Program on Technology Innovation: Using Hyperspectral Imagery and Artificial Intelligence (AI) to Detect Stressed and Dead Trees

3002022770



Photo credit: William M. Ciesla, Forest Health Management International, CC BY-ND 2.0, source

Program on Technology Innovation: Using Hyperspectral Imagery and Artificial Intelligence (AI) to Detect Stressed and Dead Trees

3002022770

Technical Update, October 2021

EPRI Project Manager

B. Madsen

EPRI
3420 Hillview Avenue, Palo Alto, California 94304-1338 ▪ PO Box 10412, Palo Alto, California 94303-0813 ▪ USA
800.313.3774 ▪ 650.855.2121 ▪ askepri@epri.com ▪ www.epri.com

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THE FOLLOWING ORGANIZATION, UNDER CONTRACT TO EPRI, PREPARED THIS REPORT:

Dr. Zachary Tane

This is an EPRI Technical Update report. A Technical Update report is intended as an informal report of continuing research, a meeting, or a topical study. It is not a final EPRI technical report.

NOTE

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

© 2021 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute and EPRI are registered marks of the Electric Power Research Institute, Inc. in the U.S. and worldwide.

ACKNOWLEDGMENTS

The following organization, under contract to the Electric Power Research Institute (EPRI), prepared this report:

Dr. Zachary Tane

Principal Investigator
Dr. Zachary Tane

This report describes research sponsored by EPRI.

This publication is a corporate document that should be cited in the literature in the following manner:

Program on Technology Innovation: Using Hyperspectral Imagery and Artificial Intelligence (AI) to Detect Stressed and Dead Trees. EPRI, Palo Alto, CA: 2020. 3002022770.

ABSTRACT

This report:

- Summarizes the spectral changes to trees caused by drought, stress, and mortality.
- Reviews major considerations and state-of-the-art academic papers on detecting trees that have died using high spatial resolution imagery (0.3m – 2.4m).
- Provides an overview of the science of detecting trees that are undergoing various forms of stress using high spatial resolution imagery, reviewing what is commercially possible now, and looking forward to future advances.
- Reviews major vendors that provide remote sensing data for utility vegetation management.

Keywords
Artificial intelligence (AI)
Hyperspectral imagery
Rights of way (ROW)
Satellite imagery
Tree health
Vegetation management
Wildfire

Deliverable Number: 3002022770

Product Type: Technical Update

Product Title: Program on Technology Innovation: Using Hyperspectral Imagery and Artificial Intelligence (AI) to Detect Stressed and Dead Trees

PRIMARY AUDIENCE: EPRI utility member company staff involved in issues of wildfire risk, vegetation management, or technology innovation; and EPRI technical staff interested in the application of artificial intelligence (AI) and imagery data to wildfire risks and vegetation management in the electric power sector.

SECONDARY AUDIENCE: Other utility, commercial sector, and government staff involved in issues of wildfire risk, vegetation management, or technology innovation.

KEY RESEARCH QUESTION

What are existing applications of AI to imagery data (satellite data, hyperspectral imagery data) to identify dead and stressed trees?

What vendors are providing services to the themes noted above?

RESEARCH OVERVIEW

Information for this report was obtained by a review of relevant scientific literature, as well as contacting major commercial vendors related to using fine-scale remote sensing on trees. This report:

- Summarizes the spectral changes to trees caused by drought, stress, and mortality.
- Reviews major considerations and state-of-the-art academic papers on detecting trees that have died using high spatial resolution imagery. In this paper “high spatial resolution” generally refers to resolution in the range of 0.3m – 2.4m.
- Provides an overview of the science of detecting trees that are undergoing various forms of stress using high spatial resolution imagery, reviewing what is commercially possible now, and looking forward to future advances.
- Reviews major vendors that provide remote sensing data for utility vegetation management.

KEY FINDINGS

- Identifying tree mortality after a tree’s leaves are no longer green using high spatial resolution imagery is a well-established area of inquiry in remote sensing.
- The two main approaches for identifying visibly dead trees are to separate healthy trees from dead trees using vegetation indices (generally paired with some thresholding or clustering approach) and using a machine learning algorithm.
- While many modern machine learning algorithms can achieve good results, convolutional neural networks (CNNs) are one of the most cutting-edge machine learning methods to do this with high-resolution imagery. CNNs take into account spatial patterns in the vicinity of the analysis area.
- For identifying stress in visibly healthy trees, the spectral patterns are generally more subtle. The success of remote sensing devices that scan broad spectral ranges in identifying stress is mixed and dependent on the agent of stress. The loss of leaf area and discoloration of vegetation is likely identifiable, whereas identifying drought (moisture) stress in trees and early bark beetle infestation is not reliably detectable. With existing technology, the most accurate approach available is to track leaf area during drought and assume declines in leaf area are indicative of stress.

.

- Imaging spectroscopy (often known as hyperspectral remote sensing) wherein channels have very narrow spectral ranges (<10 nm) are generally more sensitive to the spectral signals associated with tree stress. The key spectral ranges associated with moisture stress appear to be at the red edge (680-720 nm) and some near-infrared wavelengths (around 940 nm and 1250 nm). Although there is no commercial spaceborne high-resolution imaging spectrometer currently, there may be soon. Currently, high-resolution imaging spectroscopy is possible over narrow spatial regions by aircraft, although this is typically done at a higher cost than satellite-based approaches.
- Several vendors provide high-quality products that may meet the specific needs of utility companies for vegetation management.

WHY THIS MATTERS

Vegetation maintenance around utility lines is a major expense for utility companies. As droughts, pests, and diseases all impact tree health, utilities need cost-effective tools that identify trees that pose a risk to utility infrastructure. High spatial resolution remote sensing is one technology utilities can employ to proactively manage at-risk trees, minimize costly power outages and decrease the cost of utility vegetation maintenance.

Dead and highly drought-stressed trees also create unique fuel behavior in wildfires. Understanding the location and spatial scale of these fuels and the likely near-term shifts in fuel types from live trees to dead trees will be important information in predicting and mitigating extreme wildfire behavior that is becoming increasingly common with climate change.

HOW TO APPLY RESULTS

The study provides EPRI members with a foundational understanding of applying remote sensing and machine learning techniques to satellite acquired high spatial resolution remote sensing to identify dead or stressed trees.

