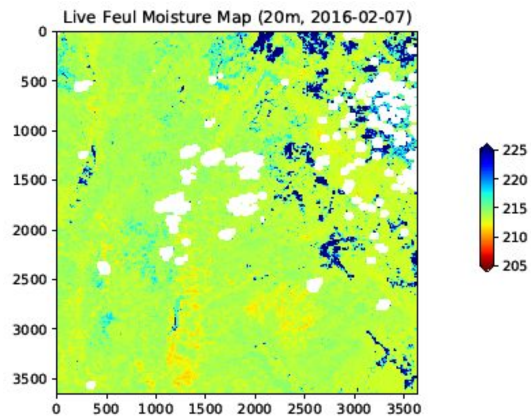

From Table 5, predicted values of the LFMC (histogram column) vary significantly and generally have a greater spread than either the AFMS system or the Globe-FMC data. A few of the models also predict negative LFMC values, which are considered physically impossible, implying that the scale of typical LFMC values has not been well captured. Despite the large variance in values across models, the morphologies of the predicted maps are in reasonable agreement with the AFMS map, meaning that the models are able to detect differences in LFMC corresponding to differences in the landscape, such as dry grasslands and wetter forests.

Discussion

We compare the high-resolution LFMC prediction from our models to the MODIS-derived estimate used by the [Australian Flammability Monitoring System](). As mentioned above, predicted

LFMC values are considerably different between models and the ground truth. Some of the distributions appear reasonable in shape, but have negative values that are physically impossible. We believe that this is largely because of the sparse data available for analysis and the wide temporal gap between most ground-truth measurements and satellite passes. In short, the models are not well constrained by the current datasets.

However, we believe the method is sound and is worth exploring further. In particular, the maps based on the S1 and S2 data show fine structures corresponding to ridges and valleys that appear reasonable. On average, the LFMC values of the two maps are similar, but the variance across the AFMS map is much greater. This difference in variances is partly due to very different LFMC estimates on some vegetation types.

One reason behind such differences is the lack of available data spanning various vegetation types. No ground-truth measurements have been taken for fields, which limits the generalisation power of the model. However, generalisation across vegetation types may not be important when looking at LFMC, when vegetation type itself is a better indicator of flammability than LFMC (compare sclerophyll forest to pasturage).

We saw mixed results about whether including S1 data alongside S2 resulted in better model performance. More work is required to understand the contribution that S1 data is making, and whether it can be enhanced by additional feature engineering, such as that used by Rao et al., 2020.

Next Steps and Recommendations

Before further pursuing higher-resolution LFMC maps, we recommend assessing the importance and uncertainty of each data layer (LFMC vs topography vs vegetation type) to bushfire management. If LFMC resolution (rather than other LFMC uncertainties) is one of the bottlenecks to improving management, the defined resolution challenge can be further developed.

As mentioned in the Data Gaps section, the biggest factor that affects the performance of the model across varying environments and conditions is the availability of LFMC data. If this research project is to be further developed in the future, the availability of ground-truth data is the first item that needs to be addressed. LFMC data should be collected over a wider variety of geographical locations and bioregions in dry to wet conditions. The biggest drawback to collecting more field data is that measuring LFMC is a long, rigorous process (collecting and weighing samples, retrieving to the lab, oven-drying at ~100℃ up to 2 days). Using the same collection method in the future is only likely to increase the availability of data 2 or 3 fold at most. While this will certainly be valuable, it might not lead to desired model performance. An alternative method would be to sacrifice some data quality for quantity. A citizen science project could be set up where people collect their own LFMC measurements. Other advantages of such an approach include the potential of continuous data gathering, including in active fire seasons. Potential LFMC proxies more easily measurable could include: leaf electrical conductance, leaf dielectric permittivity, leaf colour relative to colour templates, leaf mass to size ratio relative to templates and leaf-burning tests.

References

Rao et al., 2020. **SAR-enhanced mapping of live fuel moisture content.** Remote Sensing of Environment, Volume 245, https://www.sciencedirect.com/science/article/pii/S003442572030167X

Yebra et al., 2018. **A fuel moisture content and flammability monitoring methodology for continental Australia based on optical remote sensing.** Remote Sensing of Environment, Volume 212, pp. 260 - 272, https://www.sciencedirect.com/science/article/abs/pii/S0034425718302116

Yebra et al., 2019. **Globe-LFMC, a global plant water status database for vegetation ecophysiology and wildfire applications.** Nature Scientific Data, Article number 155, https://www.nature.com/articles/s41597-019-0164-9

Appendix: Technical Requirements

The processing power required by the machine learning pipeline is very low, owing to the small dataset and model simplicity. As such, it can be quickly run on off-the-shelf machines regardless of operating system.

The training script was built using Python 3.8.5, and used the following libraries:

- numpy==1.19.1
- matplotlib==3.3.0
- pandas==1.0.5
- scikit-learn==0.23.1

The disk space required for storing the S1 and S2 training data is less than 2 GB.

FIRE RISK MAPS AND FIRE PROGRESSION

BUSHFIRE DATA QUEST 2020

Technical Memorandum

Predicting the burned area of bushfires using machine learning

Fuel Assessment Team II

Dr Ehsan Abbasnejad (Australian Institute for Machine Learning, University of Adelaide, Adelaide)
Martyn Eliott (University of the Sunshine Coast, Sunshine Coast)
Mahdi Kazemi (Australian Institute for Machine Learning, University of Adelaide, Adelaide)
Dr Yuri Shendryk (CSIRO Agriculture and Food, Brisbane)
Sam Van Holsbeeck (Forest Research Institute, University of the Sunshine Coast, Sunshine Coast)

Science and Data Leads

Dr Marta Yebra (Fenner School of Environment & Society, Australian National University, Canberra)
Dr Chedy Raïssi (Ubisoft, Singapore)
Prof Yarin Gal (University of Oxford, United Kingdom)

Abstract

This paper presents an efficient machine learning approach to predict the burn probability and burn area of a bushfire as soon as the first ignition points are detected. This task is crucial in early mitigation of fires before they develop into large-scale disasters. In order to address the challenges of this long-term prediction, we propose to define the task as an image segmentation problem. We use two variants of encoder-decoder convolutional neural networks known as UNet and BASNet for this purpose. To increase the accuracy of our prediction in the absence of reliable fire ignition point ground truth, we propose to sample random points from a burnt area. We show that by effectively fusing various data sources such as Sentinel-2 imagery, different weather variables, historic fire data and others, we can achieve a high burn area prediction accuracy of 88% on an unseen validation set and satellite images. Our method is one of the early successful attempts to use image segmentation neural networks for the task of burn area prediction by utilising fire ignition points. This method can be used as a lightweight tool to provide real-time predictions to help address resource management during the bushfire season.

1. Introduction

Wildfires, also known as bushfires, have become more frequent in the Australian landscape during the last 50 years [1,2]. This increasing trend is often related to various human activities and global warming. With continuing climate change, population growth, and increased interaction between people and the landscape, researchers have witnessed an increase of flammable fuels, and thus, a higher risk of more catastrophic fire events. The impact of wildfire on human health, lives and property, and the natural environment can be devastating and long-lasting [1,3,4]. Every year in Australia, billions of dollars are spent on preventing wildfires, mitigating the effects of fires, and on fire management activities [5–7]. However, traditional firefighting methods are not sufficient to battle the biggest fires. The 2019-2020 fire season in Australia's southeast was one of the worst the world had ever seen. Hundreds of fires burnt throughout the summer and the lives of 34 people were lost. Over 3 billion animals were affected while over 18 million hectares of the land were destroyed. Detecting fires soon after they form and predicting their behaviour are critical tasks for managing bushfires, alongside mapping fire risk in the landscape in response to ignition events. Understanding fire behaviour will allow monitoring of influential parameters that make ignition more likely and lead to the rapid expansion of the fire front, flame height, intensity, and the overall severity of the wildfire The risk associated with wildfires is a reflection of the availability of ignition sources, the likelihood of vegetation to ignite, meteorological conditions, and the rate of spread after ignition [9].

The spread of fire is uncertain but is explained by the conditions of the vegetation, the weather, and features of the landscape. Research has shown that the difference in moisture content over the landscape is one of the most influential parameters in the behaviour of fire [9,10]. This relationship helps explain the rate at which the fire moves over the land and enables researchers to estimate the possible extent of the burn area [11]. The burn area, or fire scar, delivers not only important information about the severity of the fire [12], but also helps us to understand why a particular area is affected and what led to this result. Thus, predicting the probability of the burn area in the landscape can provide helpful information to apply better forest management, predict the behaviour of a fire, and appropriately respond to fires [11,13–15]. Learning from historic fire events and analysing the vegetation, weather and terrain circumstances in the lead-up to these fires helps us understand past events and predict bushfires in the future.

Satellites and remotely sensed images can help us detect and map burn areas. Several contemporary satellites have onboard dedicated sensors to monitor vegetation and changes in the landscape [16]. However, because of their location above the atmosphere, there are limitations on the resolution of the images, interference from clouds and gaps in coverage due to the time satellites need to circle the globe. [17–19]. Different satellites deliver different data products and combining these can improve the performance of remotely sensed data products that aim to detect wildfires rapidly and predict the behaviour of fires. Knowing the possible burn area of fire from the moments it ignites is useful information in a targeted approach of fire response teams.

Unfortunately, it takes time for a satellite to orbit the world and detect a fire. It also takes time for scientists to predict the path of the fire, often using outdated fire behaviour models. The delay is compounded by the time taken for fire crews to reach the fire front by land or air - valuable time we often do not have. That is where the use of machine learning (ML) and artificial intelligence (AI) plays a crucial role. Handling data derived from specialist satellites, interpreting weather patterns and accounting for the effect of terrain, and vegetation is time-consuming and needs more than a well-trained eye to identify anomalies. Thus, wildfire science and management rely on improved statistical capability, extra compute power, and the ability to identify complex relationships among data inputs [4,16]. With ML we can program a computer to learn to predict future fire events and how they behave [16]. Only recently has ML been deployed in the prediction of the burn area and fire occurrence in the landscape [11,20–24].

The research described in this document aims to deliver a proof-of-concept method for the prediction of the burn area as a result of wildfires in Australia. Thereby, it utilizes satellite-derived images as predictors of the vegetation, digital elevation models as predictors of the topography, meteorological databases for weather predictors, and lastly, historic wildfire events (burnt areas) to train an ML model to make predictions of future burn area with high accuracy. The model distinguishes between the probability of the burn area prediction in case of fire ignition.

1.1. *Identified Needs and Opportunities*

In this study, we identify the need to develop a fast and accurate prediction of the burn area in the event of fire ignitions using ML techniques and advanced data processing. In particular, we aim to correlate the occurrence of historic fire scars with vegetation, topographic and weather variables and hope to interpret these circumstances to predict the burn area of fires using ML in the future. Being able to predict the possible impact of fire if an ignition were to happen based on Sentinel-2 images, weather elements, and elevation aspects of landscape, offers rapid insight into the possible damage, vulnerabilities, and direction for fire expansion. This information is currently missing in the Australian context and can improve targeted fire-response and pre-fire hazard reduction.

Current state-of-the-art has explored the use of ML for burn area prediction in the landscape based on susceptibility [20,25,26] or under the conditional occurrence of ignition of one or more wildfires [4,22,24,27–30]. However, none of these approaches has been applied in the Australian context which has, compared to the rest of the world, a unique combination of weather patterns and vegetation types. The direct use of satellite images in ML for burn area prediction is limited to applications of MODIS [11,20,25], Landsat 8 [31], or Landsat 7 [15]. The satellite products are often low in spatial resolution (MODIS) or have a low temporal resolution by which they only revisit a single location every 8-16 days (Landsat 7-8). Taking the knowledge from these studies [11,15,20,25,31], and linking them with Sentinel-2 imagery, delivers higher combined spatial (20 m) and temporal resolution (5 days).

The concept designed in this research adds promising knowledge to the existing literature and enables future research opportunities. The outcome of burn area prediction can be improved with the inclusion of more parameters influencing wildfires. In particular, including the spread of fire ignition points would add significant accuracy to the prediction of the burn area. The step-by-step development of fires between the first ignition and the total burn area predicted in this research will allow researchers to look more closely at the circumstances during the evolving fire, and thus the behaviour. In this study, the ability of multispectral satellite imagery, elevation, historic fire frequency and climate data to predict burn probability is tested.

2. Data description

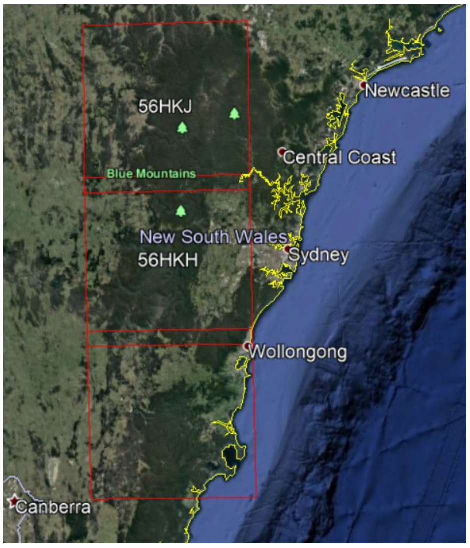

Figure 1: The area under study comprises three 100 km x 100 km Sentinel-2 tiles covering a region immediately west of Sydney.

The study area was located in the eastern part of New South Wales (NSW) and covered an area of approx. 30,000 km^2 (see Figure 1). The satellite imagery consisted of Sentinel-2 multispectral imagery and was downloaded from the Sentinel Australasia Regional Access (SARA) Data Hub. Sentinel-2 imagery has a spatial resolution of 10 - 20 m and temporal cadence between satellite revisits of up to 5 days. Sentinel-2 Level-2A data products (i.e. bottom-of-atmosphere reflectance) were downloaded for the period from December 2018 to July 2020 for three 100 km x 100 km tiles (labelled T56HKJ, T56HKH and T56HKG). The 10 spectral bands of Sentinel-2 imagery were cloud and cloud shadow masked and resampled to 20 m spatial resolution.

Shuttle Radar Topography Mission (SRTM)-derived digital elevation model (DEM) at 30 m resolution was downloaded from the NSW Department of Planning, Industry and Environment Data Portal. The DEM was further resampled to 20 m to match the Sentinel-2 imagery.

The fire history layer was extracted from the NSW Department of Planning, Industry and Environment Data Portal in vector format. This dataset contains information on fire extent and temporal length of all recorded wildfires, and prescribed burns, between 1902 and 2020 across NSW. The fire history between 1902 and 2018 was rasterized to create a fire frequency layer and was also resampled to 20 m resolution, matching the Sentinel-2 imagery.

Climate data were downloaded from the database of Australian climate data (SILO) in the period from December 2018 to July 2020. There were seven weather variables included in the analysis that were identified as potential influential factors that can increase the probability of bushfire ignition. These data at 5 km spatial resolution contained daily averages for (1) solar exposure consisting of both direct and diffuse components (in units of MJ/m^2), (2) maximum temperature (ºC), (3) Moreton's areal actual evapotranspiration (mm), (4) vapour pressure (hPA), (5) relative humidity at the time of maximum temperature (%), as well as daily totals of (6) class A pan evaporation (mm) and (7) rainfall (mm).

2.1. Data Gaps

The Data Quest is exploratory in nature and a significant outcome is to identify the 'gaps' in the input data. In other words, is there a need for specific (and realistically attainable) data that would drastically improve the outcomes of the project?

1) Temporal maps of ignition points would have greatly improved the predictive powers of this work. A large ignition point dataset was previously accessible via the NSW Department of Planning, Industry and Environment Data Portal. However, this was unfortunately not available at the time of this study as it had been withdrawn from the public domain due to issues with data quality and

currency. The Visible Infrared Imaging Radiometer Suite (VIIRS) 375 m thermal anomalies / active fire product available from 20 January 2012 to the present through NASA's Fire Information for Resource Management System was also investigated in this study. Unfortunately, it was deemed unsuitable due to low spatial resolution (i.e. 375 m) and inability to differentiate ignition points from individual hotspot/fire pixels.

2) The wind is a major driver of fires meaning that accurate and high-resolution data on wind speed and direction is a critical input to models of fire spread. Wind information is accessible on request from the Bureau of Meteorology but was not utilized in this research because of time constraints. Additionally, the resolution of the available gridded wind data is low, likely due to the very sparse distribution of weather stations in the landscape. The wind can change direction and speed multiple times over the course of the fire, making it difficult to choose a single measurement for each individual fire. Rather, models should take into account the temporal weather changes and leverage these for inference.
3) A finer grid of climate measurements could have been generated by downloading climate information from the database of Australian climate data (SILO) for individual stations spread across the study area and interpolating them through co-kriging, in conjunction with DEM information to a finer grid than 5 km. However, this was not attempted because of time constraints.

3. Methodology

Our approach was to train image segmentation models using a data-stack of Sentinel-2 images, rasterized weather and climate measurements, gridded DEM data and historic burn area outlines. The burn area data was used to make 'ground-truth' binary masks that were the target variable when training the models. The two models chosen for this work were UNet (https://arxiv.org/abs/1505.04597) - a well-studied semantic segmentation convolutional neural network (CNN) - and BASNet (https://ieeexplore.ieee.org/document/8953756) - a newer refinement of the UNet architecture that focuses on correctly predicting boundaries.

3.1. Workflow overview

Figure 2 shows a schematic of the inputs used to train our models. Note the difference between the two model's inputs and output predictions.

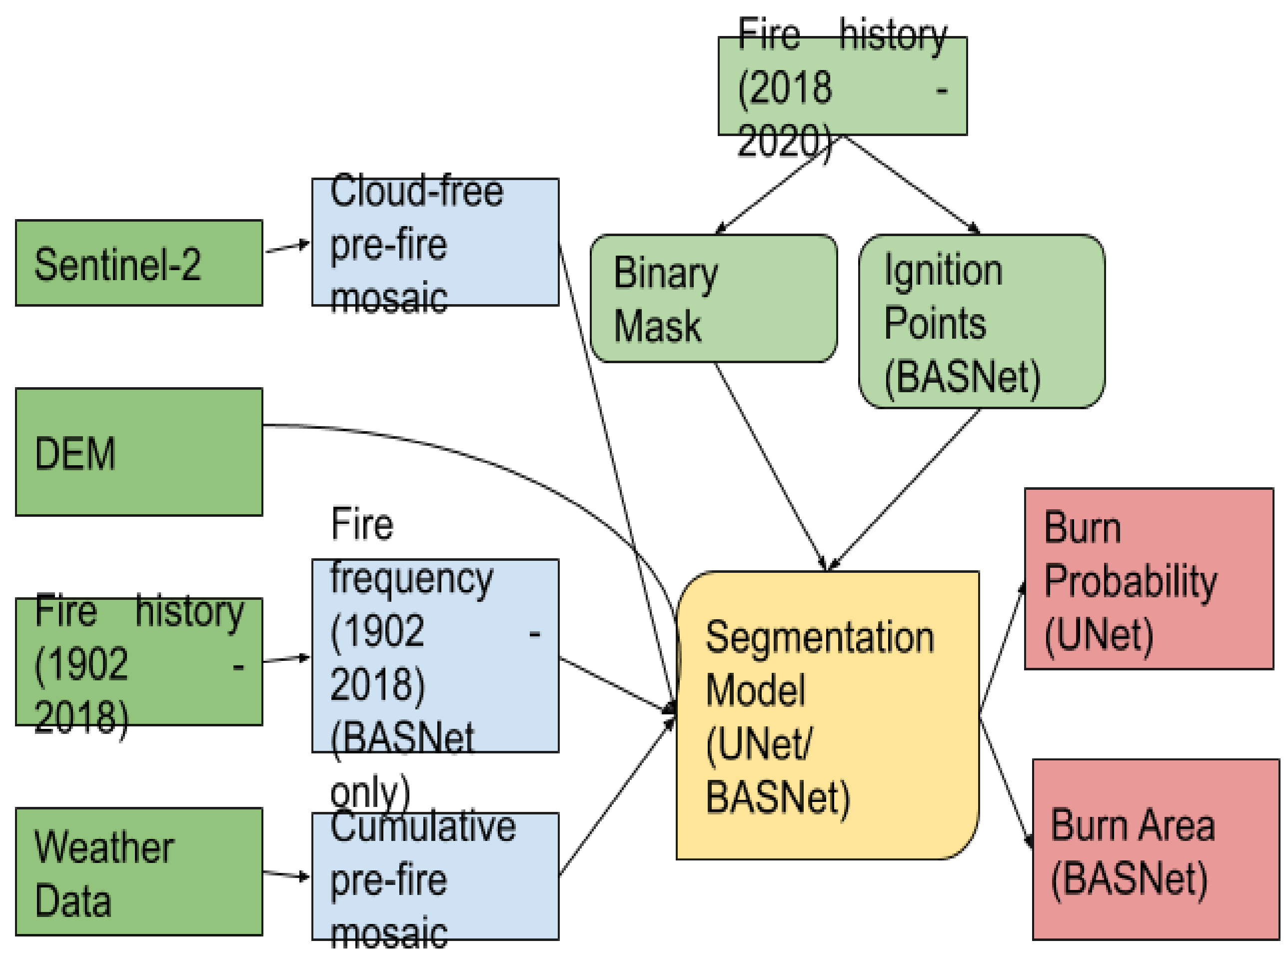

Figure 2: Flow diagram showing the training inputs to the BASNet segmentation model. Sentinel-2 imagery from 28 days prior to training fired is assembled into a cloud-free pre-fire mosaic. Fire history data is transformed in a fire frequency heat map and weather data is assembled into a cumulative pre-fire mosaic giving average conditions on the day of the fire. The response variable is a binary mask of burnt / non-burnt areas generated from the fire history dataset. The model outputs a map of 'burn probability'.

3.1.1. Response variables

Individual burn scars within each Sentinel-2 tile extent (T56HKJ, T56HKH, T56HKG) between 16 December 2018 and 3 July 2020 were rasterized at 20 m resolution to generate binary masks of non-burned (0) and burned (1) areas and were used as response variables. Within the study area, there were in total of 130 burn scars that originated from wildfires and prescribed burns in this period.

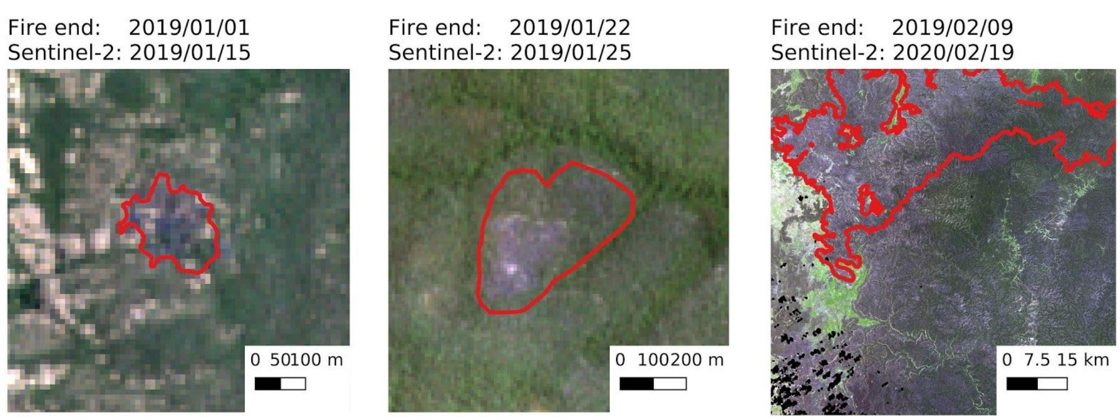

Figure 3: Examples of fire scar data, distributed as closed polygons that outline the burnt area. Note that the scar outline sometimes does not correspond to the burnt area (which is larger) because multiple fire events merged to form a blended scar area (right-most panel)

3.1.2. Predictor variables

Both the fire frequency layer and SRTM-derived DEM were directly used as predictors of burn probability. In contrast, to generate cloud-free composites of satellite imagery, cumulative mosaics of Sentinel-2 imagery were created using all images 28 days prior to each fire event. Pixel values feeding into each mosaic were chosen from the time of maximum normalized difference vegetation index (NDVI). That is, a Sentinel-2 mosaic was derived by selecting pixel values of each band that corresponded to the maximum NDVI of all Sentinel-2 images in a 28 days period prior to each fire event.

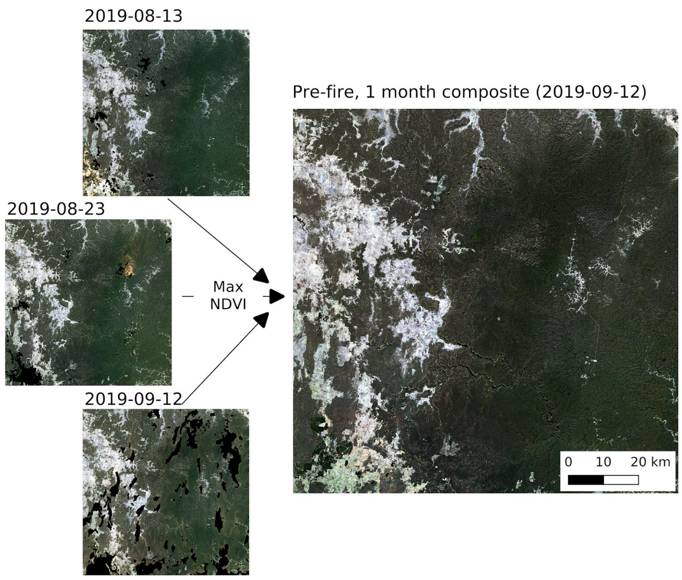

Figure 4: Illustration of how images from a 28-day window are used to build a cloud-free image of the landscape, as seen by Sentinel-2.

Similarly, to take temporal information into account, daily averages of climate data were further cumulatively averaged, while daily totals were cumulatively summed in the period of 28 days prior to each fire event.

3.2. Burn Probability Prediction

In order to predict a probability (e.g. risk) of burning for every single 20 m x 20 m pixel in the Sentinel-2 images, we build upon the UNet segmentation architecture. We modify the input layer to feed a stack of 14 channels of data: 10 channels from the Sentinel-2 image, relative humidity, evapotranspiration, maximum temperature within a defined period of time before the start of the fire and digital elevation. We train the model using Cross-Entropy loss to predict the scars associated with each input stack. We use Adam optimization to train our model.

3.3. Burn Area Prediction

The previous model is able to correctly learn to associate a burn probability to each pixel based on its different features related to the amount of fuel, the topography of the region and pre-fire history of weather variables, however, we empirically observed that it has a limited capacity to predict the exact boundary of the burn area. This limitation is mainly due to the loss function used to optimize the model. The loss function treats all the wrong predictions equally in them irrespective of their location and proximity to the actual burn area. In order to address this, we used a further developed version of UNet known as BASNet. The new model is relatively similar in terms of architecture to the UNet model. The main difference, however, is the incorporation of Structural SIMilarity (SSIM) and Intersection over Union (IoU) objectives that help limit the burn area prediction as a single unified area. We use the new model with the same input stack to predict the burn area.

3.4. Ignition Point Sampling

There are other confounders that influence the burn area of a fire. For instance, the human intervention, weather variables during the fire and most importantly the points where the fire starts from, e.g. ignition points. The prediction of the final burn area in the absence of these confounders is an uncertain prediction. In order to address the lack of the above-mentioned data sources, we propose to sample ignition points from the burn area. Intuitively, we assume every single point in the fire scar has an equal probability of being an ignition point. During training, we uniformly sample five ignition points and feed the points as an extra channel added to the input stack of images. This way we condition our model's burn area prediction on the configuration of these ignition points. In practice, we can use real-time detection of ground-truth ignition points to replace the pseudo-ignition point created by our model for more accurate burn area prediction.

3.5. Multi-Step Burn Area Prediction

Conditioning our model on the ignition points not only increases the burn area prediction accuracy but it also endows extra capabilities to it. One common scenario in a real bushfire is that new ignition points are detected as the fire front develops. Therefore, it is desirable for any predictive model to be able to update its predictions over time given the new contextual information. This context in our case is the configuration of the observed ignition points. By adding the newly detected ignition points overtime to the set of previously observed ones and rerunning the model we can produce new burn area predictions. In order to showcase this capability, we simulate such a scenario by using a sampling method. In our method, we assume a Gaussian distribution of unit standard deviation around each observed ignition point. As we observe more ignition points, we add new distributions to the mixture of Gaussians. We continue sampling new ignition points from the previous distribution and adding the new sample to the previous set of points. This way we simulate a real scenario by assuming that the points closer to one of the previous ignition points have a higher probability of burning. Using this method, we can generate a multi-step fire development map.

In practice, we can add further priors to the ignition sampling method like the proximity to fire combat units, urban areas, water resources etc. Additionally, we can feed actual ignition points as we detect them using other methods, whether using other remote sensing measures or on the ground observation units.

4. Results and Discussion

Figure 5 shows the results produced by the UNet model when applied to a previously unseen geographic location. In this figure, we can see that our model is able to associate a high burning probability to the areas (mainly in the middle) that were mostly burnt after the fire.

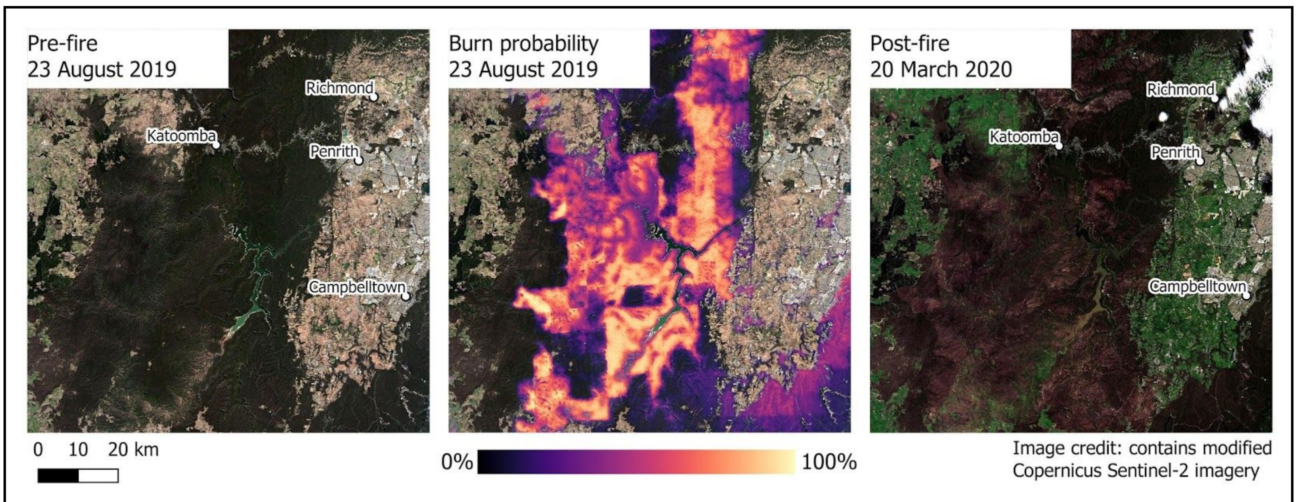

Figure 5: The *left panel* shows the true-colour Sentinel-2 image of a region on the outskirts of Sydney that was not included in the training data. This pre-fire image is used to make the prediction. The *middle*

panel shows the same image overlaid by the burn probability map produced by the UNet model. The pattern of squares seen in the map is an artefact of the coarsely gridded climate data. The *right panel* shows the post-fire Sentinel-2 image. Note the extensive burn scar (brown) and that green vegetation has sprung up after consecutive rain events.

Figure 6 below shows the results produced by our final model, e.g. BASNet including the ignition point sampling method. The model is able to accurately predict the burn area. We achieved an overall pixel-wise accuracy of 88% on an unseen validation set.

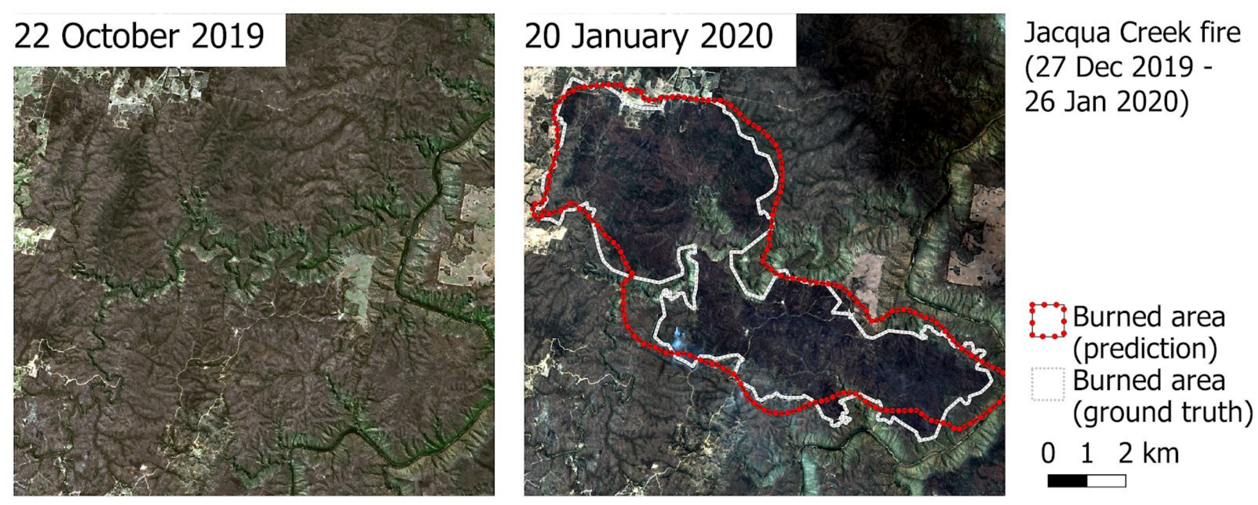

Figure 6: The *left panel* shows the true-colour Sentinel-2 image of a region on the outskirts of Sydney that was not included in the training data, before the fire. The *right panel* shows the same image overlaid by the predicted burn area (red dotted boundary) compared to the ground-truth burn area (white boundary). Our ignition sampling method in conjunction with BASNet is able to produce accurate burn area predictions.

5. Next Steps and Recommendations

In order to improve the burn probability and burn area prediction results, the following list of recommendations is discussed:

1) Incorporate fire ignition data from VIIRS or other data sources.
   In this study, it was highlighted that both in the data collection as well as the workflow, additional ignition data would be helpful. In the current research, ignition points were assumed and randomly sampled for the burn area prediction. There is no reliable data on ignition points, which would improve the accuracy of predictions and enable further research into the spread and of fires over the landscape. A classification of ignition data would also be helpful in understanding the nature of bushfires. Lightning strikes form a major component of ignition data, which are currently not available in the public domain.
2) Clean/Verify the fire history data.
   As indicated in Figure 3, some of the historic fire scars do not match the corresponding Sentinel-2 imagery. This was most prevalent in the small fires. Larger fires were often segmented in geospatial format and needed extra processing. To enhance the readiness of the data, historic fire events should undergo additional verification with satellite imagery to cross-check the boundary and level of overlap and spacing when segmented.
3) Incorporate wind information during fire events.
   In this study, we acknowledge the complex behaviour of wind as a predictor weather variable. It is highlighted that additional models would be required to investigate the direction and speed of wind during fire events. The wind is an important parameter that highly influences the speed and

direction of the moving fire front. Measures of wind will influence the prediction of the burn area and can further enhance fire behaviour models.

4) Validation of predictions with other burn area indicators.
Even though high accuracy in the prediction was achieved using the above-mentioned methods, extensive validation of the models should be performed before deploying the models in field situations. For example, as a proxy for vegetation, satellite-derived data was used, however, no additional data-processing was performed to indicate the dryness of fuel on the forest floor. To validate the accuracy of satellite data, it is worth comparing imagery derived fuel moisture content versus field measures of moisture content. These measures are also highly seasonal and vary through the different layers in the forest.

As a general comment, we can highlight that increased accuracy in the data would increase the accuracy in the prediction. Moving forward, the next logical step would be in-field validation of the prediction. Thereby we could compare the prediction of a burn area if an ignition were to happen, with field measurements of moisture content and flammability. As we expand the model beyond the three sentinel tiles selected for this study, we will require additional validation with natural fire boundaries such as water bodies and infrastructure, as these will influence the burn area. By expanding the study area, to the whole of Australia, the accuracy of the models will need to be validated with changes in the vegetation. Currently, the models are tested on forested areas within the "bushfire" topic, however, the likeness of fires to originate and expand over other vegetation types such as savannah and grassland should be incorporated. Thus, additional time and resources will be required to further validate and implement the tool.

The applied methods and possible application hold great promise for future research. Not only, were promising results discussed, and were first steps into the validation made. The tool has obvious utilization, pre- and during fire events. Pre-fire, high-risk areas can be assessed on the basis an ignition would happen at any time from the latest satellite image. This can inform forest management crews ahead of possible fires about the possible damage caused by fires and the area it could consume. During fires, when the ignition point is recorded, the model can more accurately predict the direction, extent and boundary of the fire within seconds. This can inform fire crews where to interfere to prevent further spreading. Additionally, the tool can be integrated with fuel moisture content models, using improved moisture content measures in the field to enhance the predictions. With the correct ignition data at the start and during the fire, the methods can be employed to predict the behaviour of fires during the fire, with an accurate prediction of the direction, speed and flame height of fires based on the current conditions. Additional data sources such as lightning strikes and wind data are critical.

6. Appendix: Technical Requirements

We use the PyTorch deep learning package for all our models. Both our models are trained on patches of 256x256 pixels on 20 m resolution. We use a batch size of 20, the learning rate of $1e^{-4}$ and weight decay regularisation of $1e^{-4}$ for training. We crop the input tiles into patches with a stride of 64 pixels. For inference, we use a larger patch size of 512x512 pixels. Finally, we stitch the small patch predictions using a sliding window averaging method combined with boundary Spline weighting.

Our model takes 10~12 hours to train on the whole training set of ~100 Sentinel-2 tiles using 2 V100 GPUs. Source code is available on Gitlab.

References

[1] Borchers Arriagada N, Bowman DMJSJS, Palmer AJ, Johnston FH. Climate Change, Wildfires, Heatwaves and Health Impacts in Australia. Extrem. Weather Events Hum. Heal., Springer, Cham; 2019, p. 99–116. https://doi.org/https://doi.org/10.1007/978-3-030-23773-8_8.

[2] Etchells H, O'Donnell AJ, Lachlan McCaw W, Grierson PF. Fire severity impacts on tree mortality and post-fire recruitment in tall eucalypt forests of southwest Australia. For Ecol Manage 2020;459:117850. https://doi.org/10.1016/j.foreco.2019.117850.

[3] Kasin JA, Papastathopoulos I. A spatial Poisson hurdle model with application to wildfires 2020.

[4] Coffield SR, Graff CA, Chen Y, Smyth P, Foufoula-Georgiou E, Randerson JT. Machine learning to predict final fire size at the time of ignition. Int J Wildl Fire 2019;28:861–73. https://doi.org/10.1071/WF19023.

[5] Yu P, Xu R, Abramson MJ, Li S, Guo Y. Bushfires in Australia: a serious health emergency under climate change. Lancet Planet Heal 2020;4:e7–8. https://doi.org/10.1016/S2542-5196(19)30267-0.

[6] Ashe B, McAneney KJ, Pitman AJ. Total cost of fire in Australia. J Risk Res 2009;12:121–36. https://doi.org/10.1080/13669870802648528.

[7] Ximenes F, Stephens M, Brown M, Law B, Mylek M, Schirmer J, et al. Mechanical fuel load reduction in Australia: a potential tool for bushfire mitigation. Aust For 2017;80:1–11. https://doi.org/10.1080/00049158.2017.1311200.

[8] Storey MA, Price OF, Bradstock RA, Sharples JJ. Analysis of Variation in Distance, Number, and Distribution of Spotting in Southeast Australian Wildfires. Fire 2020;3:10. https://doi.org/10.3390/fire3020010.

[9] Rao K, Williams AP, Flefil JF, Konings AG. SAR-enhanced mapping of live fuel moisture content. Remote Sens Environ 2020;245:111797. https://doi.org/10.1016/j.rse.2020.111797.

[10] Yebra M, Dennison PE, Chuvieco E, Riaño D, Zylstra P, Hunt ER, et al. A global review of remote sensing of live fuel moisture content for fire danger assessment: Moving towards operational products. Remote Sens Environ 2013;136:455–68. https://doi.org/10.1016/j.rse.2013.05.029.

[11] Maffei C, Menenti M. Predicting forest fires burned area and rate of spread from pre-fire multispectral satellite measurements. ISPRS J Photogramm Remote Sens 2019;158:263–78. https://doi.org/10.1016/j.isprsjprs.2019.10.013.