The first section “An Overview of Using High Spatial Resolution Imagery to Predict Tree Stress and Mortality During Drought” provides foundational information on identifying tree death and stress via remote sensing, as well as information specific to using high spatial resolution remote sensing. It may be best targeted for GIS specialists and vegetation managers without expertise in high spatial resolution remote sensing.

The next section “Identification of Red Attack and Gray Attack Trees using High-Resolution Imagery” focuses on summarizing the science behind identifying dead trees with remote sensing. The section after “Identification of Green Attack Trees and Tree Stress Using High-Resolution Imagery” focuses on identifying trees that are stressed, but are not visually dead. These two sections will be best used by GIS, remote sensing, and machine learning professionals. Both of these sections include a high level summary followed by an annotated bibliography.

The final section, “Evaluation of Commercial Vendors” summarizes the services of five relevant vendors in the utility vegetation management and high spatial resolution remote sensing space.

LEARNING AND ENGAGEMENT OPPORTUNITIES

- EPRI Wildfire Mitigation Planning: Research and Applications, EPRI report #3002022665, 2021
- Seeing Utility Infrastructure from Different Viewpoints (infographic), EPRI report #3002017724, 2020
- 3002017281

- Identification of At-Risk Trees Using Satellite Imagery, EPRI report #3002019050, 2020
- Results at a Glance: Vegetation Indexes for Hazard Tree Management, EPRI report #3002019574, 2020
- Vegetation Indexes for Hazard Tree Management, EPRI report #3002017448, 2020
- Vegetation Management Inspection Technologies Applications and Research Needs, EPRI report #3002017810, 2019

EPRI CONTACT: Becca Madsen, Senior Technical Leader, bmadsen@epri.com

PROGRAM: Technology Innovation

CONTENTS

ABSTRACT ..... **v**

EXECUTIVE SUMMARY ..... **vii**

1 **AN OVERVIEW OF USING HIGH SPATIAL RESOLUTION IMAGERY TO PREDICT TREE STRESS AND MORTALITY DURING DROUGHT** **1-1**

Tree Stress and Mortality During Drought from a Remote Sensing Perspective 1-1

Practical Considerations when using High-Resolution Sensors and State of the Art Methods 1-3

Further Reading 1-4

2 **IDENTIFICATION OF RED ATTACK AND GRAY ATTACK TREES USING HIGH-RESOLUTION IMAGERY** **2-1**

Summary Findings 2-1

Annotated Bibliography 2-3

3 **IDENTIFICATION OF GREEN ATTACK TREES AND TREE STRESS USING HIGH-RESOLUTION IMAGERY** **3-1**

Summary Findings 3-1

Annotated Bibliography 3-2

4 **EVALUATION OF COMMERCIAL VENDORS** **4-1**

AiDash 4-1

LiveEO 4-2

Salo Sciences 4-2

Satelytics 4-3

UP42 4-4

LIST OF FIGURES

Figure 1-1 Green attack, red attack, and grey attack phases of tree death. *Photo credit: J. Runyon, USDA Forest Service, source.* 1-2

1
AN OVERVIEW OF USING HIGH SPATIAL RESOLUTION IMAGERY TO PREDICT TREE STRESS AND MORTALITY DURING DROUGHT

Note: For an overall introduction to the report and approach, see the Executive Summary.

This section first provides some background connecting the physical processes of tree stress and death during drought with the spectral signal the remote sensing device is receiving. It also focuses on some particular challenges related to using and understanding high spatial resolution remote sensing. In this paper "high spatial resolution" generally refers to resolution in the range of 0.3m – 2.4m.

Tree Stress and Mortality During Drought from a Remote Sensing Perspective

As droughts grow hotter and more extreme as a result of climate change, tree mortality driven by water stress also grows. During droughts, there are two primary causes of tree mortality, excluding abiotic forces such as fire. The first major cause is an embolism, which occurs when excess air and water vapor in a tree's xylem[1] preventing successful water transport. The second major cause of mortality is tree pests, including both insects and pathogens. While some pests attack healthy trees, many insects specifically attack stressed trees compounding the effects of water stress. Water stress also diminishes the ability of trees to produce sap and defensive chemical compounds. In both cases, water stress is a primary driver behind tree mortality.

As plants transition from non-water stressed to increasingly water-stressed states, there are detectable changes in the spectral profiles[2] of plant leaves. Reflectance changes are caused by several factors including 1) chemical changes within leaves such as dropping moisture levels; 2) changes in leaf pigments, mainly in the relative levels of chlorophyll, carotenoids, and anthocyanins; and 3) changes in plant behavior, primarily in the timing and length of stomatal opening.

While there are many approaches to detecting reflectance changes via remote sensing, the most common method is to take two or more spectral regions that are known to be sensitive to these factors and algebraically combine them into a single number, called a spectral index. Practically, the ability of vegetation indices to correspond with the physical properties of the plants depends on several factors including 1) the spatial resolution of the sensor; 2) how much of a pixel is occupied purely by the plant's leaves; 3) the spectral resolution of the sensor (how narrow of a

[1] Xylem and phloem are the two critical components of a plant's vascular system, responsible for transporting water and nutrients between the leaves and roots.

[2] The reflectance of an object across many wavelengths, generally visually displayed as a chart with wavelength on the x-axis and reflectance percentage on the y-axis.

spectral region the sensor can sample); and 4) the bidirectional reflectance distribution function (BRDF).[3]

In the remote sensing-based scientific literature, it is often convenient to divide tree death into different unique spectral phases. The three phases most frequently referenced in the literature are green attack, red attack, and gray attack. These phases are primarily based on the model of a coniferous tree dying after a bark beetle attack, although some or all of these phase names are often applied to different types of trees and various causes of tree death.

In the remote sensing-based scientific literature, it is often convenient to divide tree death into different unique spectral phases. The three phases most frequently referenced in the literature are green attack, red attack, and gray attack. These phases are primarily based on the model of a coniferous tree dying after a bark beetle attack, although some or all of these phase names are often applied to different types of trees and various causes of tree death. Figure 1-1 below shows green attack (A), red attack (B), grey attack (C), and a mosaic of phases (D).

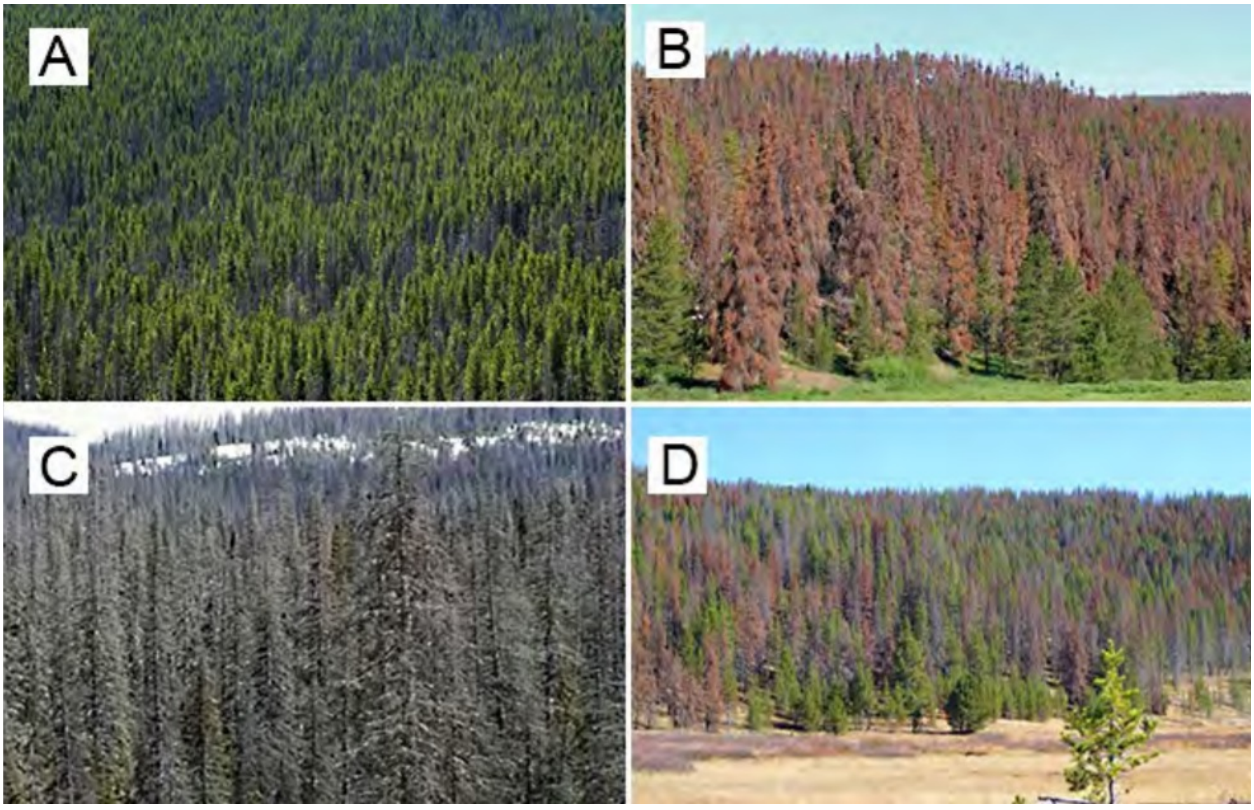


Figure 1-1
Green attack, red attack, and grey attack phases of tree death. *Photo credit: J. Runyon, USDA Forest Service, source.*

[3] A theoretical function that represents how light moves from a source, reflects off a target object, and reaches a sensor. Practically, in remote sensing, this refers to the angle and direction of sunlight, the atmospheric composition, the manner a target surface reflects light, and the angle and distance of the sensor from the target.

Bark beetles kill trees by creating tunnels through a tree's phloem system. When a tree has been killed by bark beetles, the response in the leaves is not instantaneous. Leaves can remain visibly green for months after the death of the tree. This phase is called the "green attack" phase. To the naked eye, the tree is indistinguishable from a healthy tree; however, there are subtle spectral shifts that can be detected by a sufficiently sensitive spectral sensor. After the "green attack" phase, the "red attack" phase begins, this is when a tree's leaves begin to turn yellow and then red. Spectrally, this phase is the most unique from other common targets such as healthy trees, ground cover, and soil. Due to the color change in the leaves, the "red attack" phase is spectrally the easiest to identify. Finally, once the tree has lost all its leaves, but is still standing, it has entered the "gray attack" phase. This phase can be more challenging to detect than the red attack. From the perspective of a satellite, a greater percentage of a pixel containing gray attack trees will also be occupied by other objects no longer blocked by leaves, such as understory trees, ground cover, and bare soil, diminishing the spectral signal of the dead tree.

Practical Considerations when using High-Resolution Sensors and State of the Art Methods

Proper georeferencing of high spatial resolution images will be critical if time series analysis is to be performed on the images, comparison with any ancillary data is to occur, or multiple images are to be mosaicked together to create a large scene for analysis. Although many commercial satellite products come with orthorectification and georeferencing already performed, additional testing should be taken to assure georeferencing accuracy is acceptable.

Since many high-resolution satellites like WorldView are pointable[4] off-nadir imaging is likely. Care should be taken as off-nadir imaging will cause non-linear spatial distortions across the scene.

If multiple images are mosaicked together for a single analysis, another consideration is the BRDF will vary across images. This means the same surface could have a different spectral profile, and the differences will be systemic across the multiple scenes. There are several ways to correct this. When training a supervised machine learning algorithm, if the algorithm is trained from images ranging across several scenes, it may be more robust to variations in spectral reflectance. Another approach is to implement a radiometric normalization algorithm to smooth the between image scenes.

The amount of light in each wavelength range a surface reflects from the sun to the sensor will determine each pixel with an image's brightness. The brightness of the pixels with a single tree's canopy can vary widely using high spatial resolution images, based on the slope of the branches, direction of the sun, and location of the sensor. This can be especially true on extreme topographic slopes, which causes the incidence angle of sunlight to vary wildly. Topography can cause the same target to be extremely bright in places where the sun is hitting the target directly and extremely dim where the topography shades the target from the sun. Semantic segmentation,[5] where each pixel is examined individually for its class is a very common task in

[4] In contrast with non-pointable satellites that will only image what's directly below them, pointable satellites are able to be tasked to collect images in specific areas and pointed at them in a slightly off-nadir direction.