[12] Keyser AR, Westerling ALR. Predicting increasing high severity area burned for three forested regions in the western United States using extreme value theory. For Ecol Manage 2019;432:694–706. https://doi.org/10.1016/j.foreco.2018.09.027.

[13] Flannigan MD, Wotton BM, Marshall GA, de Groot WJ, Johnston J, Jurko N, et al. Fuel moisture sensitivity to temperature and precipitation: climate change implications. Clim Change 2016;134:59–71. https://doi.org/10.1007/s10584-015-1521-0.

[14] Podschwit HR, Larkin NK, Steel EA, Cullen A, Alvarado E. Multi-model forecasts of very-large fire occurences during the end of the 21st century. Climate 2018;6:1–21. https://doi.org/10.3390/cli6040100.

[15] Pu R, Gong P. Determination of burnt scars using logistic regression and neural network techniques from a single post-fire Landsat 7 ETM+ image. Photogramm Eng Remote Sensing 2004;70:841–50. https://doi.org/10.14358/PERS.70.7.841.

[16] Jain P, Coogan SCP, Subramanian SG, Crowley M, Taylor S, Flannigan MD. A review of machine learning applications in wildfire science and management 2020.

[17] Chuvieco E, Mouillot F, van der Werf GR, San Miguel J, Tanasse M, Koutsias N, et al. Historical background and current developments for mapping burned area from satellite Earth observation. Remote Sens Environ 2019;225:45–64. https://doi.org/10.1016/j.rse.2019.02.013.

[18] Plank S, Karg S, Martinis S. Full-polarimetric burn scar mapping–the differences of active fire and post-fire situations. Int J Remote Sens 2019;40:253–68. https://doi.org/10.1080/01431161.2018.1512768.

[19] Chuvieco E, Aguado I, Salas J, García M, Yebra M, Oliva P. Satellite Remote Sensing Contributions to Wildland Fire Science and Management. Curr For Reports 2020;6:81–96. https://doi.org/10.1007/s40725-020-00116-5.

[20] Mayr MJ, Vanselow KA, Samimi C. Fire regimes at the arid fringe: A 16-year remote sensing perspective (2000–2016) on the controls of fire activity in Namibia from spatial predictive models. Ecol Indic 2018;91:324–37. https://doi.org/10.1016/j.ecolind.2018.04.022.

[21] De Bem PP, De Carvalho OA, Matricardi EAT, Guimarães RF, Gomes RAT. Predicting wildfire vulnerability using logistic regression and artificial neural networks: A case study in Brazil's Federal District. Int J Wildl Fire 2019;28:35–45. https://doi.org/10.1071/WF18018.

[22] Xie Y, Peng M. Forest fire forecasting using ensemble learning approaches. Neural Comput Appl 2019;31:4541–50. https://doi.org/10.1007/s00521-018-3515-0.

[23] Mote T, Singh A, Prasad M, Kalwar P. Predicting burned areas of forest fires: An artificial intelligence approach. Int J Tech Res Appl 2017:56–8.

[24] Castelli M, Vanneschi L, Popovi A. Predicting burned areas of forest fires: an artificial intelligence approach. Fire Ecol 2015;11:106–18. https://doi.org/10.4996/fireecology.110106.

[25] Archibald S, Roy DP, van Wilgen BW, Scholes RJ. What limits fire? An examination of drivers of burnt area in Southern Africa. Glob Chang Biol 2009;15:613–30. https://doi.org/10.1111/j.1365-2486.2008.01754.x.

[26] Arnold JD, Brewer SC, Dennison PE. Modeling climate-fire connections within the great basin and upper Colorado River Basin, Western United States. Fire Ecol 2014;10:64–75. https://doi.org/10.4996/fireecology.1002064.

[27] Cortez P, Morais A. A Data Mining Approach to Predict Forest Fires using Meteorological Data. Proc 13th Port Conf Artif Intell 2007:512–23.

[28] Toujani A, Achour H, Faïz S. Estimating Forest Fire Losses Using Stochastic Approach: Case Study of the Kroumiria Mountains (Northwestern Tunisia). Appl Artif Intell 2018;32:882–906. https://doi.org/10.1080/08839514.2018.1514808.

[29] Liang H, Zhang M, Wang H. A Neural Network Model for Wildfire Scale Prediction Using Meteorological Factors. IEEE Access 2019;7:176746–55. https://doi.org/10.1109/ACCESS.2019.2957837.

[30] Mitsopoulos I, Mallinis G. A data-driven approach to assess large fire size generation in Greece. Nat Hazards 2017;88:1591–607. https://doi.org/10.1007/s11069-017-2934-z.

[31] Zald HSJ, Dunn CJ. Severe fire weather and intensive forest management increase fire severity in a multi-ownership landscape. Ecol Appl 2018;28:1068–80. https://doi.org/10.1002/eap.1710.

EARLY DETECTION

DATA QUEST 2020

EARLY DETECTION OF FIRE IGNITION

BUSHFIRE DATA QUEST 2020

Technical Memorandum

Detecting Fires Earlier Using Image Stacking and Super-resolution

Katharine Melnik (Scion Crown Research Institute, NZ)

Ilze Pretorius (Scion Crown Research Institute, NZ)

Alex Codoreanu (Swinburne University Gravitational Wave Data Centre)

Jack White (Swinburne University of Technology)

Science and Data Leads

Dr Ruth Luscombe (Fireball International Pty Ltd)

Dr Marta Yebra (Australian National University)

Dr Chedi Raïssi (Ubisoft, Singapore)

Abstract / Executive Summary

Early fire detection is vital for effectiveness and ultimate success of fire emergency response. Geostationary satellites such as Himawari-8 provide data at a high temporal resolution (at least every 10-minutes across the entire disk and more frequently in special regions) necessary for identifying fire ignitions soon after they occur. However, the spatial resolution of Himawari-8 images is 1km × 1km or lower (compared to spatial resolutions in the order of ~400 m from polar-orbiting satellites), and there is a need to process the lower-resolution geostationary satellite data in a way that will increase the detectability of wildfire ignitions.

We explored two techniques for improving Himawari-8 image quality with the aim to increase the effectiveness of current fire detection methods:

1. image stacking coupled with image subtraction – averaging several previous images to form a "baseline" image and subtracting it from the newest available image to make temperature anomalies clearly visible; and
2. super-resolution – a machine-learning-driven technique for increasing the resolution of an image.

We stacked the pre-fire images by calculating the average value across the images for each pixel, and subtracted the resulting image stack from the ignition image (the image obtained just after ignition. After this we used a previously developed super-resolution network to increase the resolution of the pre-ignition stacked image, the ignition single image, and the resulting subtracted image. The Orroral Valley fire hotspot is more clearly visible in the subtracted image compared with the pre-ignition image alone, and is even more distinctive in the super-resolved subtracted image. This technique, coupled with the current state-of-the-art fire detection models, promises to increase the speed and reliability of ignition detection and fire location accuracy, boosting the effectiveness of fire emergency personnel by allowing them to fight fires early while they are at a manageable size and require less resources.

Introduction

Hundreds of fires are ignited by lightning during the Australian bushfire season. This often happens in remote areas where fires can become large and unmanageable by the time they are detected.

Bushfires can significantly impact ecosystems and society, and often require extensive management to mitigate negative effects. Early fire detection is vital for the effectiveness and ultimate success of fire emergency response, and satellite data sometimes provides the only way to identify fire ignitions when they occur in remote areas. There are two types of satellites: polar-orbiting satellites that orbit the Earth and collect high-resolution information of each location as they pass over it (usually twice a day), and geostationary satellites that travel at the speed of Earth's rotation and collect lower-resolution data of the full disk (the whole Earth as visible from the satellite location) at predetermined intervals. Since the long return interval of polar-orbiting satellites makes them unsuitable for early fire detection, there is a need to process the lower-resolution geostationary satellite data in a way that will increase the detectability of wildfire ignitions.

In the present work, we investigate the ability of image stacking and machine-learning techniques to detect fires from orbit soon after ignition using fast imaging. Detecting and localizing fires earlier will enable small teams to access blazes before they get out of control thereby reducing the damage caused by fires, potentially saving lives and decreasing fire management costs.

Identified Needs and Opportunities

The existing Himawari-8 fire detection product has a thresholding approach, where the difference between band-7 and band-14 brightness temperature is calculated and if it's greater than 30 Kelvin (K) during the day, and 15 at night, a fire detect is registered (Xie et al 2018, Xu and Zhong 2017). This characteristic active fire signature is the result of the enormous difference in 4- and 11-Angstrom blackbody radiation emitted at combustion temperatures as described by the Planck function. The threshold values (30 and 15 K for day and night, respectively) are conservatively chosen to reduce the number of false positive detections (Xie et al 2018). However, a consequence of the thresholding approach is that small/early fires are often missed, as these often fall below the chosen threshold value (Xie et al 2018, Lin et al 2016). Therefore, there is a need to improve the current fire detection pipelines so reliable detections of small/early fires are possible.

We identify the following three opportunities to improve on current methods using machine learning techniques and advanced processing:

1) Increase the signal to noise ratio in the detected image
2) Perform image subtraction to identify fire pixels at lower threshold values without increasing false positive detections
3) Increase the resolution of the image to identify the fire location and/or extent with increased precision

Interestingly, we decided on the above approach before reading a paper by Pennypacker *et al*. (2013), which recommends a very similar approach for a proposed satellite system. The fact that we independently arrived at the same conclusions as the authors of this paper testifies to the validity of this approach. To our knowledge, this is the first time this methodology was applied to geostationary satellite data for fire detection.

Data description

Figure 1 shows a typical true-colour view of Australia and New Zealand from the Advanced Himawari Imager (AHI) on the Himawari-8 satellite. In order to validate our approach we decided to pull data for all 13 bands of the AHI data on the first day of our target fire (the Orroral Valley Fire, near Canberra), along with the previous two days of data from 0900 to 1600. The central wavelength, resolution and bandwidth of each AHI band is described here and we pulled all 13 bands in order to:

- make true color images of our area of interest,
- investigate stacking strategies for each band to determine the point at which stacking stops adding any useful information,
- validate our subtraction hypothesis on all bands, and
- reproduce previous detection strategies (Xie et al 2018).

We ended up with wavelength coverage from 0.455μm to 13.30μm of the Orroral Valley fire ignition area from 0900 to 1600 for the 25th, 26th and 27th of January 2020 at 10 minute intervals. This allows us to make an on-the-ground comparison before, at and after the first reported ignition, which occurred at 13.50 on the 27th of January.

We created 2 domain regions, a *large* domain to be used for visual inspection of true color images and a smaller *science* domain to be used for the proposed detection pipeline. The *large* domain is 200 × 200 pixel square defined by an upper left corner at 147.66°, -34.35° and lower right corner at 150.28°, -36.89°. The *science* domain is a 60 × 64 pixel rectangle defined by an upper left corner at 148.54°, -35.18° and lower right corner at 149.33°, -35.99°.

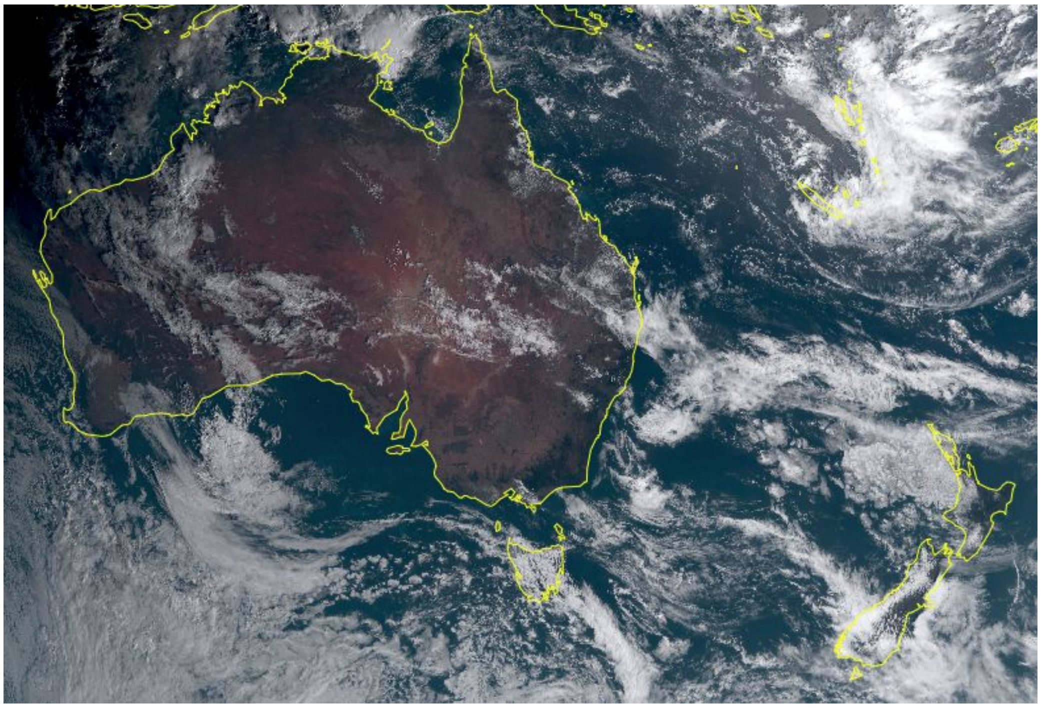

Figure 1. True-colour RGB image of Australia and New Zealand on December 4th, captured from the JAXA Himawari Monitor at https://www.eorc.jaxa.jp/ptree/.

Data Gaps

In this study we used the full disk imagery, which arrives every 10 minutes. However, Himawari-8 also offers additional observation modes. These can increase the acquisition frequency to one image every 30 seconds (Figure 2). This higher frequency can have a significant positive impact on anomaly detection algorithms, such as ours.

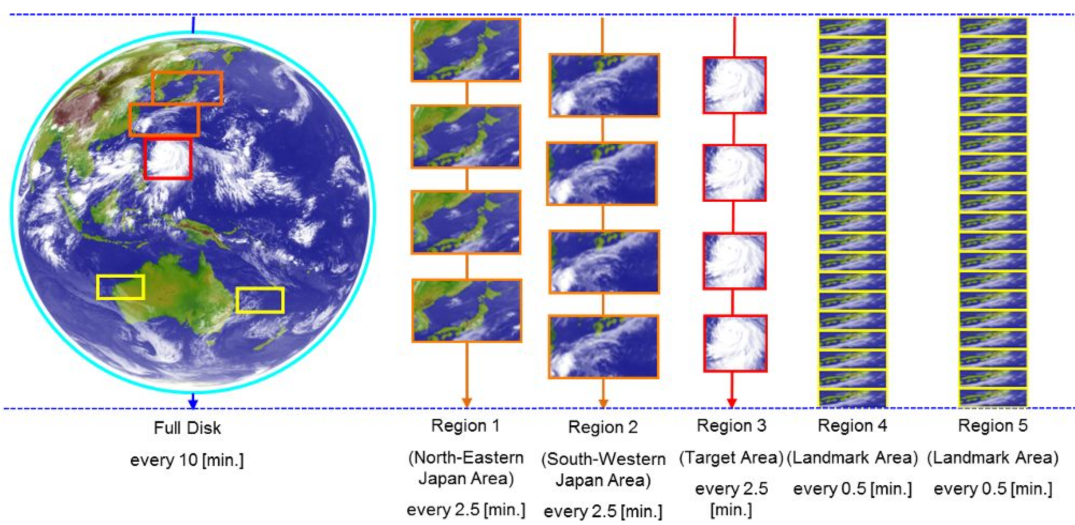

Figure 2. Diagram showing the locations of the three regional frames that are acquired every 2.5 or 0.5 minutes, compared to every 10-minutes for the full-disk images. *Image credit: https://spaceflight101.com/spacecraft/himawari-8-and-9/*

For the purposes of this study, we visually inspected the data in order to identify similar non-cloudy days. Additional work and data would be needed in order to automate both the cloud-masking and plume vs. cloud identification.

Ground-based sub-mm radar with coverage over remote areas can also provide an important complementary data set to facilitate the validation of satellite based triggers by directly probing the resulting smoke plumes.

Data Exploration

Image stacking is a crucial component of our novel detection strategy and one of the first questions we wanted to answer was: how many images would be enough and would this number be dependent on the band considered?

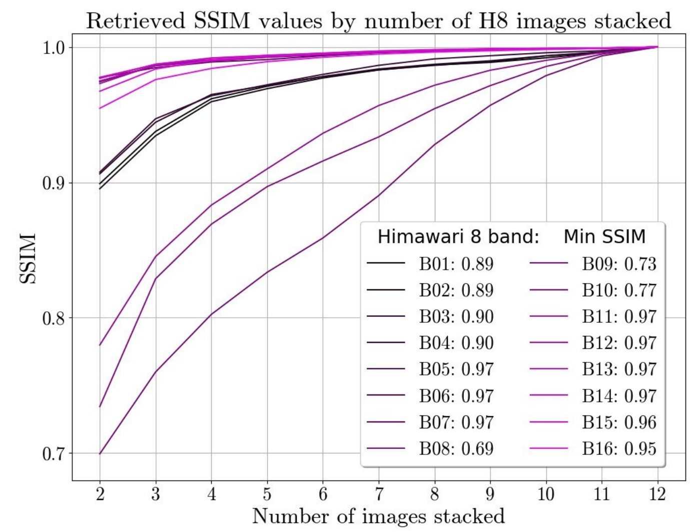

Figure 3. Structural Similarity Index as a function of number of images used in the averaging stack.

We decided to investigate the above question by iteratively stacking up to 12 images. We used this final stacked image as the 'best' image to be compared to each subsequent stack. In order to quantify the relative similarity for each comparison, we computed the *structural similarity index* (SSIM). The results are presented in Figure 3.

We found that the improvement in similarity is not uniform across the bands, and the curves flatten out for all bands. Interestingly, the greatest gain in quality is across bands 8 through 10. We do not use these bands but simply highlight this gain. For the purposes of our study, we decided to focus on band 7 images, and to stack up to 6 images together.

Methodology

Workflow overview

Our workflow consists of four individual steps:

- create a preprocessed stacked image for each 10-minute interval composed of images from an hour of data;
- extract the relevant stamp from the incoming full disk image;
- subtract the incoming stamp from the appropriate stacked image; and
- identify hotspots, if present.

A description of the workflow with reference to a case study of the early detection of the Orroral Valley Fire in a remote area in south-eastern Australia is given in the following sections, and illustrated in Figure 4.

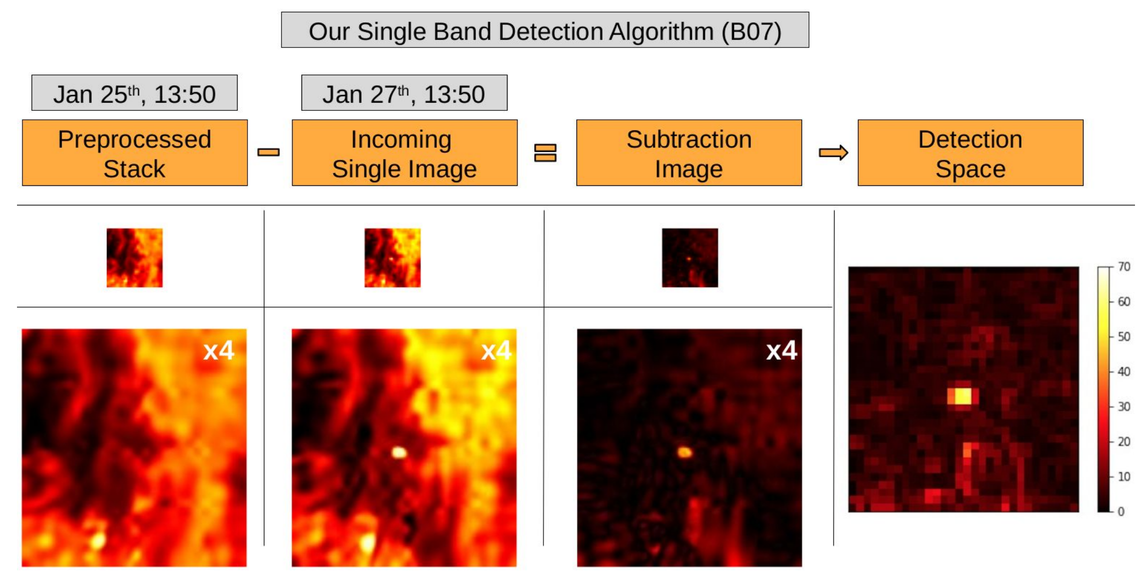

Figure 4. Illustration of the workflow used to create the detection image from Band-7 (B07) data.