[5] A machine learning task where every pixel in an image is placed into a class. There is, however, no notion of objects in semantic segmentation.

remote sensing. Since high spatial resolution remote sensing will have pixel by pixel brightness differences that are almost always more extreme than at coarser spatial resolutions, semantic segmentation can have more noise. One means of overcoming the problem of within target pixel variability is using a convolutional neural network[6] (CNN). CNNs are becoming increasingly common in industry and academia and are extremely well suited to leverage the high amount of contextual information provided by high spatial resolution imagery. Rather than viewing each pixel in isolation, CNNs allow contextual information across an analysis window to be considered. CNNs are an extremely active area of computer vision research, and a diverse set of network architectures[7] that meet various goals and constraints are rapidly being developed and proliferating. With tasks such as identifying tree canopy location, tree species, and tree mortality, leveraging contextual information can greatly reduce the noise in outputs.

Further Reading

Hall, R. J., Castilla, G., White, J. C., Cooke, B. J., & Skakun, R. S. (2016). Remote sensing of forest pest damage: a review and lessons learned from a Canadian perspective. The Canadian Entomologist, 148(S1), S296-S356.

Huang, C. Y., Anderegg, W. R., & Asner, G. P. (2019). Remote sensing of forest die-off in the Anthropocene: From plant ecophysiology to canopy structure. Remote Sensing of Environment, 231, 111233.

Kattenborn, T., Leitloff, J., Schiefer, F., & Hinz, S. (2021). Review on Convolutional Neural Networks (CNN) in vegetation remote sensing. ISPRS Journal of Photogrammetry and Remote Sensing, 173, 24-49.

Masaitis, G., Mozgeris, G., & Augustaitis, A. (2013). Spectral reflectance properties of healthy and stressed coniferous trees. Iforest-biogeosciences and Forestry, 6(1), 30.

Senf, C., Seidl, R., & Hostert, P. (2017). Remote sensing of forest insect disturbances: Current state and future directions. International Journal of Applied Earth Observation and Geoinformation, 60, 49-60.

Wulder, M. A., Dymond, C. C., White, J. C., Leckie, D. G., & Carroll, A. L. (2006). Surveying mountain pine beetle damage of forests: A review of remote sensing opportunities. Forest Ecology and management, 221(1-3), 27-41.

[6] A neural network (abbreviated NN) is a type of machine learning algorithm. A neural network has a layer of inputs (e.g. features), one or multiple layers of neurons which perform some simple algorithm based on one or multiple inputs, and an output layer. Neural networks are a supervised form of machine learning. Different NN have been developed to accomplish a wide variety of tasks including classification, object detection, and regression among others. A convolutional NN is a special class of neural networks, characterized by performing a convolution, a function that has multiple adjacent pixels as an input, on an image.

[7] Network architecture is the composition of the inputs, outputs, and intermediary layers of neurons and how they are connected.

2
IDENTIFICATION OF RED ATTACK AND GRAY ATTACK TREES USING HIGH-RESOLUTION IMAGERY

This section focuses on identifying trees that have, to the human eye, visibly died with high spatial resolution remote sensing. First, a summary of the best available science on this topic is given and followed by an annotated bibliography for the sources of the summary.

Summary Findings

Two main approaches have been taken for identifying red attack and gray attack trees. The most common approach is to use a machine learning algorithm to classify red attack and/or gray attack trees. Machine learning models have progressed along with machine learning research. Initially, simple machine learning methods such as maximum likelihood classification[8] were used with success. More recently, papers have used next-generation nonparametric algorithms such as Support-Vector Machine (SVM),[9] Random Forest (RF),[10] and Neural Network (NN) to achieve even higher accuracy. In the past couple of years, CNNs which use local context to make class determinations have achieved state-of-the-art accuracy.

Another approach to identifying red attack mortality is to calculate vegetation indices, generally the red green index (RGI)[11] and/or the normalized difference vegetation index (NDVI).[12] Generally, a baseline year is established, and then changes in the index through time are tracked. Pixel's that have spikes in RGI and declines and NDVI through time are likely to red-attack morality. Since healthy trees can have fluctuations in both indices as well, a key step is separating the amount of change that indicates red attack mortality versus a healthy tree. This has been done through iteratively testing thresholds until an acceptable accuracy has been obtained on training data. A more statistically robust approach is to use an unsupervised clustering algorithm that will automatically identify the clusters associated with healthy trees and red attack trees. Because the RGI and NDVI vegetation indices take advantage of the increased reflectance in red bands resulting from dying needles on beetle-infested trees, these indices have limited utility in gray attack identification.

---

[8] Maximum likelihood classifier is a simple supervised classification method. It assumes normal distribution of input features (which is not true of satellite imagery) and assigns the likelihood of test data falling in each class.

[9] Support vector machines are a supervised machine learning algorithm for classification and regression. The algorithm attempts to divide the entirety of the feature space into however many classes are given.

[10] Random forest is a supervised machine learning algorithm for classification and regression. It is an ensemble classification method using a user defined set of decision trees. Each decision tree is given a random subset of the feature space and training data.

[11] Red green index is a vegetation index calculated as Rred/Rgreen (R represent the reflectance of a channel or wavelength range

[12] Normalized difference vegetation index is a vegetation index calculated as (RNIR-Rred)/(RNIR+RRed). It is one of the most commonly used vegetation indices, often associated with green leaf density in a pixel and plant health.

With either a machine learning or vegetation index-based approach, known examples of the desired output classes are required both to train and validate the model. Two common methods of collecting known examples of each class are via field plots or trained technicians interpreting very high-resolution imagery. When considering the methodology for collecting training and validation data, it is important to consider the output classes. The output classes need not be limited to tree health; however, if different cover classes are desired it's important to consider how to collect training and validation data for them. Even a trained image interpreter may not be able to separate all classes from high-resolution imagery. Other classes may only be possible to collect accurately via image interpretation. In some cases, a mixture of field data and image interpretation may be appropriate. Collection of sampling data should be done in a manner that adequately covers the variability of the area of interest and considers issues of class balance. Specifically for high-spatial-resolution remote sensing, for the field-collected data to be useable, it will need to be collected using high precision GPS, which can add additional cost and training for proper usage.

Several practical considerations should be made when trying to implement these methods on high-resolution images. Shaded areas are low information parts of the image that, if not properly addressed, will yield results that are likely erroneous. If a machine-learning algorithm approach is being used, often, making an explicit shade class is the most straightforward means of addressing shadows. Another means of identifying shade is to exclude pixels from an analysis that do not have digital number values for some band or band combination above a predetermined threshold. Excluding pixels is simple, doesn't depend on using a machine learning method, and is accurate. However, the threshold value used to exclude pixels is a subjective decision and can have serious implications for the final results. CNNs mostly avoid the shadow problem since they can take into account contextual information around the shadows. Regardless, if the training and validation data was not collected from image interpretation, it is important to verify that the pixels selected are not shadow-covered.

Some studies remove portions of the test imagery from analysis before the imagery is run through the algorithm, known as masking. For machine learning-based methods, if areas of the test imagery contain cover types not included in the original training class, masking should be employed to maintain result integrity. For vegetation index-based methods, masking will likely be necessary if the images contain non-vegetated surfaces, especially man-made surfaces. The most straightforward way to mask is to remove areas if there is existing information on the location of surfaces; which is often true for man-made surfaces such as buildings and roads. If the ancillary information on the location and size of the surfaces is accurate, masking can be done without downsides.

A mixed pixel is when a single pixel contains multiple classes of interest, for example, at the edge of a tree's canopy a single pixel might contain both green canopy and soil. For pixel-by-pixel machine learning-based methods, this can be a major source of error. Generally, classifiers are trained on pure examples of each class such as only tree leaves or soil. The example pixel that's getting roughly equal reflectance contributions from both soil and green vegetation might have a spectral profile that the classifier may see as most similar to a red attack tree canopy. One method used to ameliorate the problem of mixed pixels is applying a moving window transform on the final classification map to remove some of the noise associated with mixed pixels. CNNs should mostly not be impacted by this problem as context clues should indicate the pixel is

mixed. Another direct means of addressing this problem is spectral mixture analysis, where a single pixel is allowed to contain multiple different cover types.

Studies using high spatial resolution satellite data are generally limited to a relatively small study area. This is due to the small footprint of high spatial resolution sensors, as well as their expense. Many papers studying tree mortality limit their focus to forests dominated by a single species (often lodgepole pine) that have all been killed by the same disturbance agent (often mountain pine beetles). The sensitivity of different models to differences in ecosystem types, plant phenology, solar illumination, mortality agents, among other factors has not been evaluated thoroughly in the literature. Even spectral index-based methods, wherein analysis is firmly correlated with the physical properties of plants and plant death, need to be tuned based on training data. The exact thresholds needed to obtain ideal accuracy likely vary in different conditions. In summary, users should not assume the accuracy of a model trained and validated in one set of conditions will perform similarly in a new environment.

Annotated Bibliography

Chiang, C. Y., Barnes, C., Angelov, P., & Jiang, R. (2020). Deep Learning-Based Automated Forest Health Diagnosis From Aerial Images. *IEEE Access*, *8*, 144064-144076.

STUDY AREA: Southern Scotland.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Various Scottish hardwoods, primarily oak, no major disturbance agent.

IMAGE SOURCE AND RESOLUTION: Aerial photography, precise pixel size not specified (< 1m).

METHODS USED: Mask RCNN[13].

TRAINING DATA SOURCE: Synthetically created images cutting and pasting from real images.

VALIDATION DATA SOURCE: Image interpretation.

OUTPUT CLASSES: Instance segmentation of dead trees.

SUMMARY: This paper uses Mask RCNN, a modern RCNN implementation capable of instance segmentation[14]. Because there were insufficient examples of dead trees in their training data, the authors cut out the approximately 300 dead trees from their training data and pasted the trees onto 63 backgrounds (an example of synthetic training data creation), greatly extending the number of training images available. The study focused on the best manner to train a mask RCNN to identify dead trees. They experiment with modifying several parameters, components,

[13] RCNN is a family of CNNs. Notable for being one of the first CNNs capable of object detection, by creating 2,000 regions that objects could be in an image. Mask RCNN is one of the more modern RCNN implementations, having first been documented in 2017. While a modern CNN with respectable performance, it is no longer state of the art.

[14] The identification of polygon regions containing objects. Unlike semantic segmentation, two adjacent objects of the same class would be noted as different in instance segmentation.

and methods of training for the RCNN including two different pre-training datasets (Image-Net and COCO)[15], different learning rate schedules[16], different learning rate optimizers (Adam versus stochastic gradient descent (SGD)[17], different training dataset sizes, different CNN activation functions (Relu versus Leaky Relu), and using dropout (a means of preventing overfitting) or not. They trained each model for 100 epochs[18]. Based on various experiments they found COCO pre-trained networks, using a Relu activation function, using the SGD learning rate optimizer, not using dropout, and using a dataset of 3,600 images (on the larger end of the data set sizes they experimented with) to be best. They had the best results with a fast-medium-slow learning rate setting (learning rate schedule $1E^{-4}$, then $1E^{-5}$, then $1E^{-6}$). Overall the best model had a low tendency to false positive, but relatively low recall[19] scores, only finding about half the dead trees in the scene.

While it's tough to know how generalizable the results of this paper were, it gives a good overview of some common details to be aware of when training a CNN for remotely sensing tree mortality. It also contains a high-level background on the usage of CNN generally and Mask RCNN specifically. The use of synthetically created images to boost the amount of training and validation data is interesting, although they conclude real data gets the best results. Still, in cases of extremely rare classes or extreme class imbalance, it may be a viable solution.

Dennison, P. E., Brunelle, A. R., & Carter, V. A. (2010). Assessing canopy mortality during a mountain pine beetle outbreak using GeoEye-1 high spatial resolution satellite data. ***Remote Sensing of Environment*****,** ***114*****(11), 2431-2435.**

STUDY AREA: Southeastern Wyoming.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Lodgepole pine and mountain pine beetle.

IMAGE SOURCE AND RESOLUTION: Geo-Eye 1[20], 1.5 m pixel size after pan-sharpening[21].

---

[15] NN are initialized with random weights. A common approach is to feed the neural network large but general training datasets for a while so the neural network is primed for common textures and patterns, and then later feed the neural network your specific data. This priming is called pre-training.

[16] Learning rate is a critical hyperparameter for training neural networks. Controls how quickly a model will change its weights based on error.

[17] A learning rate optimizer tweaks the base learning rate based on the perceived error on the validation data. Adam and SGD are two different optimizers which have different trade-offs in terms of speed, accuracy, and generalizability.

[18] An epoch occurs when a neural network has seen the entire input dataset once.

[19] A measure of object detection performance. Calculated as the number of true positives divided by the number of true positives plus false negatives.

[20] Launched in 2008 a commercial high spatial resolution satellite. Has red, green, blue, and IR channel as well as panchromatic channel at 1.5 m resolution.