Data inputs

For the purpose of this study, all Himawari-8 full disk images were saved to local storage and the *large* and *science* domain stamps were extracted from each one of them. This was done in order to minimise computing time/overhead as this study was a proof of concept for the novel remote wildfire detection method described in this document. All operations in the processing were applied to Band-7 data only.

Super Resolution

The *large* and *science* domain stamps were then super-resolved using a pre-trained Enhanced Deep Super-Resolution Network (EDSR) (Lim et al, 2017). EDSR is a deep residual network that reconstructs input images to a given upscaled resolution. The model can increase the resolution of an input image by up to four times the original resolution through inferring the correct pixel intensities between existing pixels. Utilising EDSR in this workflow allowed us to run the follow-up analysis in parallel, for both datasets. The super-resolved data has been upscaled so that each pixel has an on-the-ground size of 0.5 × 0.5km, compared to 2.0 × 2.0km in the original images.

The EDSR model was pre-trained on a diverse set of high-resolution/low-resolution pairs from the DIV2K dataset (Agustsson, E., & Timofte, R., 2017). Although the image content in the training data is not Earth-observation based, the generalisability of these networks are sufficient to capture general low-level features and upscale them efficiently. However, some reconstruction quality is sacrificed when the input Earth-observation image is relatively featureless (e.g., large bodies of water, desserts). With custom training data that is comprised of Earth-observation images and featuring fire affected regions, we expect that these reconstruction losses can be addressed for future integration into a fire detection pipeline.

Integrating this model into a pipeline would also warrant computational optimisation of the network for a particular resolution, as the speed of the reconstruction is dependent on the resolution of the input image. Considering the time-dependence that early fire detection requires, measuring the reconstruction speed and then exploring the resolution/reconstruction time dependence could save valuable time in a future emergency response system. Although it is worth noting that current reconstruction speeds for a single image are on the of less than a minute for an image of around 800x800 pixels being reconstructed to x4 the original resolution.

Image Stacking

Next, we created an associated stacked image for each 10-minute interval starting from 0900 to 1600 for the two days considered:

- 25th of January - 2 days before the fire with similar clear sky conditions as the ignition time
- 26th of January - 1 day before the fire with cloudy conditions over the ignition area

We stacked images by taking a simple mean of the input images, but other schemes are also possible, for example, median stacking or weighting images using a measure of noise. At this point, we had the single incoming images as they would have arrived in real-time, along with the image stacks from the previous day/s. This allowed us to effectively simulate an observation strategy where a new incoming image is super-resolved and then compared to a stacked image from a previous day, with and without similar cloud conditions. For the original-scale images, we are able to send an alert within 10 seconds of receiving the full disk image while the super-resolution step adds another 30 to 40 seconds to the total processing time. Currently, the output of our pipeline is a subtraction image that displays the relative temperature change for each pixel.

With this workflow, we would be able to identify and localise a wildfire event within one minute of receiving an incoming image.

Asides – Paths not Taken

We briefly explored the potential of detecting fires in bands dominated by signals from particular ion or gas emission lines present in flames or smoke and compiled a summary of potential emission lines (Table 1). We decided not to pursue this path due to the time constraints of a one-week research sprint, but we identified that there may be potential in this space. We highlight that there are visible traces in Band-1 which most likely trace the CO_2 released by wildfires. This suggests that there is scope for this tracer to be used as a discriminator to minimise false-positive alerts.

Table 1. A summary of wavelengths from particular ion or gas emissions present in flames or smoke.

Element/ molecule	**Peak 1 (µm)**	**Peak 2 (µm)**	**Notes**	**Main Reference(s)**
Potassium (K^+)	0.7665	0.7699	The 769.9 peak overlaps with O2 absorption in the atmosphere, therefore the 766.5 nm line would have the highest intensity when viewed through the atmosphere. K only emitted during flaming fire and can be detected through smoke	Amici, et al., 2011
Sodium (Na^+)	0.5893		The Na emission signal magnitude is substantially lower than that of potassium, presumably because of the higher	Amici, et al., 2011

			excitation energy and lower percentage dry weight of the former element in vegetation fuel	
Potassium (K^+)	see above			
Chlorine (Cl^-)	0.4526	0.7769	Chlorine has 11 spectral lines, the two included here are the most intense Couldn't find any evidence that this has been used in fire detection before.	Bolshakov and Barnes, 1997.
Sulfate (SO_4^{2-})			Couldn't find any evidence that this has been used in fire detection before.	
Carbon Dioxide (CO_2)	~1.61		There are other peaks, but some are absorbed by the atmosphere.	Ross et al, 2013.
Methane (CH_4)	~1.65			Ross et al, 2013.

Other options that showed potential, but could not be researched further due to time constraints, were:

- using Synthetic Aperture Radar (SAR) instruments on satellites to detect fires through cloud
- using ground-based weather radars to detect smoke plumes from early fires.

Results and Discussion

We demonstrated the feasibility of using single-band images (B07) to effectively identify a remote wildfire. Our detection strategy is unique in that we approach the problem from first principles and identify changes in ground emissivity (i.e., the number of photons reflected by the surface of the Earth) centered around 3.9µm. This corresponds to a blackbody with peak temperature of around 400K. In remote or urban areas, there are no variable heat emitters that could introduce transient sources in this range. Thus we are confident that our detection strategy has the potential to limit false positive triggers and contribute to a robust wildfire alert system.

As it stands, the implemented workflow also provides an improved localisation of the wildfire ignition hotspots. This is a crucial element in early response as the current 2 km positional uncertainty could lead to fire teams deploying on the wrong side of a mountain ridge, or on the opposite bank of a river. An accurate localisation can then influence the deployment plans of first responders.

Next Steps and Recommendations

The limiting factor to our research sprint was time. We had a very productive sprint and produced a successful proof-of-concept detection system. However, further work is required to automate each of the steps in the workflow, including:

- automating the detection and masking of clouds when selecting individual images that feed into the stacked reference image;
- developing and testing an automated workflow to create a reliable reference image, possibly from archive data, when images from preceding days are unusable due to clouds or other issues; and
- automating the detection and localisation of the ignition hot spot within the resulting subtraction image.

Following these developments, we would test and validate the detection pipeline by running it on historical data, and compare the results to ground-truth from previous fire seasons. This would provide an accurate characterisation of the precision and recall of our solution and provide a benchmark for converting our approach into a live tool. We note that experimental test-burns would be particularly useful for validation and testing, as ignition time and rate-of-spread are usually recorded in detail.

In more detail, the additional work we propose is:

1. ***Automatically detect and segment clouds within an image*:**
   Current and previous teams from the Frontier Development Lab network have already tackled this technical challenge (e.g., https://arxiv.org/abs/1911.04227), in addition to recently published research (e.g., https://doi.org/10.3389/fenvs.2019.00020). We estimate that this effort would require 2 weeks to incorporate within a stable framework.
2. ***Identify an appropriate detection threshold to investigate*:**
   This is a crucial step and our team is uniquely positioned to deliver on this task. The Scion Rural Fire Research Team has performed a scientific experimental burn in 2020 which was not identified as an official hotspot by the current Himawari-8 detection algorithm. We have the exact coordinates of this burn and we can calibrate our detection threshold so that this fire is identified. This step would take around 1 week.
3. ***Create a service architecture to simulate the 2019 fire season*:**
   In order to validate our novel detection algorithm we would need to scale our solution. This would require a full season of data and the development of a new computing architecture. We are an experienced team and can deliver on this task as well. DUG Technology has already offered additional computational support and we have also secured compute time on Swinburne's supercomputer, OzStar[1]. This is a significant task and would take 6 to 8 weeks to develop and run.
4. ***Quantify the precision and recall of each detection threshold*:**
   Finding a wildfire is important but it is just as important to send confident alerts. This is a pivotal component of the work as it comments on the opportunity cost incurred by end users. This analysis would follow on the previous work and would take approximately 2 to 4 weeks, depending on the level of visualisation that will be required.
5. ***Document and publish results*:**
   Finishing the technical work would in itself be a crucial milestone, but additional time would be required to fully document the work and also produce a final report. This would take an additional 2 to 4 weeks depending on the level of the final report (i.e., academic paper vs. internal document).

In closing, for an additional 12 to 19 weeks of work we could deliver a fully-functioning remote fire detection system that can provide early alerts and continue mapping hotspots as they move across the landscape.

References

Agustsson, E., & Timofte, R. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops,* 2017, pp. 126-135.

Amici, S.; Wooster, M.J.; Piscini, A. Multi-Resolution spectral analysis of wildfire potassium emission signatures using laboratory, airborne and spaceborne remote sensing. Remote Sens. Environ. **2011**, 115, 1811–1823.

Bolshakov, A. and Barnes, R., 1997.Electron temperature and radiative attachment continua in enclosed

[1] https://supercomputing.swin.edu.au/

inductively coupled plasma in argon and chlorine. Spectrochimica Acta, Part B (52), 2127-2150.

Lin, L.; Meng, Y.; Yue, A.; Yuan, Y.; Liu, X.; Chen, J.; Zhang, M.; Chen, J. A spatio-temporal model for forest fire detection using HJ-IRS satellite data. Remote Sens. 2016, 8, 403.

Pennypacker, C. R., M. K. Jakubowski, M. Kelly, M. Lampton, C. Schmidt, S. Stephens, and R. Tripp, 2013: FUEGO - Fire urgency estimator in geosynchronous orbit - A proposed early-warning fire detection system. Remote Sens., 5, 5173–5192, doi:10.3390/rs5105173.

Ross, A.N.; Wooster, M.J.; Boesch, H.; Parker, R. First satellite measurements of carbon dioxide and methane emission ratios in wildfire plumes. Geophys. Res. Lett. 2013, 40, 4098–4102.

Xie, Z., W. Song, R. Ba, X. Li, and L. Xia, 2018: A spatiotemporal contextual model for forest fire detection using Himawari-8 satellite data. Remote Sens., 10, 1–24, doi:10.3390/rs10121992.

Xu, G., and X. Zhong, 2017: Real-time wildfire detection and tracking in Australia using geostationary satellite: Himawari-8. Remote Sens. Lett., 8, 1052–1061, doi:10.1080/2150704X.2017.1350303.

B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee, "*Enhanced deep residual networks for single image super-resolution*" in Proceedings of the IEEE conference on computer vision and pattern recognition workshops, 2017, pp. 136–144.

Appendix: Technical Requirements

Python libraries used:
pandas, gdal, pyproj, numpy, cv2, fiona, afine, matplotlib

SIGNATURES OF EXTREME FIRE

BUSHFIRE DATA QUEST 2020

Technical Memorandum

Predicting extreme fire behaviour from smoke plumes

Todd Ellis (PhD Student, University of Tasmania)

Tom McCavana (Serenitec)

Stephane Mangeon (CSIRO)

Anna Matala (University of Tasmania)

Science and Data Leads

Dale Hamilton (Northwest Nazarene University)

Marta Yebra (Fenner School of Environment & Society, Australian National University, Canberra)

Abstract / Executive Summary

Extreme fire behaviour turns the bushfire from bad to catastrophic. They are unpredictable and cause more damage than a non-dynamic but fierce bushfire. Their occurrence cannot be predicted accurately until very close to the event which puts the firemen in danger. An experienced fire-fighter can see the beginning of a pyroCB event from the change in smoke plume colour and texture; would it be possible to analyse the plume features to detect these changes before they are visible to an eye? We use plume and haze masking to identify the smoke, and use methods of unsupervised machine learning to analyse the plume features of Gosper Mountain Fires in 2019-2020. We hoped to see clustering of the results, clearly marking the difference between a normal fire day and a pyroCB event. By using Himawari-8 data from all bands, we observed some visual changes in the plume as well as small differences in the principle components between different kinds of days. We were able to identify the frequency bands that marked the clear difference in case of an extreme fire event. One of the challenges was the spatial resolution of Himawari-8 data which was not ideal for detecting relatively small size events such as pyroCB's. This can be improved for the future either by using different satellite data or selecting the Himawari-8 bands with higher spatial resolution. We are also planning to improve the data set by adding weather data, as well as adding temporal gradients to smooth the data. For better analysis in the future we are considering using object detection and image segmentation.

Introduction

Fires with extreme - or dynamic - fire behaviour contribute to disproportionate damage statistics. These fires are uncontrollable and unpredictable, and can often lead to loss of lives (Filkov et al. 2020). An experienced firefighter can observe the change in fire behaviour by the smoke plume colour and texture. Unfortunately the observations are not visible until very close to the actual change, or they occur in the upper layer of the plume while the smoke closer to ground covers the vision.

In the present work we investigate whether satellite data combined with machine learning could simulate the observations of an experienced firefighter with the benefit of seeing above the smoke. If we were able to identify the features of smoke plumes during dynamic fire events, could we use this to take a few steps back in the history and predict dynamic fire changes in advance?

In this proof-of-concept we concentrate purely on smoke plumes and weather based dynamic fire behaviour (e.g. pyroCB's). Other factors such as fuel and topography play an important role in some other types of extreme events and they should be taken into account in the future research. This study concentrates on one of the largest bushfires in Australia, the Gospers mountain fire during summer 2019-2020. Several pyroCB's were detected during that fire which makes it an ideal target for our study.

Identified Needs and Opportunities

With fire seasons lasting longer and causing more damage than ever before there is a growing need to automate the detection and monitoring of wildfire behaviours. With the formation of the Australian space agency in 2018, we look to the next generation of satellites (currently under review?) and indeed their instrumentation to enable this. There is a strong need to identify spectral features or signatures that would indicate changes in fire behaviour and mark extreme events.

Satellites often face resolution limitations either spectrally, spatially or temporally. When identifying a satellite data

We identified an opportunity to use smoke plumes as a proxy for fire behaviour. The goal being to detect and ideally predict events through subtle changes in the spectral features so that response coordinators and commanders can most effectively manage the resources at their disposal. Most of the current satellites available for monitoring fires have good spatial resolution, but long revisit times. Smoke plumes are often overlooked when trying to monitor a fire.

Data description

We use data from the Himawari-8 satellite in the region of the Gosper mountain fire as well as regions of south coast NSW (the bounding box between longitudes 148.16 to 152.00 and latitudes -30.81 to -37.41) from the beginning of November 2019 until the end of February 2020. This time period covers almost the entire duration of the Gosper mountain fire and should include most if not all dynamic fire behaviours and extreme events.

Himawari-8 data enables an excellent temporal resolution of 10 minute which is critical for observing rapid changes in the fire behaviour. Note that we take hourly samples of the data for the entire period and then finer samples of 30, 20 and 10 minute samples around certain extreme fire events. To help detect plumes we chose to concentrate on daylight hours, roughly between 8 am and 6 pm of local time, which corresponds to 9 pm and 7 am of GTM (satellite time).

We used all 16 spectral bands of the Himawari data, however this meant that we were limited to 2 km spatial resolution. We found this reasonable for the proof-of-concept, but ideally should be improved for the final application. Given our objective of detecting changes in behaviour through smoke plumes, we felt that temporal and spectral resolution were more important than spatial.

Data Gaps

According to our PCA (results explained in detail later) the most significant bands of Himawari-8 for analysing smoke plumes are the visible ones and near-IR. However, these bands are significantly affected by the time of day. Coincidentally, it was noticed that a big part of the reported pyroCB events occurred during the early hours of a day when it was still dark. This is something that should be taken into account in the future studies.

Another important factor is the spatial resolution. The size of a pyroCB plume is not large and the 2 km resolution of Himawari-8 does not provide enough pixels for fully characterising the smoke plume features.

Data Exploration

We expected to see some significant change in the smoke plume colour and texture and therefore we started the process by implementing true colour images and testing different plume and haze masks for the data. We used the ideas of Qin et al. 2019 and Shang et al. 2017 when implementing the masking.

The greatest discovery we made during the data exploration was that by multiplying two cloud indexes we can highlight the plume even better. Our final haze masking uses bands 1 (blue) and 2 (green), corrected with the temperature (band 14). Multiplying the blue and green mask, and setting limits to the range that is included we get the final mask we use. The equations are shown below.

```
cloud_g = ((373.15 - B014) / 100) * B02

cloud_b = ((373.15 - B014) / 100) * B01

plume = cloud_g * cloud_b

plume_02_07_mask = (plume >= 0.02) & (plume <= 0.07)
```

Methodology

Workflow overview

The high level workflow with different options is shown in Figure 1. The work can be summarised to three main tasks:

1. Data processing,
2. Application of ML methods, and
3. Visualisation of the results.

The details of each step are discussed further in the following subsections.

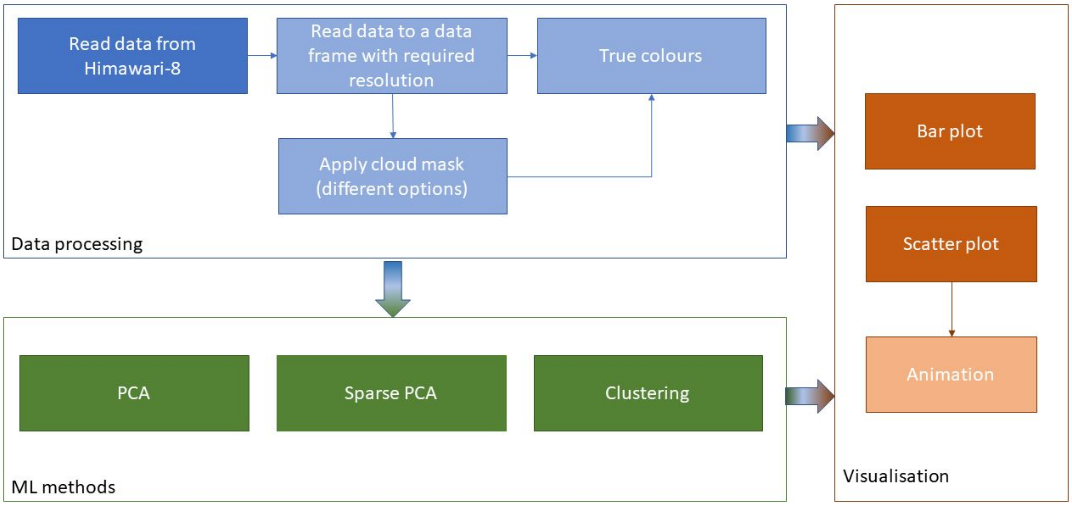

Figure 1. *Overview of the workflow of this project.*

Data processing

We created a flexible data processing pipeline for Himawari-8 in order to curate data sets that focus on a given region, times and set of spectral bands. The pipeline allows you to select a set of bands, spatial coordinates and time sampling pattern and consolidates the data into a single netcdf file. The time sampling pattern allowed us to, for example, take hourly data for the entire fire season, but subsample to 10 minute intervals around events of extreme fire behaviour.

The basic steps in the pipeline are as follows…

1. Get the data
   a. Pull the data from the NCI storage facility
   b. Check for gaps or corrupt files. Try pulling the data again.
   c. Remove missing files and leave them as gaps in the data.
2. Open a netcdf file for writing the output
   a. Transform the desired coordinates to the Geostationary reference system.
   b. Project to the nearest coordinates in the image and use it as a template for cropping.
   c. Create the netcdf data group ready for writing and copy any metadata (Coordinate Reference System, Units etc)
3. Iterate through the data cropping and stacking into the output file.

The pipeline used parallel processing as much as possible and allowed us to consolidate the Himawari 8 data for the Gosper's Mountain fire in the 2019-2020 fire season from 1.4TB of data in 61,000 files into a single stacked data cube in a single 10GB file.

One thing to note is that we decided not to reproject the satellite images into a human friendly coordinate reference system as we were concerned about the potential artifacts this might create. added.

The pipeline is in a fairly usable state (however not fully productionised or supported) and can be found in our gitlab repository.

Plume identification

Important part of the pre-processing was applying cloud and haze masks. Unlike in normal satellite data processing we only wanted to pick the pixels of the smoke plumes. The cloud masking process is explained in more detail in section Data Exploration.

Application of the ML methods

Given the brevity of the Quest, and the exploratory nature of our work investigating the ability to predict fire behaviour based on smoke signature. Therefore, we focus our Machine Learning

workflow on unsupervised learning. More specifically, we apply Principal Component Analysis (PCA) and Sparse-PCA to identify the most important bands to the spectral characterisation of haze from fires. PCA is a well-established method to reduce data to a lower dimensionality, Sparse PCA expands on the method by penalising the number of input channels used (Wall et al., 2003). See our results for an illustration of this. One pathway for predictions is through clustering analysis. To illustrate this, we investigate the applicability of Density-Based spatial clustering (see Schubert et al., 2017). This particular approach to clustering focuses on identifying areas of high density separated by low densities, a more appropriate approach when dealing with highly-connected clusters, seemingly the case with satellite observations.
All ML methods were applied using the scikit-learn python library (https://scikit-learn.org).

In parallel, the team from Northwest Nazarene University explored the use of an entropy-based decision tree (ID3), to identify which spectral band held most information. The ID3 algorithm is a supervised classification algorithm which utilizes Information Gain to recursively build a decision tree, by identifying the band with the highest Information Gain using a set of training pixels labeled as Plume or Not Plume (e.g. surface or cloud). Subtrees are then recursively created from subsets of the training pixels partitioned on the band value split point for which the band was found to have the highest Information Gain. Bands that are more useful for classifying Plume vs Not Plume will appear at the top of the decision tree, closest to the root. Bands that are not as useful will be found further down the tree, farther from the root. In the resulting Decision Tree, Himawari-8 bands 1 and 10 were found to have the highest Information Gain, followed by bands 8, 5, 16, 6, 2 and 4. While the steps followed allow us to better understand the data and to select which spectral bands will contain the most information about fires and their smoke, they are not currently predictive. However, cross validation accuracy exceeding 75% from very preliminary training sets shows promise for being able to use a decision tree to classify smoke plumes from Himawari-8 data.

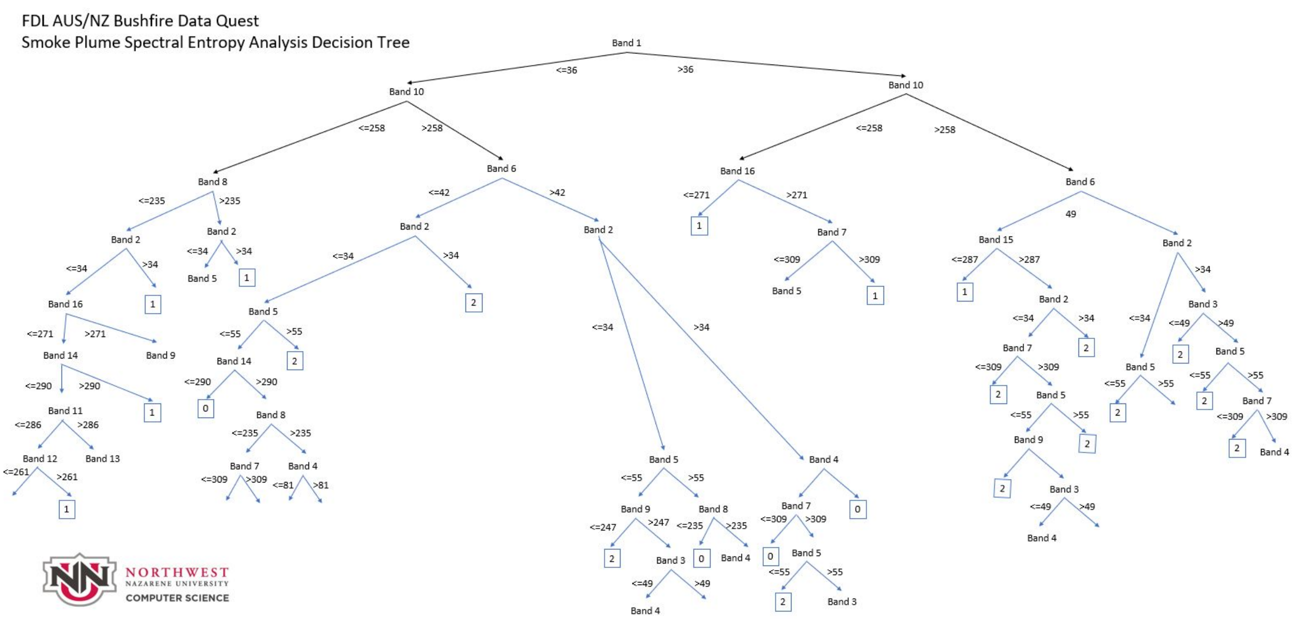

Visualisation

Visualisation is not only for producing nice colourful images but to evaluate the results. We hoped to see clusters forming during the extreme fire events to separate them from the normal (bad) fire behaviour. We did visualisation in three parts: (a) True colour and plume mask animations directly from the satellite images, (b) scatter images and animations of PCA and sparse PCA results (animations showing clustering as the time changes), and (c) barplots of the principal components.

Asides – Paths not Taken

We also explored density-based clustering to try and identify subsets of fire behaviours within the himawari data. Due to time constraints we did not get to explore this fully, however there were some promising initial results.

Our approach was to try and validate clusters by comparing them with PyroCB events. We used two subsets of data

1. The first was all plume pixels across the whole season as identified by our plume mask.
2. A subset of pixels centred on known PyroCB events for 2 days leading up to the event and 1 day after.

As feature sets for the unsupervised learning we took all 16 bands of the himawari 8 data along with the time of day. In some cases we tried to incorporate the plume spatial density by a couple of different approaches. All features were normalised.

We looked at minimum group sizes equal to the number of events and a variety of different distance metrics readily available in the sklearn library. We found that epsilon values (an algorithm parameter that characterises the distance between related points) between 0.01 and 0.2 tended to give healthy cluster populations to examine.

The biggest problem we faced with this approach was validating or identifying interesting groups. The best data we had for validating this in an automated way was PyroCB event information. However, this still seemed to lack the accuracy required for a detailed comparison (accuracy in time and position, the actual certainty of the occurrence and the fact that many of the events occurred during the night or during twilight).

Our approach to validating clusters was to firstly exclude groups that 1) related spatially to pixels and 2) occurred at a single epoch/day (i.e. we need the group to have members over several times/fire events). We then looked at groups with the lowest time offset from the PyroCB events and then those with most consistent offsets from events in case there was a pre or post signature for PyroCB events.

The preliminary results from the density-based clustering indicated that indeed there were groups that corresponded to events on particular days and circumstances. However, there were none that consistently matched our catalogue of PyroCB events. As mentioned above, our search was not exhaustive and this approach would serve well from a data exploration point of view (given some more time).

Results and Discussion

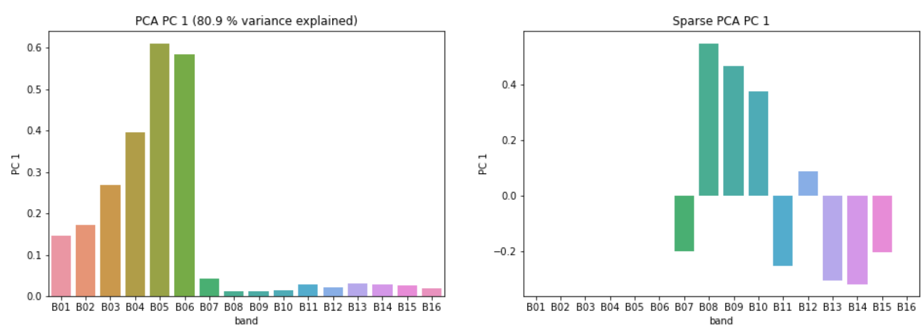

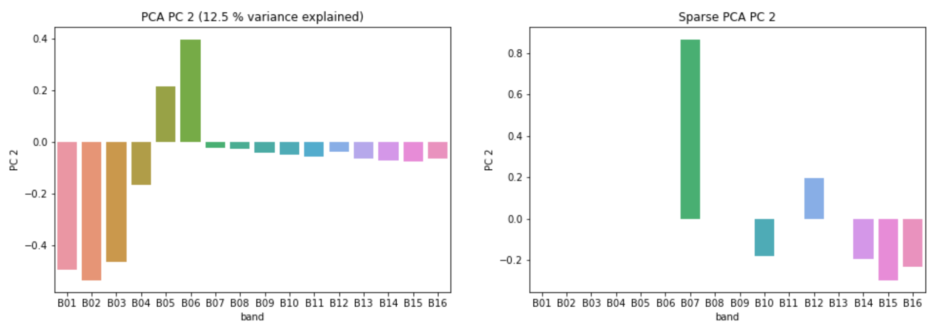

Figure 3. *The Principal Component Analysis (PCA) Himawari spectral band components, applied onto all pixels within the Gosper Mountains fires that were identified by our plume mask. The standard PCA (right) contrasts with the Sparse-PCA (left), which aims to minimize the dimensionality of each component.*

Figure 3 shows the components associated with each spectral band. The data consists of hourly observation for all dates from 1/11/2019 until 29/2/2020. While the visible and short-IR bands (bands 1 to 7) dominate the standard PCA, thermal IR dominates the Sparse PCA. The analysis included both cloud-masked (left) and all (right) data but interestingly the results were similar for sparse PCA. This suggests that the sparse PCA implementation ignores the spectral variariance of pixels we identified as plume.

The first component (PC1) explains the vast majority of the variance in the data, with 80.9 and 86.1% respectively for masked and all pixels. The second component (PC2) explains about 12% of the remaining variance. The general trend is similar, with PC1 focusing on the visible and short-IR bands (bands 1 to 6), all with positive components. In contrast PC2 distinguishes between the short-IR (bands 5 and 6), and the visible and near-IR (bands 1 to 4). The thermal IR seems to hold less information, although PC2 seems to put more emphasis on the far-IR (bands 14 to 16) than PC1. The Sparse PCA highlights a different picture, with the first 6 bands reduced to 0, which suggests the information they contain is highly correlated with other bands. Interestingly the Sparse PC1 distinguishes between mid and far-IR (bands 8 to 10 and 13 to 16 in particular).

Our PCA supports the idea that when it comes to fire smoke (masked pixels), more information is contained in the Infrared channels than in the visible. This echoes the findings of many previous research papers (e.g. Sowden and Blake, 2020)

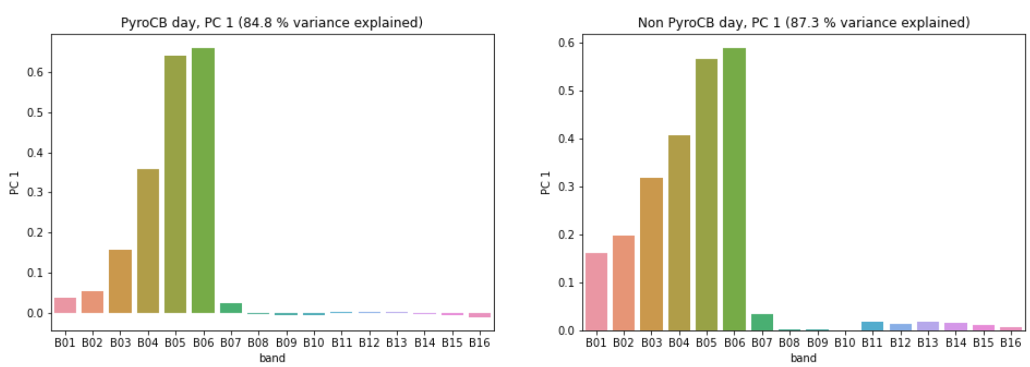

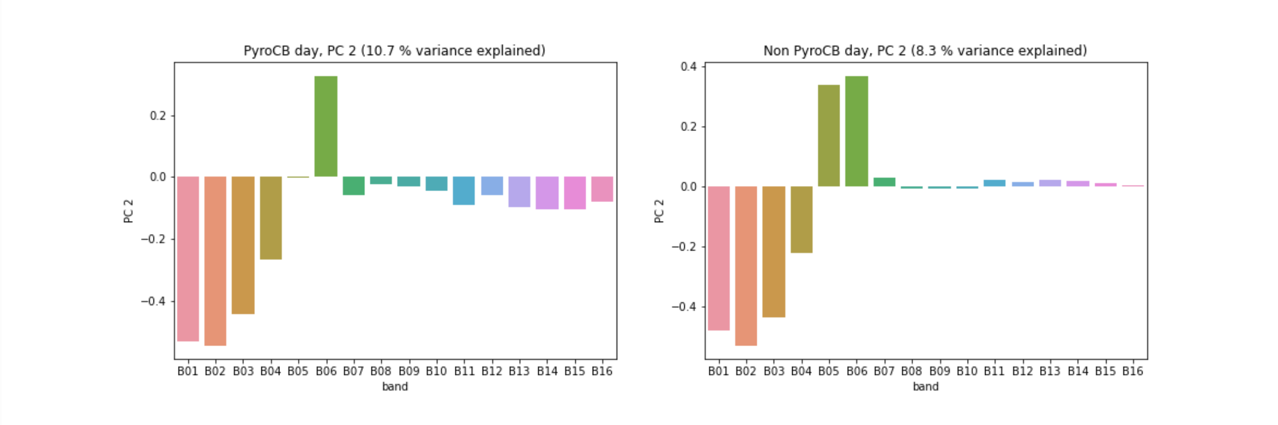

Figure 4. *Comparison of principal components for a pyroCB day (28.-29.12.2019) and non-pyroCB day (25.-26.12.2019) with 10 minute resolution.*

Figure 4 shows PCA for a pyroCB day (several confirmed pyroCB's in the area) and a normal bad fire day without dynamic fire behaviour. Both time periods are from December 2019 for ensuring that other conditions (e.g. daylight time) are approximately the same. For the principal component 1 (explaining about 85 % and 87 % of the variance for pyroCB day and non-pyroCB day, respectively), it is noteworthy that the visible bands (1-3) are less significant for pyroCB days than for the day without dynamic fire behaviour. For the PC 2, the difference can be observed in band 5 (for a pyroCB day not significant while for a non-pyroCB day very significant), and in the higher bands 11-16 that show more activity during pyroCB day than for a normal bad fire day.

The next three pages (Figure 5 a/b/c) shows a comprehensive visualization of the Gosper mountains fires on the 28th of December 2020, a known PyroCb day. We recommend visualizing the companion animation for a more representative and compelling assessment of the temporal changes in the spectral characteristics of the fires' smoke. More importantly, the animation shows the process between the snapshots we discuss below is gradual, which highlights the potential for predictions.
We will start by discussing the general picture provided by the true color images (along with our plume masks), before moving to the spectral signatures. Mid-IR bands, and band 7 in particular, are often used for fire detection (hotspots). We plot Band 7 as an indicator of hotspots, with colours ranging from low (blue) to high (red) temperatures, hotspots are particularly striking at 1pm (Figure 2 c). Meanwhile the differences between Mid-IR (Band 7 and 8 here) and Far-IR (Band 12 and 14 here) is commonly used for cloud and PyroCb identification.

At 11am (AEST), the fires' haze is blanketing the area, including Sydney. There is relatively low cloud cover. Half an hour past noon, the picture changes, with some clouds intruding from the South West, but also a lot of cloud activity around the area of interest (green). We posit these clouds may be fire-related and potential PyroCb events. At 1 pm, the cloud activity around the fires seems to have diminished, while the South-western clouds have progressed further into our area of interest.
The spectral visualization also shows a likely sign of PyroCb expulsion: pixels within the area of interest (green), move towards lower Far-IR temperatures while their Mid-IR remains similar. A similar separation occurs in our first principle component (PC1), with high values in the second (PC2). In contrast, the clouds that intrude at 1 pm have lower temperatures across Mid and Far IR.

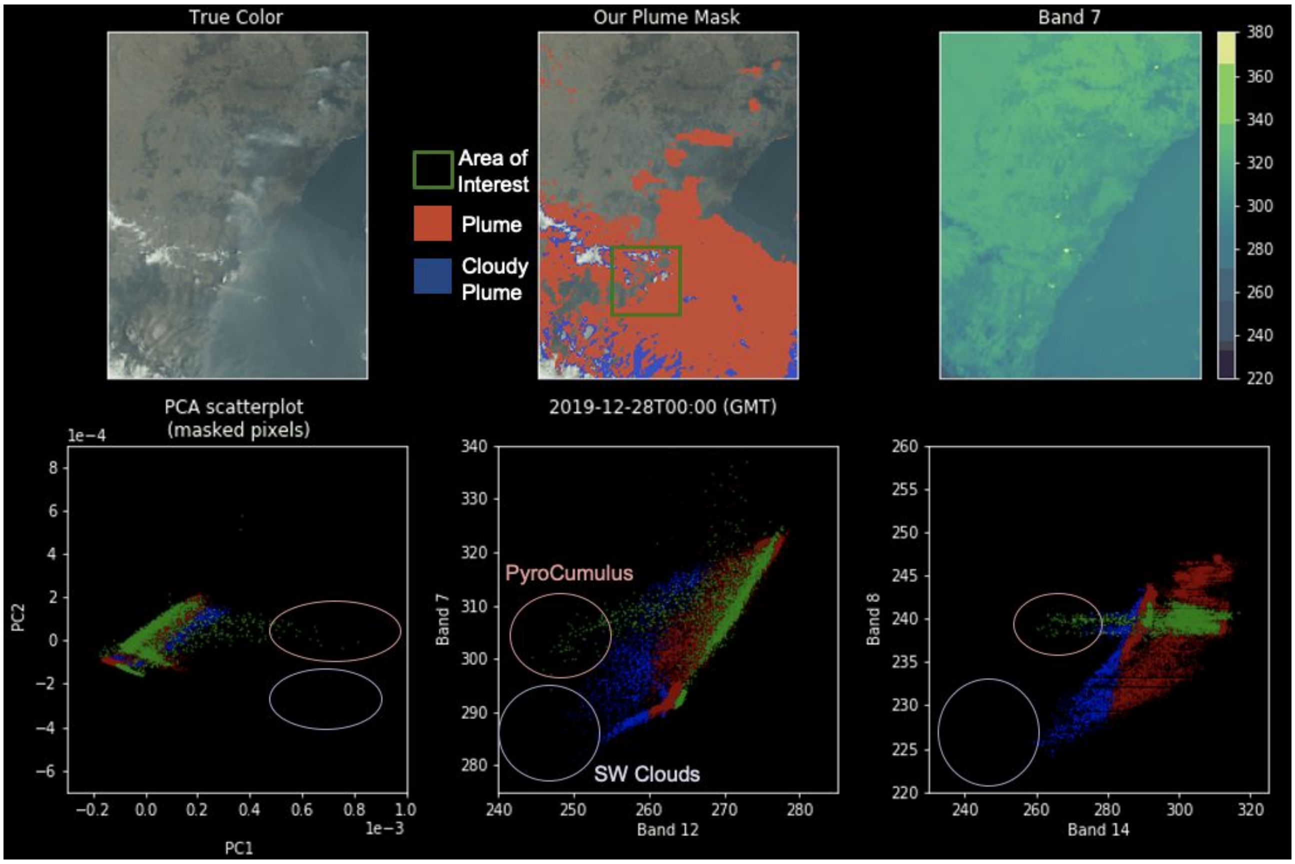

Figure 5 (a). *11am AEST on 28/12/2019. A comprehensive visualization of the Gosper mountains Fires. Blue and green shows pixels within our plume masks, while green shows all pixels within the bounding box around known PyroCb detections. Likely before a major PyroCb event.*