[21] A common technique where the higher resolution panchromatic channel is used to infer what the slightly lower resolution multispectral images would look like at slightly higher spatial resolutions. The accuracy of these predictions vary, however for most high resolution sensors the process is thought to be fairly accurate and is commonplace.

METHODS USED: Maximum likelihood classifier.

TRAINING DATA SOURCE: Image interpretation (of a high-resolution aerial image).

VALIDATION DATA SOURCE: Field Plot.

OUTPUT CLASSES: Healthy canopy, red-attack canopy, gray-attack canopy, and shadow.

SUMMARY: It takes a novel approach to validating the output classes by using field plot data containing tree health status and a tree's diameter at breast height to estimate the canopy cover within a plot. Then, for each plot, the classifier-generated coverage of the four classes was estimated. The shade coverage, which is not represented in the field data, is removed and the percentage for the three other cover classes is compared with field data. The correlation between the field plot estimates and the classifier estimates was strong (R^2 values ranged between 0.67 and 0.77 for the three cover types). Aggregating the different plots and cover types obtained estimates for total coverage that was within 1.7% of field data. Since all validation plots were in forested areas, the problem of mapping non-forested areas was ignored. This study is notably the first peer-reviewed journal article to assess gray-attack trees using high spatial resolution imagery. It also takes interesting approaches to validating a classification map with field data and dealing with shade.

Garrity, S. R., Allen, C. D., Brumby, S. P., Gangodagamage, C., McDowell, N. G., & Cai, D. M. (2013). Quantifying tree mortality in a mixed species woodland using multitemporal high spatial resolution satellite imagery. *Remote Sensing of Environment*, *129*, 54-65.

STUDY AREA: Northern New Mexico.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: The study Spans several species and disturbance agents, however primarily piñon-juniper woodlands.

IMAGE SOURCE AND RESOLUTION: Both QuickBird[22] and WorldView-2[23] imagery, ~ 0.6 m pixel size after pansharpening.

METHODS USED: To classify vegetated from non-vegetated cover types a supervised classification using the Genetic Imagery Exploitation (GENIE) algorithm, which combines a genetic algorithm[24] and linear discriminant analysis (LDA)[25] was used. To classify live and dead trees, yearly differenced NDVI and RGI values for pixels were used. A Gaussian mixture model (GMM) was used to cluster randomly selected pixels by their NDVI/RGI values. To best determine the optimal number of clusters, the Akaike Information Criteria was used.

---

[22] Launched in 2001 one of the oldest commercial high resolution sensors. Contains four multispectral bands (red, green, blue, and near infrared) at 2.9 m pixel size and a panchromatic band 70 cm pixel size. Geoeye was turned off in 2015.

[23] A commercial multispectral satellite launched in 2009. Scans 8 multispectral bands spanning the color and near-infrared spectrum, also scans a panchromatic band. Multispectral bands are 1.8 m pixel resolution and the panchromatic band is 50 cm pixel resolution (actual resolution may vary if images were taken off-nadir).

[24] A heuristic search algorithm, inspired by natural selection.

[25] A supervised machine learning algorithm for regression which seeks to find a linear combination of the available features to make predictions.

TRAINING DATA SOURCE: Image interpretation (of a higher resolution aerial image).

VALIDATION DATA SOURCE: For the supervised classification, interpretation of aerial image was used. For the assessment of the cluster's ability to match live and dead trees both interpretation of aerial image and a single 1 ha plot.

OUTPUT CLASSES: Clusters representing live and dead trees were identified.

SUMMARY: Unlike many studies using high spatial resolution imagery, the study area spans several different coniferous ecosystems and several different causes of tree mortality. This study uses a time series of three different images, two QuickBird images were acquired in 2002 and 2006 and a WorldView image was acquired in 2011. Images in the different time series were radiometrically normalized using an iteratively re-weighted Multivariate Alteration Detection (IR-MAD) algorithm. Areas that ancillary data indicated were not occupied by trees were masked out. Then a supervised classification was done to remove all parts of the unmasked image area that were not live tree crowns in 2002 images. After images were masked to only contain areas that are live tree crowns in 2002, NDVI and RGI were calculated for each image in the series. The values were then differenced between the images (both 2006 and 2011 values were differenced against 2002 images), and 20% of the pixels in the image were then randomly selected and clustered into four clusters in an unsupervised manner. Clusters were then determined to be representative of the live or dead trees. Extremely high accuracies were obtained (97.9% and 98.5%) for predicting tree mortality. The change detection analysis did not distinguish different stages of tree death and easily detected red attack, gray attack, and fallen trees. Since the study area spans a broad set of classes and the methodologies are fairly general this analysis may be a framework for a larger scale general solution using relatively simple approaches and time series analysis.

Kislov, D. E., Korznikov, K. A., Altman, J., Vozmishcheva, A. S., & Krestov, P. V. (2021). Extending deep learning approaches for forest disturbance segmentation on very high-resolution satellite images. *Remote Sensing in Ecology and Conservation*.

STUDY AREA: South Sakhalin Island and South Kunashir Island.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Various coniferous and broadleaf forest types; bark beetles and windthrow.

IMAGE SOURCE AND RESOLUTION: Images WorldView-2, WorldView-3[26], Pleiades-1A, Pleiades-1B[27], 0.3-0.5 m pixel size after pansharpening.

---

[26] Launched in 2014 as the next generation WorldView satellite. Contains 8 multispectral bands and a panchromatic band. The multispectral bands have ~1.3 m resolution and the panchromatic band has 40 cm resolution (actual resolution may vary if images were taken off-nadir.

[27] The Pleiades satellites constellation are high resolution commercial satellites. They contain 4 multispectral bands and a panchromatic band. The multispectral bands have 2 m resolution and the panchromatic band has 0.5 m resolution.

METHODS USED: U-net-like CNN[28], AdaBoost[29], K-nearest neighbor[30], Naive Bayes, Quadratic Discriminant Analysis[31], RF, SVM.

TRAINING DATA SOURCE: Image Interpretation.

VALIDATION DATA SOURCE: Image Interpretation (different images than train).

OUTPUT CLASSES: The probability pixel was windthrow or bark beetle mortality.