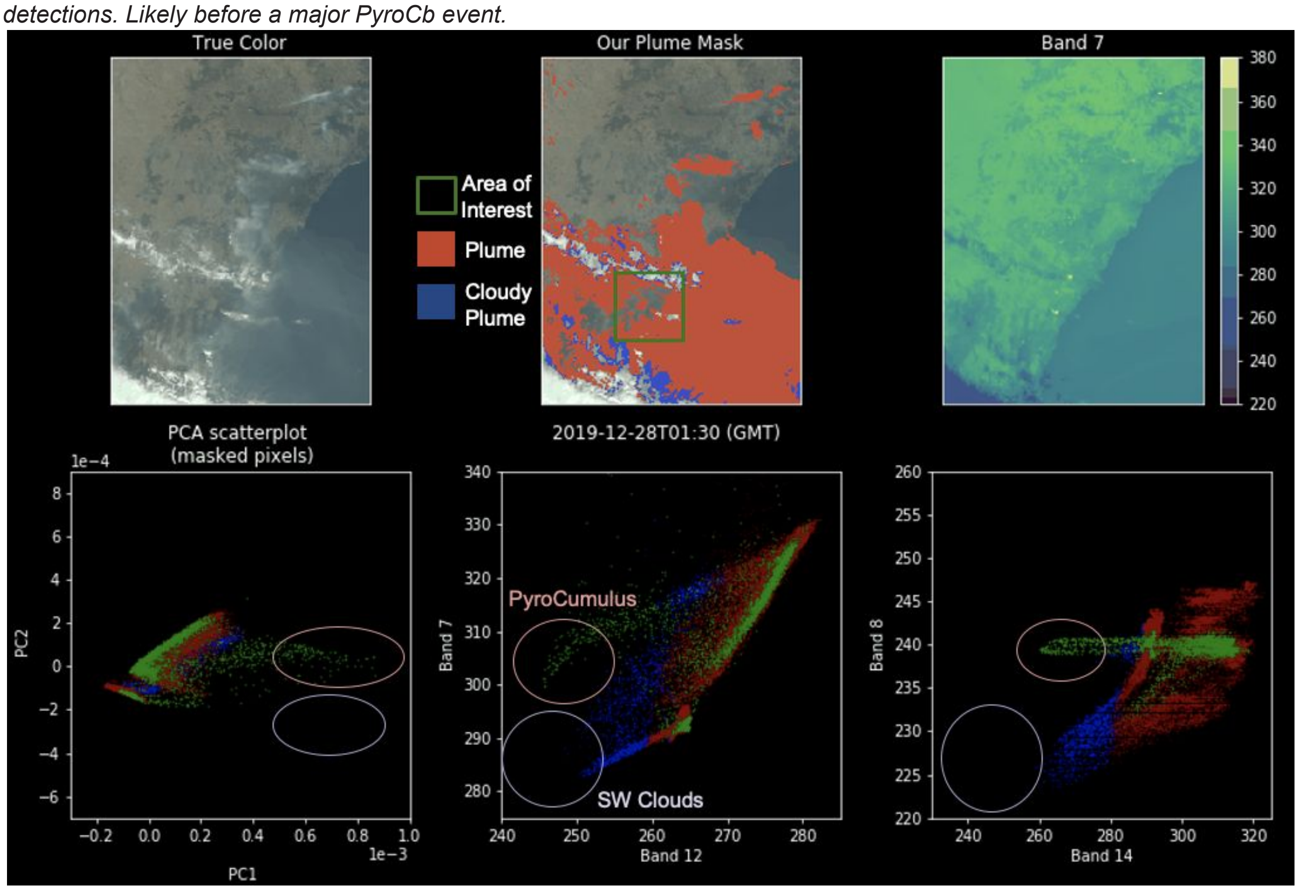

Figure 2 (b). *12.30pm AEST on 28/12/2019. A comprehensive visualization of the Gosper mountains Fires. Blue and green shows pixels within our plume masks, while green shows all pixels within the bounding box around known PyroCb detections. Likely as a major PyroCb event is occurring. Note how the PyroCb does not appear to be masked as smoke.*

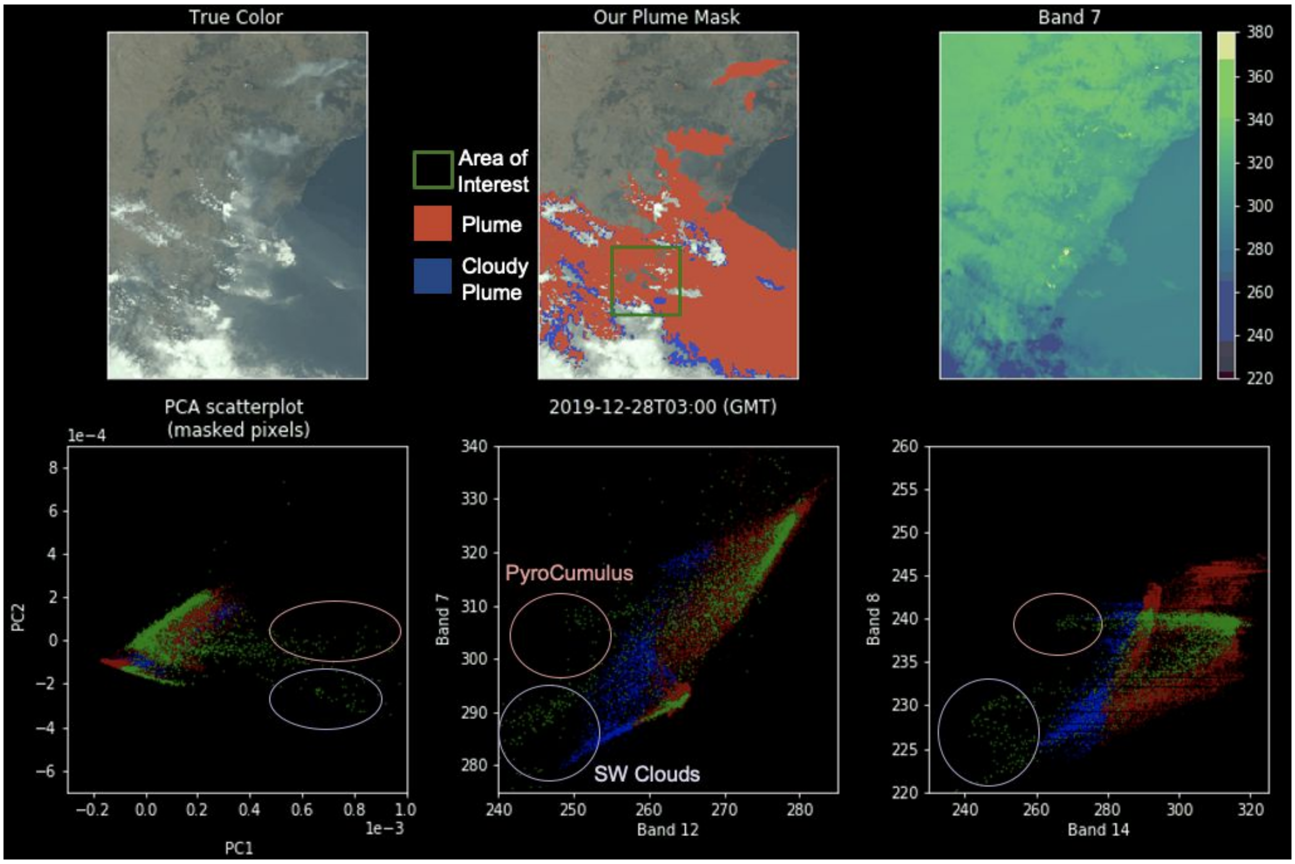

Figure 2 (c)*. 2pm AEST on 28/12/2019. A comprehensive visualization of the Gosper mountains Fires. Blue and green shows pixels within our plume masks, while green shows all pixels within the bounding box around known PyroCb detections. Likely after a major PyroCb event is occurring. Note the clouds coming from the south-west display a spectral signature distinct from previous snapshots (green points in the lower-left part of our band scatterplots, the bottom figures).*

Next Steps and Recommendations

If this work was to be taken further, we recommend the next steps could focus on:

1. Expanding our feature inputs by taking gradients for pixel both temporally as well as across our spectral bands. Temporal gradients would be good indicators for sudden changes in a single band measurement when it comes to expanding our unsupervised learning approach. Looking at gradients and indexes based on the combinations of bands would be more appropriate for a labelled/supervised learning set.
2. We made some attempts to incorporate spatial information as features into our clustering analysis (average localised density of smoke plumes), however were not able to correlate this with any fire events. Some of the himawari bands offer resolutions of 1km and even 500m in the case of band 3. It might be useful to interpolated a smoothed model of the 2km bands at lower resolutions.
3. This analysis would benefit greatly from weather data.

Object detection and image segmentation: object detection is a popular and fast-advancing field of ML research and applications. These methods could be used to identify specific cloud signatures, rather than spectral signature. Unlike our approach, this would take into account the size, shape and orientation of smoke plumes. While many of these algorithms can be taken off the shelf, they are commonly trained on more human topics: self-driving cars, or automated image tagging. Therefore we would estimate a few weeks to a few months of research would be necessary to transfer these algorithms to fire smoke identification and characterization, and behaviour prediction.

References

Filkov, A. et al. 2020. Frequency of Dynamic Fire Behaviours in Australian Forest Environments. Fire 2020, 3(1), 1; https://doi.org/10.3390/fire3010001

“DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. Schubert, E., Sander, J., Ester, M., Kriegel, H. P., & Xu, X. (2017). In ACM Transactions on Database Systems (TODS), 42(3), 19.

Qin, Y. et al 2019. (2019) Cloud Cover in the Australian Region: Development and Validation of a Cloud Masking, Classification and Optical Depth Retrieval Algorithm for the Advanced Himawari Imager. Front. Environ. Sci. 7:20. doi: 10.3389/fenvs.2019.00020

Shang, H. et al. (2017), Development of a daytime cloud and haze detection algorithm for Himawari-8 satellite measurements over central and eastern China, J. Geophys. Res. Atmos., 122, 3528–3543, doi:10.1002/2016JD025659.

Sowden, M., and D. Blake. "Which dual-band infrared indices are optimum for identifying aerosol compositional change using Himawari-8 data?." Atmospheric Environment (2020): 117620.

Wall M.E., Rechtsteiner A., Rocha L.M. (2003) Singular Value Decomposition and Principal Component Analysis. In: Berrar D.P., Dubitzky W., Granzow M. (eds) A Practical Approach to Microarray Data Analysis. Springer, Boston, MA. https://doi.org/10.1007/0-306-47815-3_5

Appendix: Technical Requirements

The following table includes technical requirements divided into different parts of the project. Note that many of the functions can be performed in several ways or using different programming languages such as R. We used a conda environment with python 3.6.

Project task	**Python libraries**	**Approximate size of results**	**Processing power required (in scale PC - cluster)**
Preprocessing data			
Data acquisition from Himawari	threading subprocess	1.4TB	Few CPUs. High IO
Making dataframe of specific resolution	numpy pandas xarray netcdf4 multiprocessing	10GB	10-20 cores again mostly IO bound.
Cloud masking	-		
True colours	opencv (cv2)		
Saving dataframes & other results	pickle		
General handling of data	numpy pandas xarray		
Simple ML techniques			

PCA	sklearn -decomposition -preprocessing		
Sparse PCA	sklearn -decomposition -preprocessing		
Clustering	sklearn -model_selection -preprocessing -cluster.DBscan		
Visualisation			
PCA scatter plot	matplotlib.pyplot matplotlib.patches mpl_toolkits.mplot3d		
PCA barplot	matplotlib.pyplot mpl_toolkits.mplot3d		
Animation	import imageio IPython.display		

NEXT STEPS: FROM SOLVABLE TO SOLVED

Data Quest results are at an early stage and require significant work to rise up the Technology Readiness Ladder to TRL 9 (deployment in the field). We hope to build on these promising outcomes in further research sprints, engaging new and existing partners.

Future research efforts will result in tools and methods that are ready for live trial pilot studies in partnership with users such as incident commanders in the RFS as well as land managers in charge of post-fire mediation efforts.

However, there is no silver bullet and much work needs to be done before the AI capabilities demonstrated in this document can be trusted tools for fire-fighters on the ground. In this section we imagine some use-cases for AI, discuss the requirements for effective validation and for a framework supporting full-spectrum situational awareness.

HI-RES FUEL MOISTURE MAPS

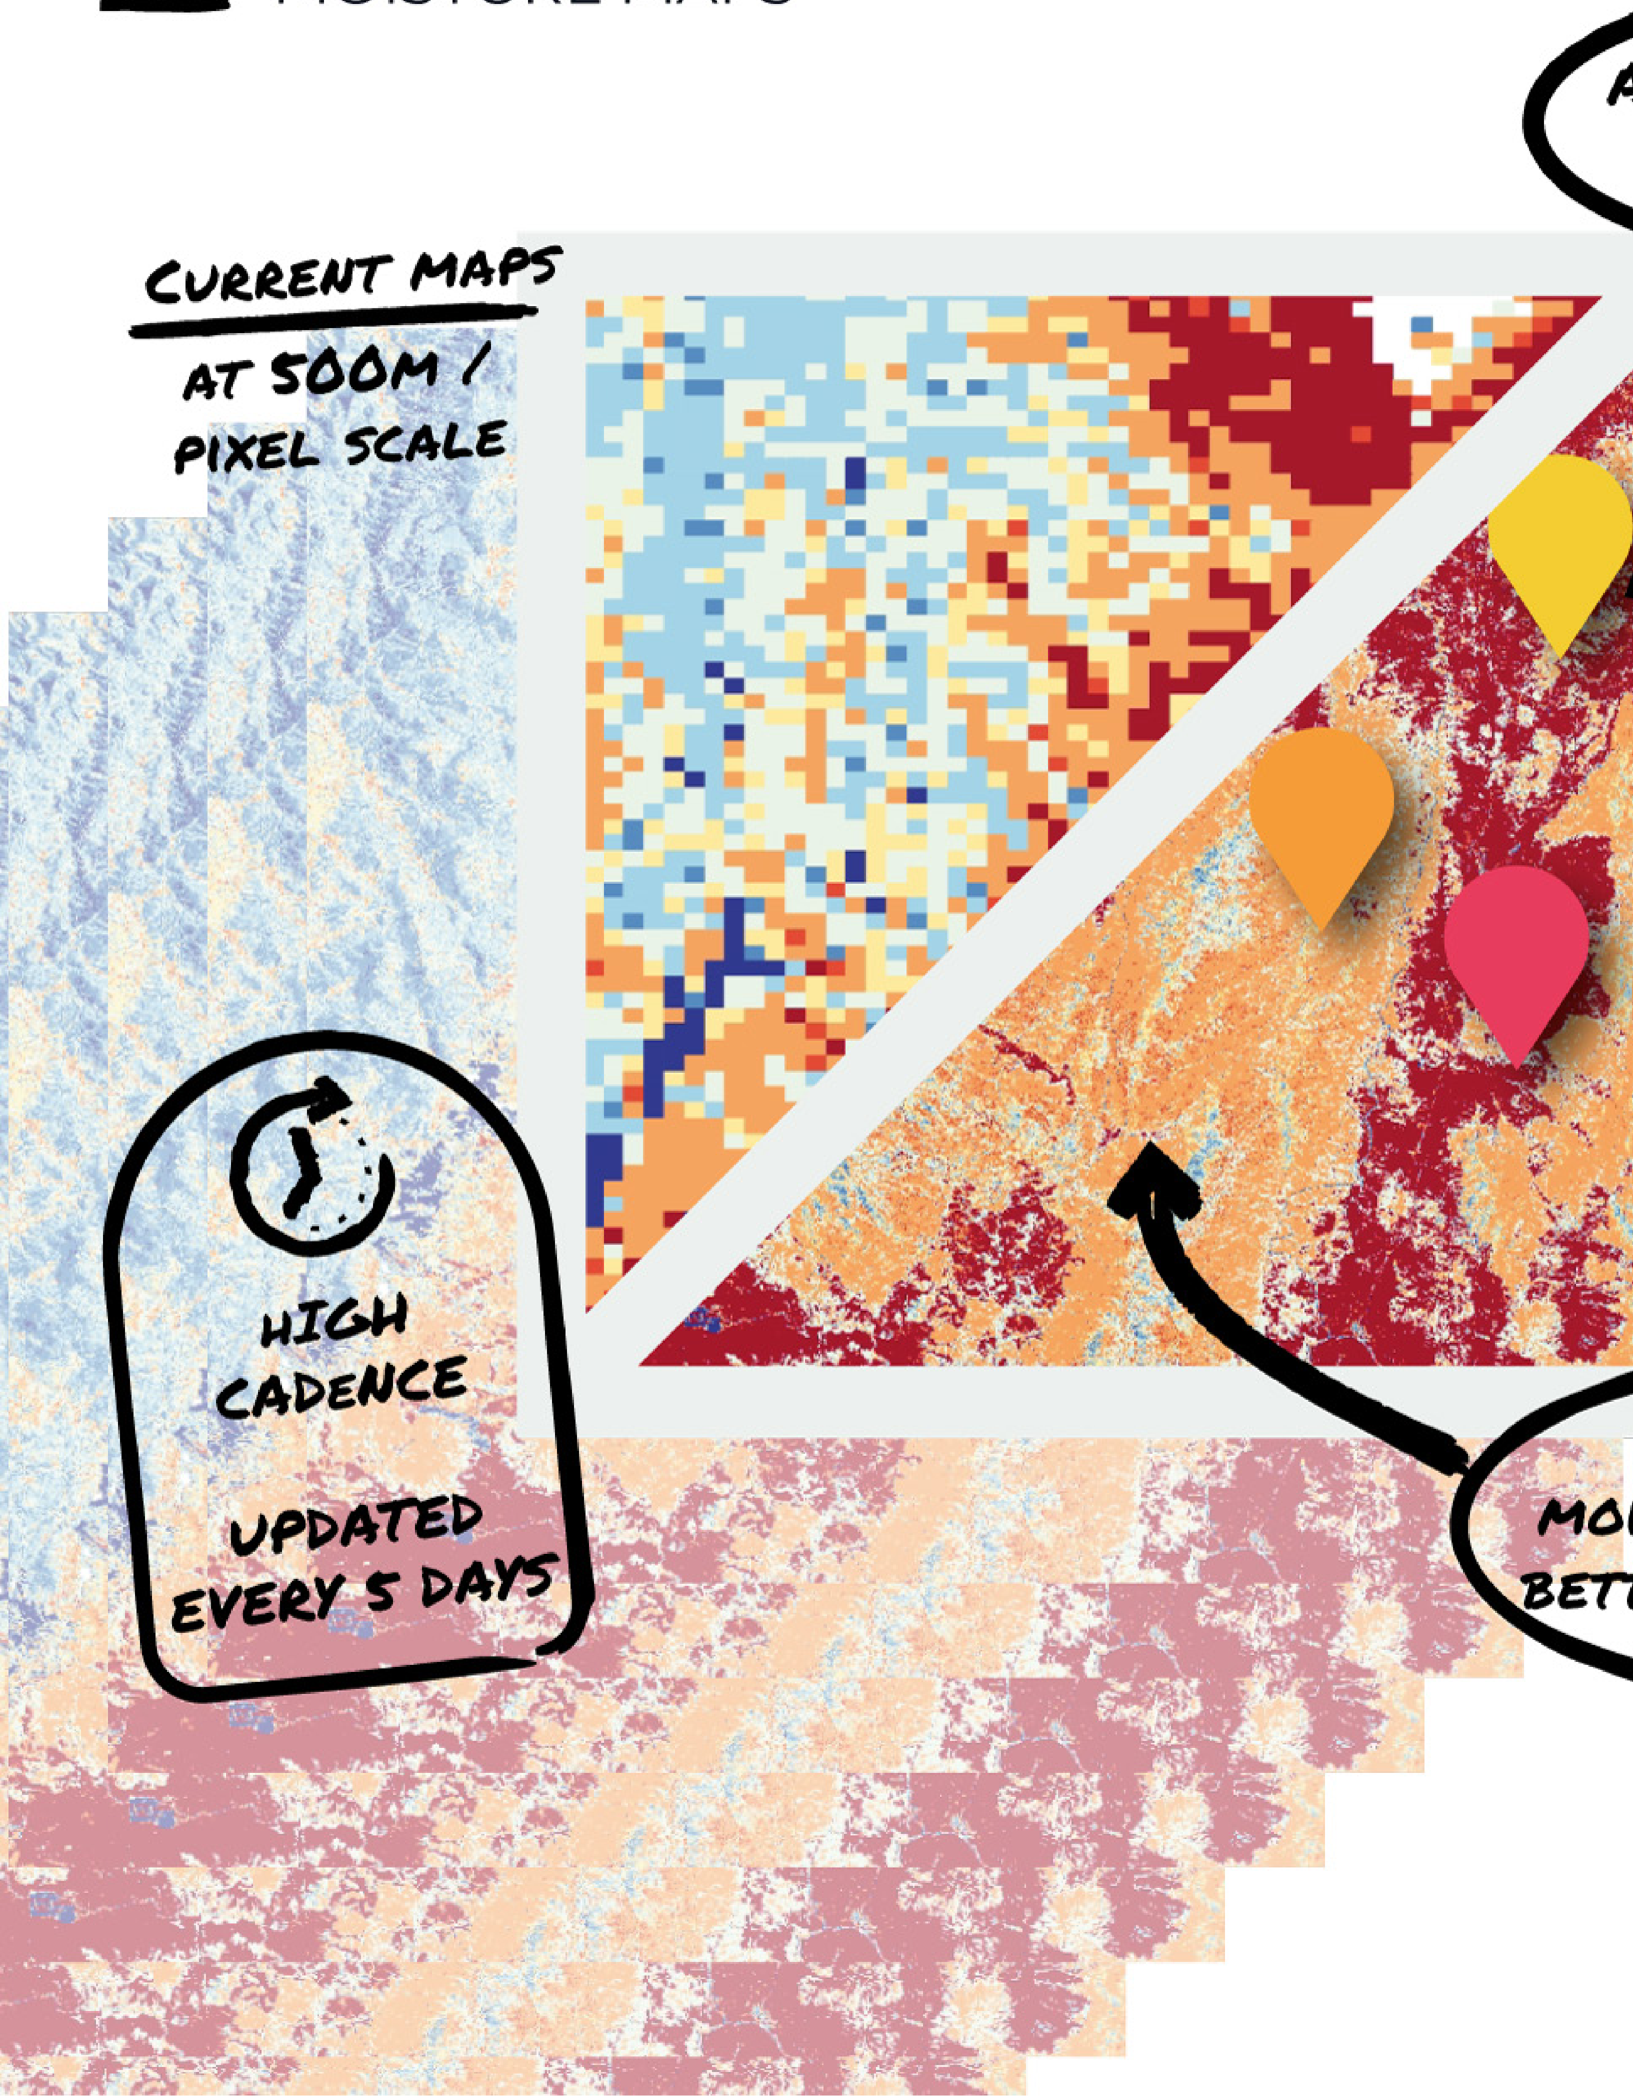

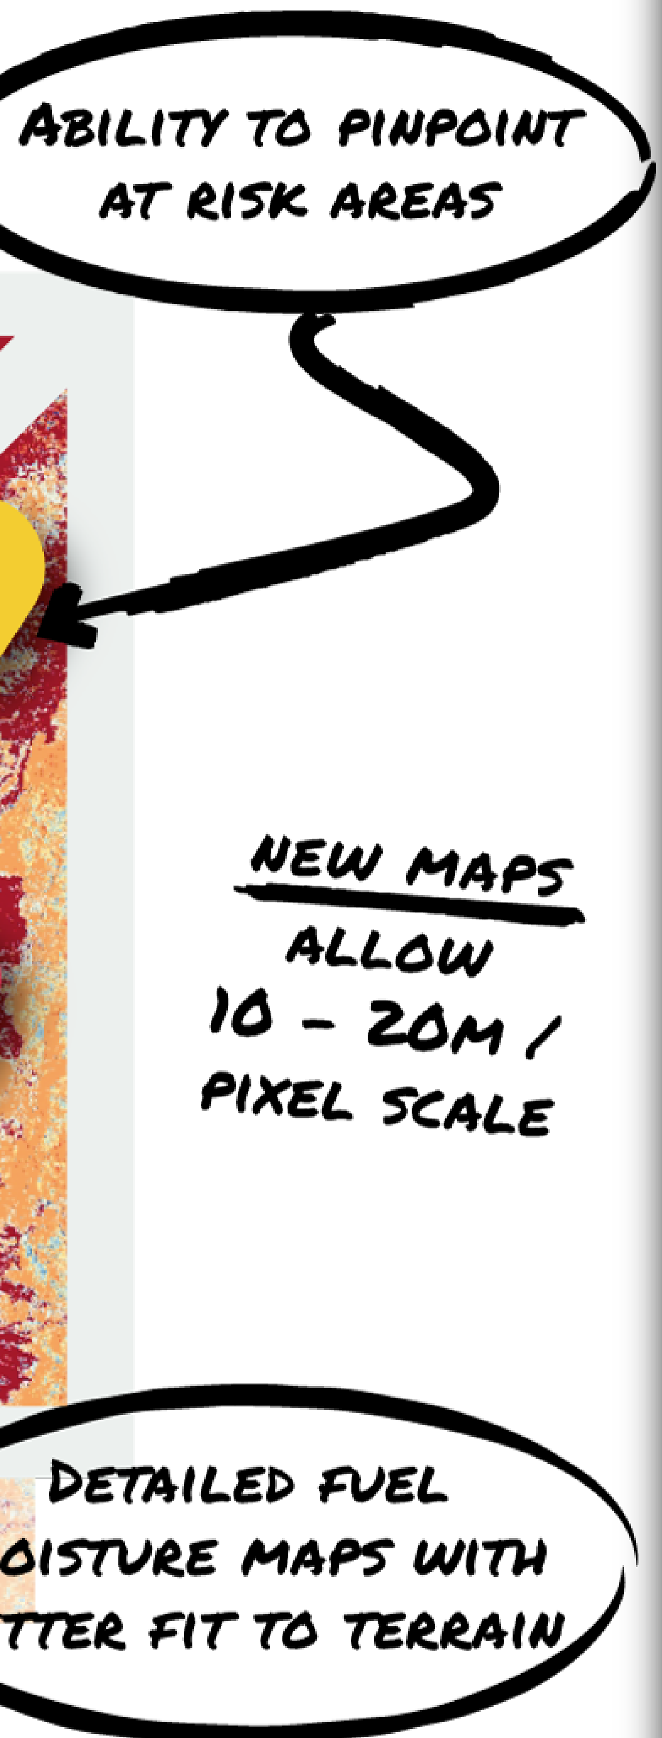

IMAGINE

A **citizen science** fuel moisture sensor kit sent to **schools and** and Scout / Guide groups **all around Australia** to improve the **ground truth** of the map.

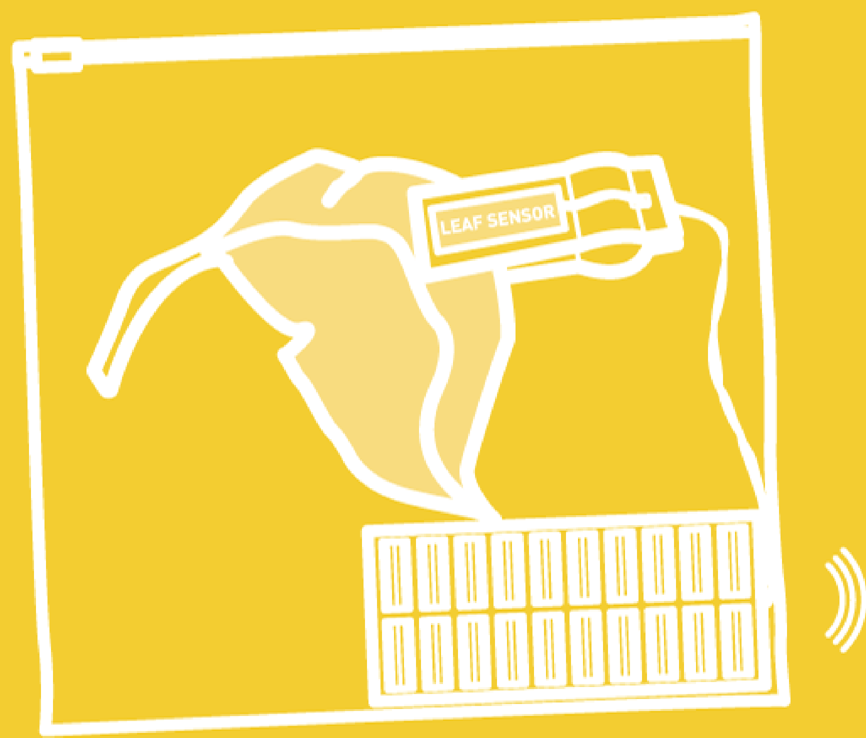

Collect leaves and **place in the sun** for two weeks - **sensor does the reporting.**

EARLY DETECTION

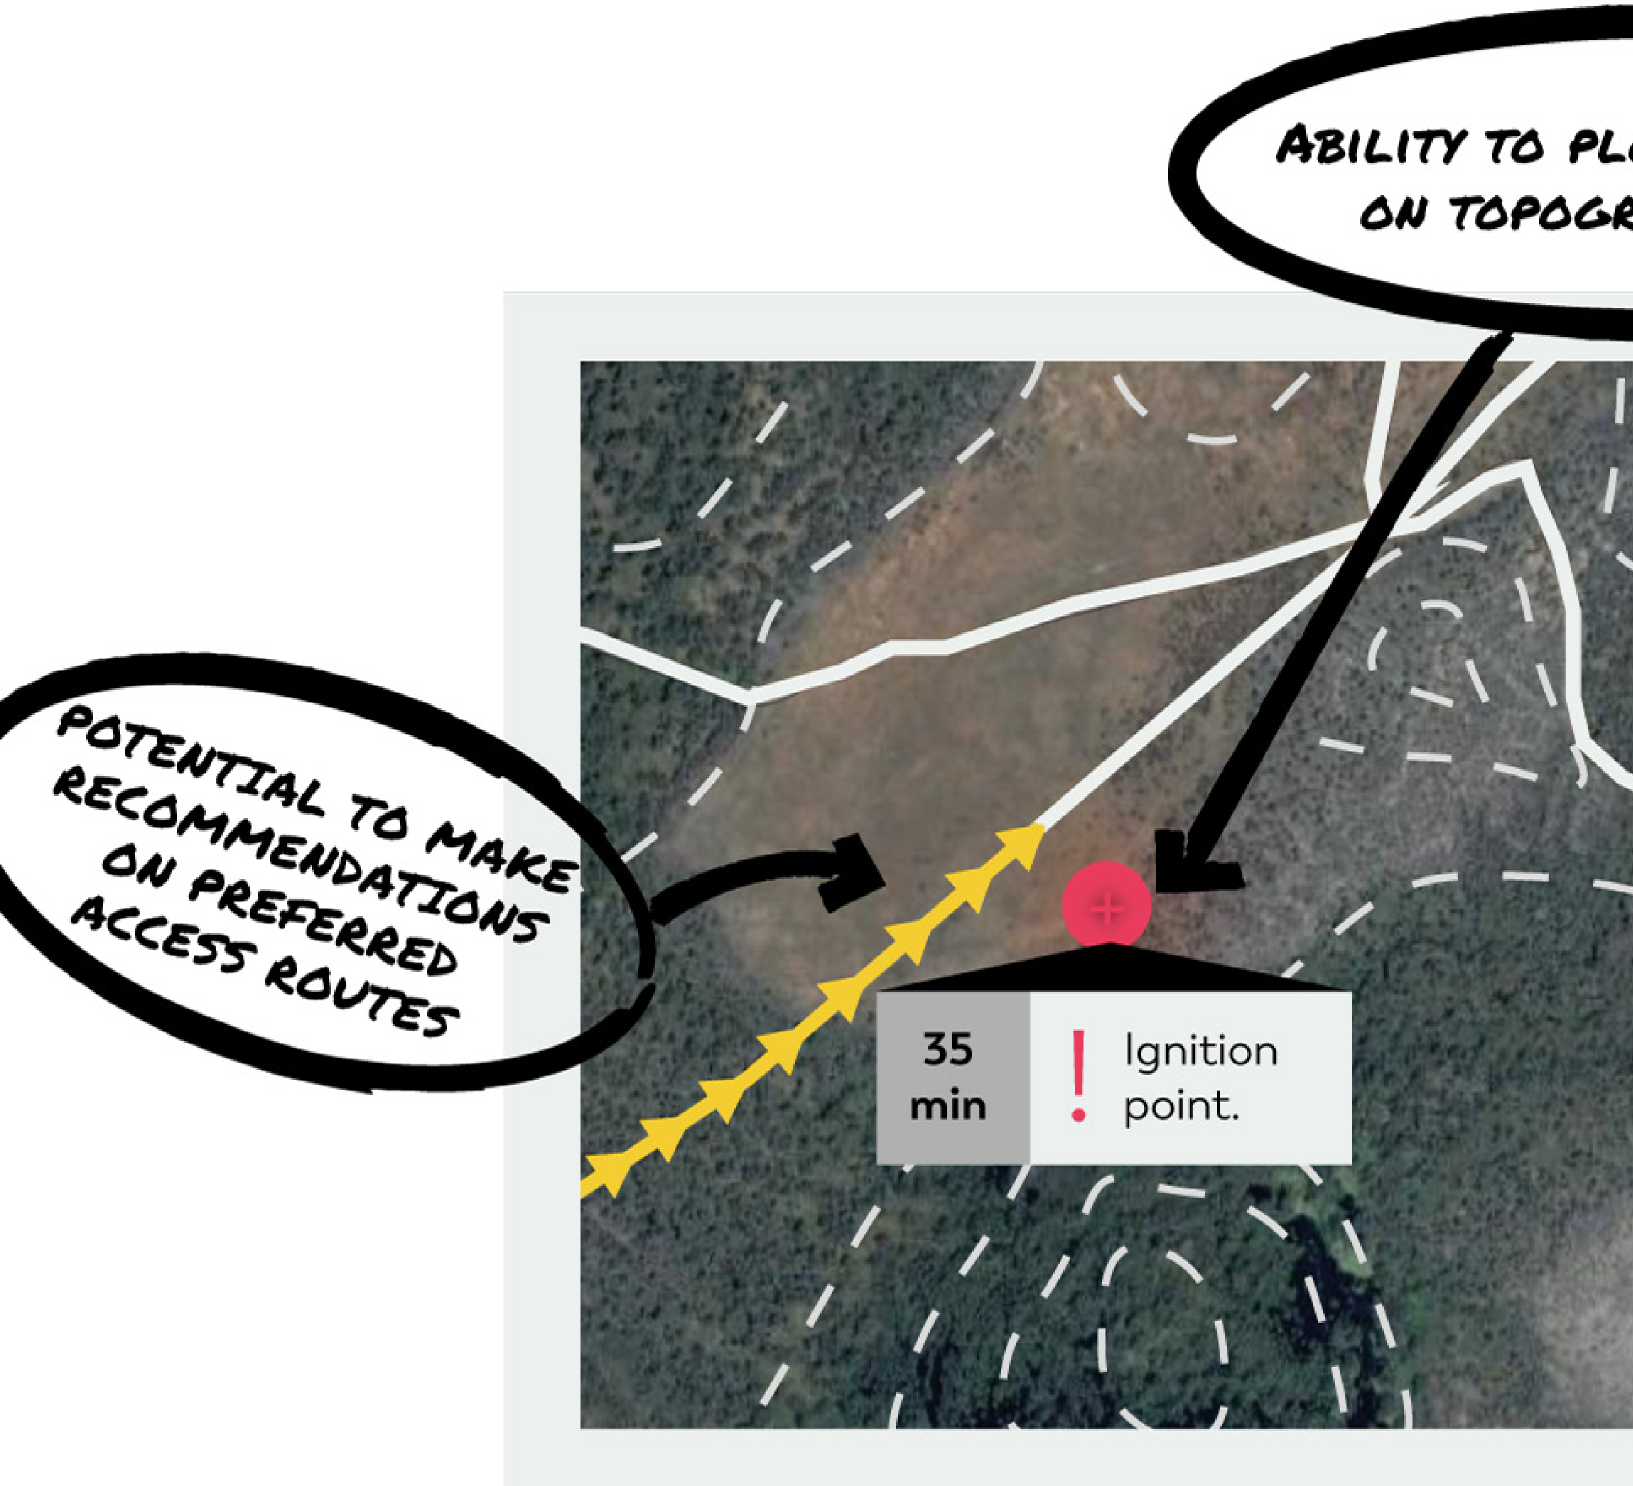

POTENTIAL FOR AUSTRALI
FIRE EARLY-WARNING SY

PLOT IGNITION POINT
GRAPHICAL MAPS

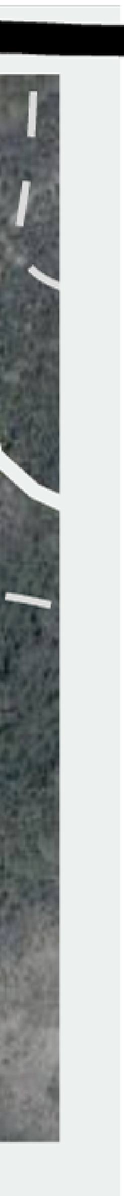

INGITION POINT
COULD BE
MERGED WITH
HIGH-RES FUEL
MOISTURE MAP

IA-WIDE
SYSTEM

IMAGINE

Detailed **operations maps** with **ignition point** and plan of **access** for firefighters.

Map of region, **ignition point, access routes, landmarks and fuel** printed and distributed to fire fighters.