SUMMARY: The study compares a CNN designed for semantic segmentation with several common modern machine learning techniques for detecting windthrow and bark beetle impacted pixels. The training processes used various pan-sharpened high-resolution scenes with just the RGB channels. For red attack bark beetle detection, most algorithms do similarly well in a validation scene that is mostly green and red attack trees; however, in scenes that present tricky false positives such as mining pits or roads with exposed soil adjacent, the CNN drastically outperforms other algorithms. The ability to interpret the spatial context and patterns is beneficial to CNN as compared to the other modern machine learning algorithms working on a pixel-by-pixel basis. One image that the CNN does produce false positives on is a mining pit that was roughly tree-shaped and sized, illustrating that the CNN is generally using patterns to prevent other false positives. For windthrow detection, there are generally large-scale spatial patterns to the impacted area, the CNN also generally outperformed other methods. While this is likely implementation-specific, using common machine learning libraries, the authors also found the CNN to train and evaluate faster on very large training sets.

On the windthrow dataset, the authors also attempted to train the other common machine learning methods with additional features that used moving window techniques to include contextual information. While traditional machine learning methods did improve when given this additional information, improvement was not significant enough to match the CNNs performance. The authors suppose that with extensive feature engineering these methods might be trained to match the CNNs performance. However, these feature engineering methods are time-consuming and a CNN can achieve similar results without this step.

Hicke, J. A., & Logan, J. (2009). Mapping whitebark pine mortality caused by a mountain pine beetle outbreak with high spatial resolution satellite imagery. *International Journal of Remote Sensing*, *30*(17), 4427-4441.

STUDY AREA: Central Idaho.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Whitebark pine and mountain pine beetle.

---

[28] A modern family of CNNs, the modifications made by the authors make it hard to know whether this implementation is getting state of the art performance, but it is likely not far from it.

[29] An ensemble supervised machine learning method similar to RF. During the training process the algorithm tweaks the weaker decision trees to be stronger.

[30] An unsupervised clustering machine learning method.

[31] Similar to LDA, except the terms can be quadratic rather than linear.

IMAGE SOURCE AND RESOLUTION: QuickBird, 2.4 m pixel size.

METHODS USED: Maximum likelihood classification.

TRAINING DATA SOURCE: Image interpretation.

VALIDATION DATA SOURCE: Image interpretation and field identified red and green trees.

OUTPUT CLASSES: Green herbaceous cover, brown herbaceous cover, green trees, and red attack trees.

SUMMARY: Whitebark pine, the focus of this study, typically grows sparsely in extremely mountainous terrain. To remove non-vegetated cover types, both a minimum NDVI and minimum red reflectance threshold were run across the images. Two vegetation indices, RGI and green reflectance were fed into the maximum likelihood classifier. Due to the extreme nature of the terrain, five separate classifiers were needed to overcome the varying solar incidence angle as determined by a USGS digital elevation model (DEM). Since whitebark crowns were considered to cover a minimum of four pixels in the imagery, if a group of fewer than four pixels was identified as red attack mortality they were mapped to the second most likely class. Overall accuracy was high (86% for the four image-derived classes and 90.9% compared to field data). The work in this study is demonstrative of the sensitivity of high-resolution remote sensing to the local variation of brightness caused by topography.

Meddens, A. J., Hicke, J. A., & Vierling, L. A. (2011). Evaluating the potential of multispectral imagery to map multiple stages of tree mortality. *Remote Sensing of Environment*, *115*(7), 1632-1642.

STUDY AREA: North Central Colorado.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Lodgepole pine and mountain pine beetle.

IMAGE SOURCE AND RESOLUTION: 4 channel aerial imagery, originally 30 cm pixel size, the imagery was also aggregated to 1.2, 2.4, and 4.2 m for analysis.

METHODS USED: Maximum Likelihood Classification.

TRAINING DATA SOURCE: Image Interpretation.

VALIDATION DATA SOURCE: Same as train evaluated by 10 fold cross-validation.

OUTPUT CLASSES: Green tree, red tree, gray tree, herbaceous, bare soil, shadow.

SUMMARY: This paper starts with extremely fine resolution 30 cm pixel size aerial photography. The authors did not attempt to mask out ground cover in any way, they just added bare soil, shadow, and herbaceous cover as classes to the classifier. While an accurate classifier was trained at the 30 cm scale, there was confusion between brightly lit healthy trees and herbaceous cover and brightly lit red attack trees and soil. The imagery was aggregated to three larger spatial resolutions and the same methods were repeated. Statistically significant improvements to the classifier's accuracy were achieved at the 2.4 m scale. The improvements were thought to be achieved because aggregation averages the most sunlit parts of tree's

canopies with less sunlit parts creating a spectral signal that was distinct from bright soil and herbaceous cover. Effectively aggregation is acting as crude means of adding local context in this study.

Stych, P., Jerabkova, B., Lastovicka, J., Riedl, M., & Paluba, D. (2019). A comparison of worldview-2 and landsat 8 images for the classification of forests affected by bark beetle outbreaks using a support vector machine and a neural network: A case study in the sumava mountains. ***Geosciences*****, *9*(9), 396.**

STUDY AREA: Sumava Mountains, the Czech Republic.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Norway spruce and the European spruce bark beetle.

IMAGE SOURCE AND RESOLUTION: WorldView-2 Imagery**,** 1.84 m pixel size.

METHODS USED: SVM and an NN.

TRAINING DATA SOURCE: Image interpretation (of an aerial photo and WorldView-2).

VALIDATION DATA SOURCE: Field data for the forest classes, photo interpretation for the non-forest classes.

OUTPUT CLASSES: Healthy Forest, affected forest (red attack / gray attack combined), regenerating forest, dead forest (stumps and trees without branches), wetlands, clear-cut forest, permanent grassland, water bodies, and artificial surfaces.

SUMMARY: Note that this paper also uses medium-scale Landsat imagery, but that part of the analysis is not part of this targeted review. This paper tests two modern machine learning algorithms, SVM and NN, the ability to classify nine different cover types across a primarily forested area. The tree mortality-related classes were a bit unusual, lumping red and gray attack into a single stage. While both algorithms achieved high accuracy scores, SVM outperformed NN (86.24% overall accuracy versus 81.65%). While the authors reach somewhat strong conclusions about the appropriateness of SVM over NN for tree mortality classification, their conclusions are, in my opinion, a bit too strong for a single data point. Many factors may have resulted in SVM outperformance, in this case, their limited training data (only consisting of 11 - 27 points for class) may favor SVM. Both algorithms achieved high accuracy when faced with the task of classifying nine separate cover classes, many of which are fairly similar spectrally, with limited training data.