FUEL AND FIRE RISK PREDICTION

DATA SOURCES INTEGRATED ONTO OPERATIONS MAP

SUCH AS WIND DIRECTION + SPEED

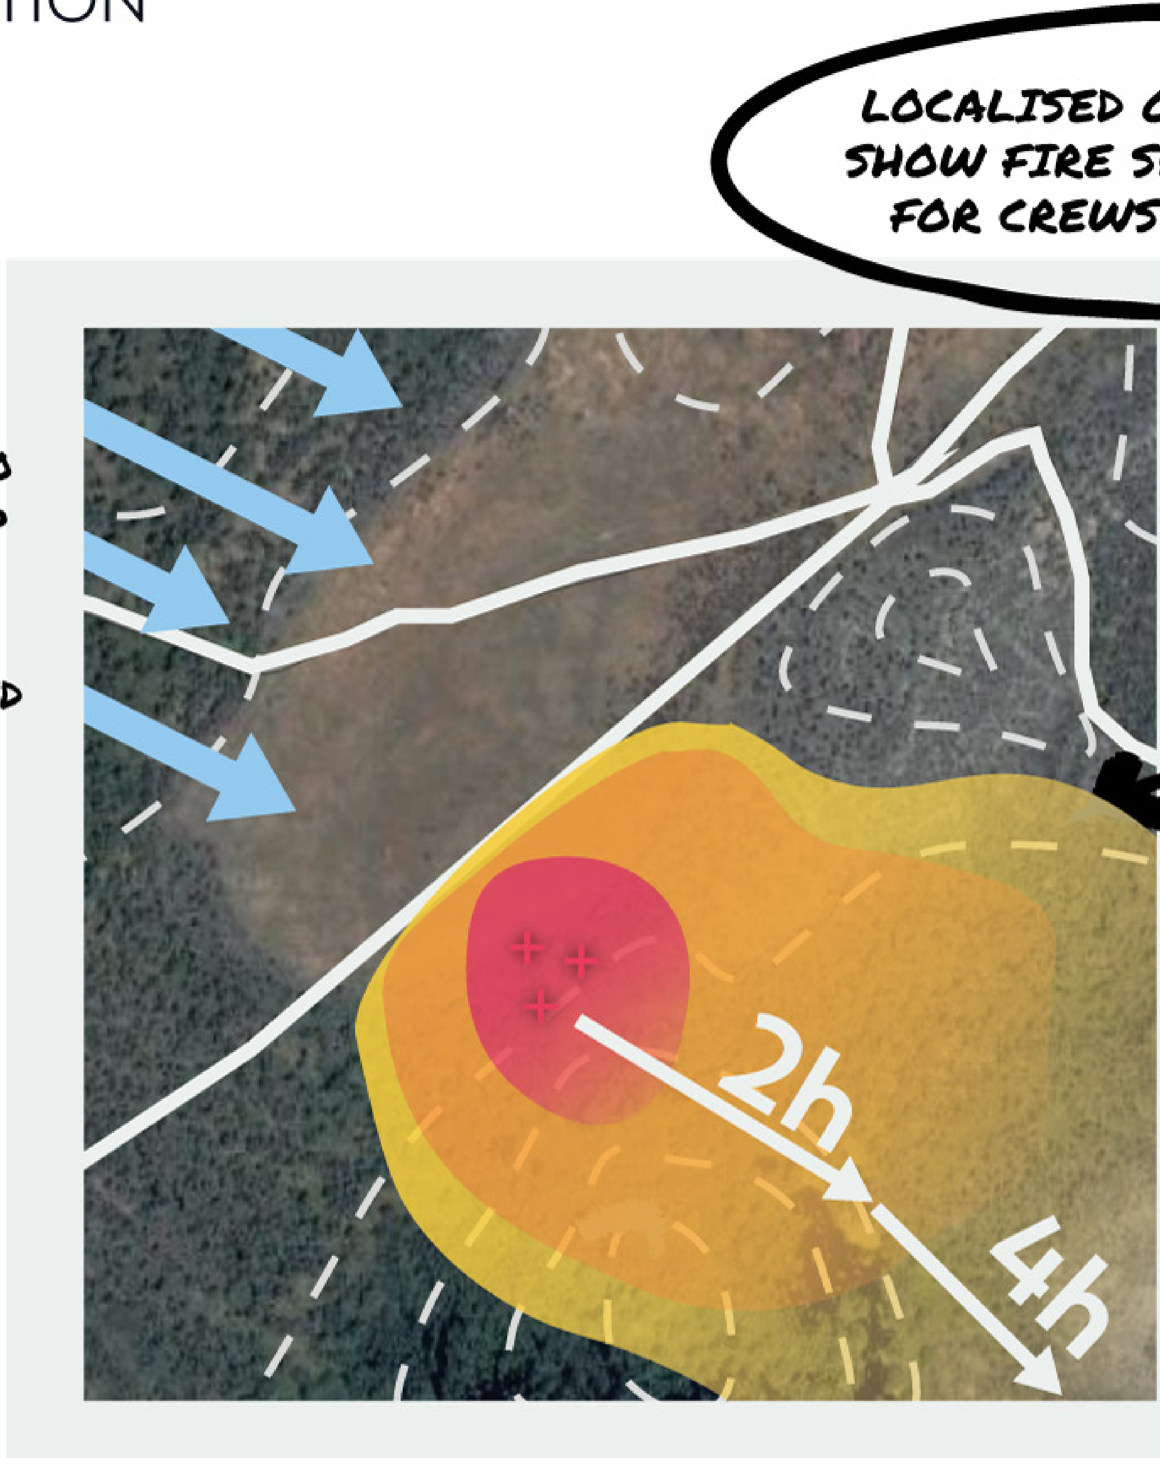

POTENTIAL FOR TIME ESTIMATES FACTORING CANOPY STRUCTURE AND WIND

IMAGINE

Fast AI based fire spread predictions on **low-cost hardware.**

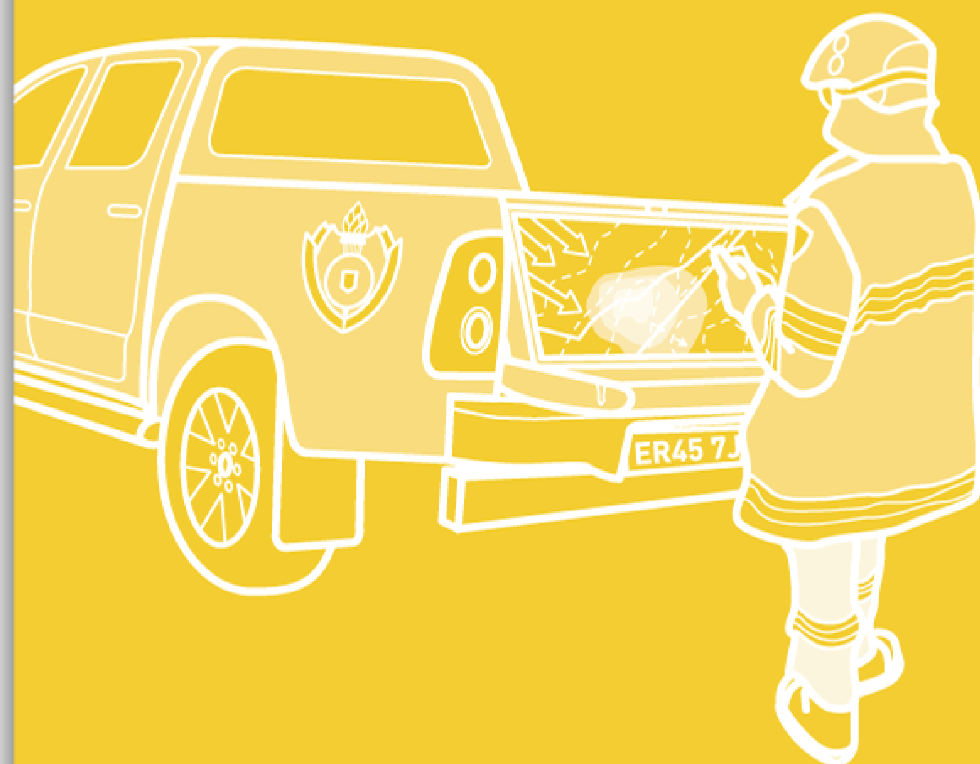

In-field operations center.

EXTREME FIRE
BEHAVIOUR

MONITOR:
-SMOKE COLOR
-CLOUD TEXTURE
-WIND DATA

POTENTIAL FOR ALERTS SENT TO FIRE CREWS
CIVILIANS IF EXTREME FIRE EVENTS DETECT

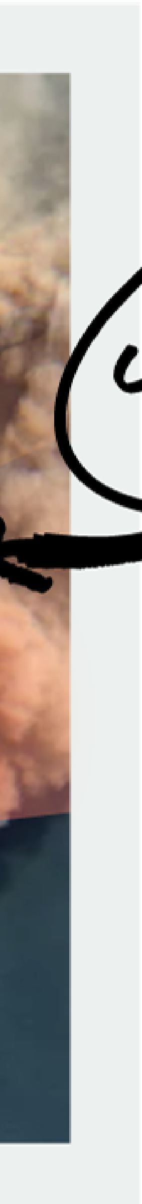

CLASSIFICATION OF SMOKE PLUMES

US AND
CTED

IMAGINE

In-situ **warning of extreme fire events** for fire fighters and civilians.

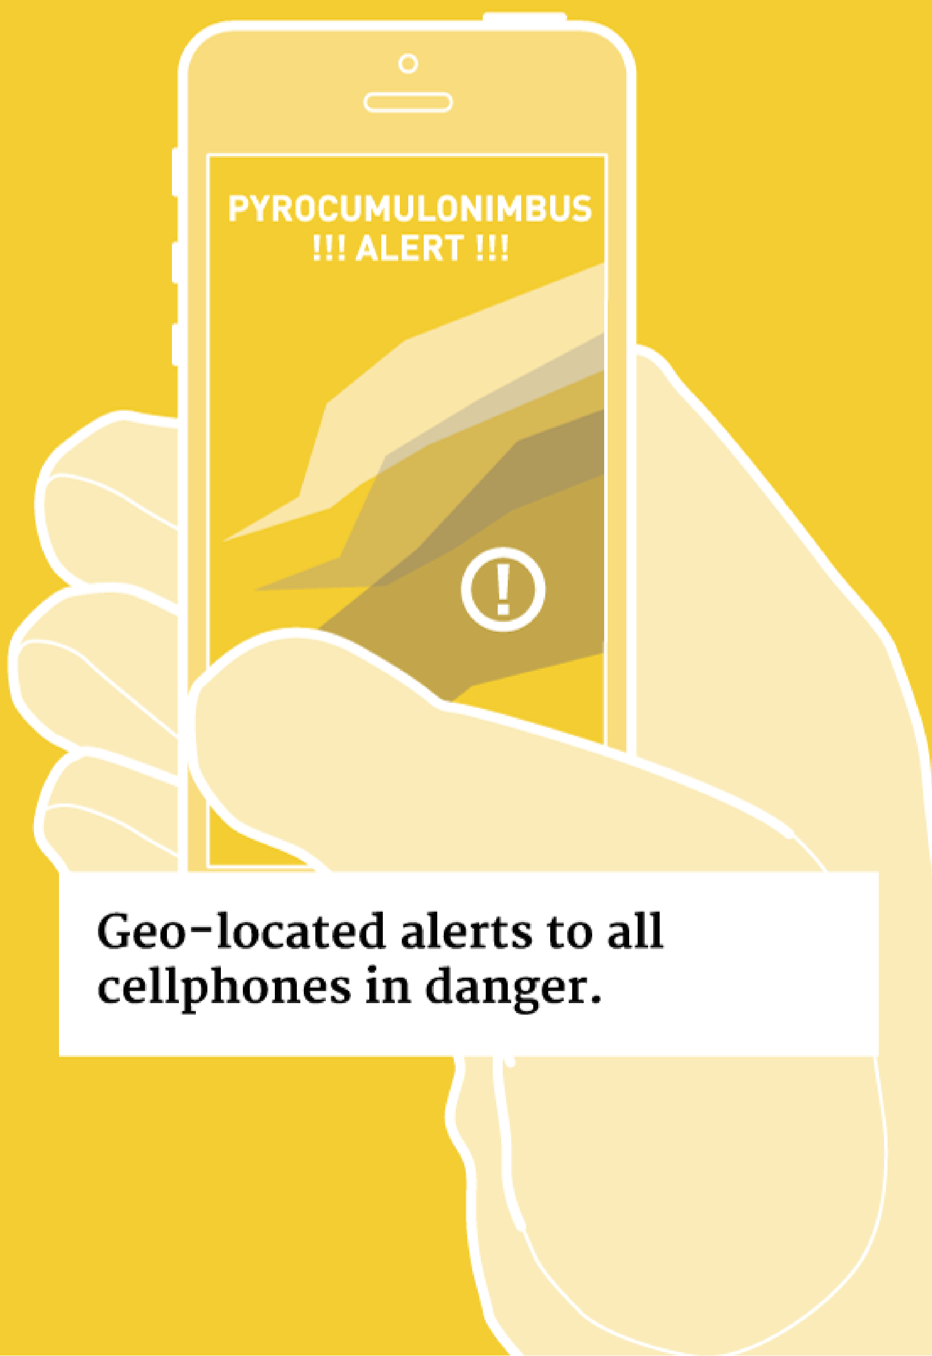

Geo-located alerts to all cellphones in danger.

THE IMPORTANCE OF VALIDATION & GROUND TRUTH

Machine learning shows great promise in delivering more accurate results on faster timescales.

However, there are also many pitfalls that can affect ML systems if improperly trained, causing them to fail under real-world conditions. For example, the Data Quest early detection workflow has great potential, but real-world validation tests are urgently needed to answer questions like "**How big does the fire need to be in different types of fuel, to be detected?**" and "**How do nearby heat sources, such a large slabs or rock, affect the detection algorithm - do they cause false positive detections or mask a new fire?**".

Another example is the Data Quest's new ML-driven fire-risk prediction workflow. This is a complex tool that uses a myriad of inputs to make a prediction: multi-band satellite data, models of terrain elevation, weather information and historic fire data. Live validation tests incorporating grids of advanced sensors are urgently needed to answer questions like "**Is a change in wind direction and speed accounted for in the prediction, and under what conditions does the prediction break down?**" and "**How does the distribution of ignition points affect the speed and intensity of the fire front?**" and finally "**Does the fuel moisture content or fuel size distribution dominate the fire risk predictions under different weather conditions?**".

Live trials are also a superb forum for training new firefighters in emerging technologies, but also for technology developers to gain insight from experienced firefighting professionals. Members and affiliates of the Australasian Fire and Emergency Service Authorities Council (AFAC), such as the Rural Fire Services (RFS) have particular needs which must drive the development of technology. Live trials and concurrent design studies that bring together users, data owners and technologists, will encourage a mixing environment for the ultimate benefit of all stakeholders.

The Data Quest also identified significant gaps in the data necessary to produce truly paradigm changing results.

Technologists, firefighting processionals, data custodians, government agencies, research institutes and private companies must all work together to fill these gaps. This can be achieved building common protocols and a common framework that guarantees interoperability of everybody's systems. We hope the partnership brought together by the 2020 Bushfire Data Quest can be the seed for that development effort.

THE FUTURE: FULL-SPECTRUM SITUATIONAL AWARENESS

The Data Quest has produced some excellent proof-of-concept technologies that promise to detect fires closer to ignition, create high resolution maps of fuel and fire-risk that are frequently updated, and detect the signatures of extreme fire behaviour. However, for maximum impact these technologies need to be deployed in a **framework that actively supports their use and integrates with dynamic, networked and robust communication tools. It is equally important that fire-fighting professionals guide the development of this system and take ownership of it as it grows. Adoption of new technology in the field will depend on good communication and technologists willing to listen - and quickly respond - to users.**

Through conversations with active firefighters we have identified that **the most immediate need is for robust communication of rich situational awareness information** between units on the ground and up the command chain. The ideal communication system should be based on **reliable network technology and underpinned by a real-time 'digital twin'** model that includes spatial information on equipment locations, unit orders, location of natural and built assets, importance of assets for protection, and information on individual firefighters such as medical data and allergies. This 'cyber-infrastructure' system should be enhanced by real-time data on the fire front, weather conditions, locations of water sources and risk-maps based on the latest fire-spread models - enhanced decedents of Data Quest outcomes.

We imagine that each Strike Team has access to a **connected device that shows an annotated topographic map** of the fireground and also serves as an **efficient multi-way communication tool**. Features on the map can be annotated and labeled by firefighters on the ground (e.g., firebreaks under construction, confirmed & potential water sources, possible helicopter landing areas, local assets and their importance) and this information is immediately synced to devices with all Strike Teams and then up the chain of command to the Incident Controller. Edit permissions would be granular, with some annotations remaining provisional until confirmed by an officer. The map would be augmented with remote-sensing data showing potential water sources, location of fire front, the latest fuel mapping information and predictive spread models .

The next diagram shows our proposed vison for developing the outcomes of Data Quest 2020.

The approach would blend more detailed understanding of user needs, existing technical and capital infrastructure, to be informed by a Megafires working group. Three more ambitious sprints would be then run looking at (1) Resilience, (2) Just-in-time insight and (3) the Response Management infrastructure to create the basis of a vision we are calling 'CATCOM.ai' (Catastrophe Command).

1. CONCURRENT DESIGN CONSULATION

Digital Earth AUSTRALIA

afac

Australian Government
Bureau of Meteorology

+ ...

TRL 0

User requirements, Data model & sources, Operation Spec

TRL 1

2. INTELLIGENT ALGORITHMS RESEARCH SPRINT

GEO-MAPPING FOR RESILIENCE

- Hydrology models
- Fuel mapping

JUST-IN-TIME PREDICTION

- Ignition detection
- Fire progression

RESPONSE MANAGEMENT

- AI-Assisted Resource Management

TRL 2

TRL 3

Common Data Channels, Integrated Modelling, Validation and Ground-truth focus

TRL 4

Integrated CATCOM Spec: Intelligence, Data, Architecture, API

TRL 5

3. ENGINEERING PARTNERSHIP, TRL 5 - 7

Full pilot applied to bushfire context in Australia

TRL 6

TRL 7

CONCLUSION

The new age of Megafires: the need for integrated Earth Observation and AI to manage bushfire resilience and response. (‘CATCOM.ai’)

During the two unprecedented and catastrophic fire seasons - the first in eastern Australia in early 2020, and the second in the USA in the middle of the year - emergency services and the information and comms infrastructure in both countries struggled to cope with the size and ferocity of the “megafires”, leading to avoidable deaths, large numbers of citizens becoming internal refugees, and massive destruction of property.

The adoption of bushfire technology is heavily influenced by government policy and the Australian authorities have just released three reports that make strong recommendations to adopt data-science methods and AI technology in support of Earth observations for bushfire management. Both the NSW Bushfire Inquiry and the Royal Commission into National Natural Disaster Arrangements make explicit a need to push available technologies harder, especially fire science, remote sensing, data science and artificial intelligence, to equip us better to understand what happens during a bush fire and respond more quickly - and recommends a data-fusion approach like those demonstrated during Data Quest. All 76 recommendations of the NSW Inquiry will be adopted by the NSW Government. The Bushfire Earth Observation Task Force (led by The Australian Space Agency, CSIRO, GA and BoM) offers even more detail and recognises the role of the private sector in driving innovation. Their report recommends securing the data pathways with international partners and supporting private industry to collaborate on developing remote sensing platforms and tools. These recommendations are aligned with state (e.g NSW) and Australian national strategies that aim to foster collaboration across the industrial and space research communities - within Australia and internationally.

However, it is our opinion that this will require a “whole-community approach”, with AI testing, validation, standards and compatibility built-in to the development cycle. By comparison, a market-driven approach may deliver disaggregated minimum-viable-products, motivated more by commercial concerns, rather than a joined-up Bushfire defense ecosystem.

We are calling this ecosystem level approach to developing a vision of an AI enabled cyber-command, ‘CATCOM.ai’

General Benefit of AI-assisted Cyber-Command for Bushfires

Effective bushfire management ideally requires an enhanced cyber-command infrastructure due to the extreme challenge and complexity of fighting fires, but also the overwhelming complexity in integrating and managing advanced geospatial, AI and communication technology. Moreover, any application of advanced technology to complex, dynamic and stressful environments context is key. Solutions must be closely tailored to the specific requirements of target users.

Speeding up the OODA loop (Orient, Observe, Decide, Act).

ML offers the ability to process and exploit large volumes of information, however there still remains a lack of clarity around what AI brings to a disaster scenario. After consultation with experienced fire-fighters and Military Capability Development (MDC) professionals, we've heard clearly that they don't want to feel removed from any decision making loop, rather they want the ability to combine different data streams and use AI to help make sense of the available information; giving users the 'right amount' of information, at the right time.

As such, the development of the CATCOM system is crucially tailored as a "decision support tool" - with numerous implications in the infrastructure architecture. A mature ML product will also produce measurable performance enhancements to the OODA loop (Orient, Observe, Decide, Act). Often attempts at these kinds of technologies are built around speeding up the OODA loop to make informed decisions at a rate that is faster than the progress of the fire. Developing infrastructure to speed up the **detect**, **collect**, **predict** and **disseminate** process, and provide a measurable metric will be of immense value in the coming decades.

CATCOM amplifies and galvanises existing Australian investments in in-space systems (e.g., satellites and bespoke instrumentation) and satellite-enabled services on earth (through Earth observation, communications, positioning and timing). As the space economy evolves, the balance of opportunity will shift from big to agile - and integrated. While global competitors in space chase macro capabilities such as launch systems and other hardware based infrastructure, the unexploited opportunity in space is to broker the system. That is: mastering autonomy, sensing, reasoning and decision making - being the brains, rather than the brawn, of humanity's future in space.

The opportunity we see is catalyse the development of a collaborative ecosystem between academia and tech providers, working with the fire-fighting services and agencies as a hub for an integrated platform like the CATCOM Cyber-command, becoming an 'integration hub' building global leadership, attracting and developing talent and inward investment.

ABOUT THE DATA QUEST FORMAT

The Data Quest is an accelerated research sprint designed to explore the solvability of challenge questions using a combination of machine learning and remote sensing data. The outcomes of the Data Quest are proof-of-concept data science workflows that show a question can be tackled using available data and state-of-the-art-algorithms.

In the parlance of NASA's Technology Readiness Level (TRL) system, Data Quest products are a level 2 outcome: Technology Concept Formulation, providing enough evidence to proceed from the concept stage to an experimental proof. At the start of the journey up the TRL ladder, the focus is on assessing the available data and its quality - hence the name 'Data Quest'.

The Data Quest has a few 'moving parts' that are key to its success:

1. Focused challenges: Tightly articulated challenge questions allow the teams freedom to explore while still setting achievable aims during the research sprint.

2. Interdisciplinary teams: Each team is carefully chosen from the brightest minds available and is a unique combination of experts suited to the research task - a 'dream team'.

3. Week-long sprint plus book-ends: A part-time on-ramp period allows the teams to spin-up and establish tight working relationships before the intense one-week research period. The off-ramp period facilitates efforts to document and finalise the research.

4. Experienced mentors: Team leads with experience in the problem domain (e.g., bushfire science or Earth observations) and machine learning guide and advise the teams. These are drawn from the partner institutions.

5. Community of partners: Data Quest, like FDL, is built on a network of partners that supply critical expertise, provide access to essential data, help set the challenge questions and inform the direction of research through use-case stories.

DATA QUEST IN NUMBERS

2020

9 Faculty

Scion Crown Research Institute, Australian National University, Northwest Nazarene University, Fireball, University of Oxford, UBISOFT, UNSW Canberra

4 TEAMS

4 FINAL PRESENTATIONS

4 RESEARCH POSTERS & TECH MEMOS

LIVE STREAMED SHOWCASE

581 LIVE VIEWERS

65 Applicants

MEDIA MENTIONS

CEDA, THE GREATER GOOD PODCAST AUG 13

DEDICATED TV SPOT ON 7-NEWS AUG 24

ITNEWS.COM.AU, NSW TURNS TO AI TO PREDICT BUSHFIRE ACTIVITY AUG 24

ABC RADIO INTERVIEW AUG 26

SCION CROWN RESEARCH ARTICLE AUG 28

17 Researchers

CSIRO, FrontierSI, Scion Crown Research Institute, Serenitec, Swinburne University, The University of Adelaide, University of Tasmania, University of the Sunshine Coast

Minderoo Foundation
Challenge Partner

dug
Technology Partner

NSW Government
Research Partner

Australian National University
InSpace ANU Institute for Space
Research Partner

TRILLIUM
TECHNOLOGIES

DATA
QUEST

MINDEROO FOUNDATION

Challenge Partner

dug

Technology Partner

NSW GOVERNMENT

Research Partner

Australian National University

InSpace
ANU INSTITUTE FOR SPACE

Research Partner