Sylvain, J. D., Drolet, G., & Brown, N. (2019). Mapping dead forest cover using a deep convolutional neural network and digital aerial photography. ***ISPRS Journal of Photogrammetry and Remote Sensing*****, *156*, 14-26.**

STUDY AREA: Southern Quebec, Canada.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: A diverse mixture of northeastern hardwood and softwood forests, the forest was not disturbed and had low levels of mortality.

IMAGE SOURCE AND RESOLUTION: 990 aerial photos containing red, green, blue, and near-infrared images at 20-cm pixel size.

METHODS USED: Convolution Neural Networks (CNN).

TRAINING DATA SOURCE: Image interpretation.

VALIDATION DATA SOURCE: Same as train.

OUTPUT CLASSES: Live conifer, dead conifer, live broadleaf, dead broadleaf, water, road, cut forest, wetland, buildings.

SUMMARY: Using a large number of aerial photos, examples of live and dead trees over a wide area and a diverse species mixture were sampled. Generally, data related to training and testing neural networks is split into training, validation, and test, where validation data is used to evaluate whether the neural network has overfitted to the training data during the training process and test data is used to evaluate the neural network at the end of the training process. In this study data was split as 60% training, 10% validation, and 30% test. Spatial adjacency constraints were also placed on the data to make sure validation and test data were not too close to training data. Since there was no specific mortality agent in this study, dead trees were significantly more rare than live ones. To avoid class imbalance problems, examples of live trees were undersampled until class parity was achieved. Several augmentations (image manipulations such as flips and rotations) were applied to the training data. A 16 layer visual geometry group (VGG)[32] CNN was used. SGD was used, a learning rate of 1e-3 was used, and each model was trained for a maximum of 150 epochs, although training was stopped early if validation loss didn't decrease in the last 20 epochs The authors ran various experiments with different window sizes (amount of context around a pixel to feed to the CNN) and different band combinations.

Several different outputs were assessed on the test data. Accuracy was fairly high across the board although decreased as the prediction classes got more complex. When comparing whether the CNN simply correctly predicted whether a tree was alive or dead, accuracy was ~95%; when comparing whether a tree was broadleaf or conifer, accuracy was ~83%; and predicting both plant functional type and health status got accuracy around ~80%. CNNs that used a 41x41 pixel window, which generally encompassed the totality of the tree's context, got the best accuracy. Between the four channels available, the exact channels used were found to be less important, although using all four combined generally got the best accuracy. The authors speculate that the fine spatial scale of the imagery is needed for the CNN to achieve good accuracy; at coarser resolutions, the necessary context is lost. The large area covered and the range of vegetation types covered demonstrate the flexibility and generalizability of a CNN for remote sensing tasks.

White, J. C., Wulder, M. A., Brooks, D., Reich, R., & Wheate, R. D. (2005). Detection of red attack stage mountain pine beetle infestation with high spatial resolution satellite imagery. *Remote Sensing of Environment*, *96*(3-4), 340-351.

STUDY AREA: Central British Columbia.

---

[32] VGG is an extremely common CNN architecture type. While the original VGG was one of the earliest CNNs to be created, it is still a common baseline CNN architecture.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Lodgepole pine and mountain pine beetle.

IMAGE SOURCE AND RESOLUTION: IKONOS[33], 4 m resolution.

METHODS USED: ISODATA[34].

TRAINING DATA SOURCE: Unsupervised classes were assigned to be red attacked trees via image interpretation of calibration sites.

VALIDATION DATA SOURCE: Image interpretation (of an aerial image).

OUTPUT CLASSES: Red attack mortality.

SUMMARY: This paper, published in 2005, is notable for being one of the earliest peer-reviewed studies to use high spatial resolution satellite imagery to identify red attack mortality. Areas in the imagery that were not forested were masked from the study. The remaining imagery was clustered using an unsupervised classification algorithm. Calibration stands within the image were then used to identify which classes are associated with tree mortality. Accuracy ranged from moderate to high, with the classification having the most trouble in areas with a low density of mortality.

[33] Launched in 1999 and decommissioned in 2015 IKNONOS is the oldest commercial satellite of higher than 5 m pixel size. It had 4 multispectral bands at 4 m resolution.

[34] An unsupervised classification method similar to K-means. Refines k-means in that it can automatically iteratively merge and split clusters.

3
IDENTIFICATION OF GREEN ATTACK TREES AND TREE STRESS USING HIGH-RESOLUTION IMAGERY

This section focuses on identifying via remote sensing trees that are experiencing some form of stress or have died but the leaves have not yet visibly changed. A summary of the best available science is given first followed by an annotated bibliography for the summary.

Summary Findings

In very controlled environments it's possible to precisely measure a plant's water deficit over time and track how that impacts the plant's reflectances; however, it's a challenge to verify this in the field. In the absence of some other physiological change (leaf loss or discoloration) the spectral impact generally involves subtle shifts among narrow wavelengths, particularly along the red edge of plants and in the near-infrared. Given the spectrally narrow shifts that occur, imaging spectroscopy (wherein the imaging devices detect spectrally narrow regions (5-10 nm) with many channels), is thought to typically better address these problems.

Even with imaging spectroscopy, differences in species, plant physiology, brightness, and mixed pixels all confound efforts to detect stress changes in leaves. Limiting the feature space by isolating only the tree crowns in a separate analysis step is a logical beginning, given the subtle signal of vegetation stress. Species classification is another frequent precursor to assure that the signal of different species does not overwhelm the stress signal. To make stress predictions, vegetation indices are often used as the input features of the machine learning model. One advantage of vegetation indices is they are an easy way to normalize for differences in brightness. There are other valid methods to normalize brightness differences though, such as continuum removal or analyzing the first derivative of the spectral profile.

One line of research in tree stress is detecting green attack trees killed by bark beetles. Practically, this makes for a relatively simple study design since green attack trees can be reliably identified in the field and will reliably change needle color within a year. Still, the reliable separation of green attack trees and healthy trees from a single image remains elusive. Even high spatial resolution imaging spectroscopy, which should be best suited to detect subtle spectral changes from bark beetle attack, has only been shown to have moderate accuracy in separating green and red attack trees. Worldview-2 imagery is only moderately better than chance at separating green attack trees from healthy trees.

Another line of research is detecting trees that are being attacked by an insect or pathogen that causes gradual leaf defoliation or discoloration. Here WorldView-2 imagery can more reliably extract the information, especially when a time series of images can be used to establish a spectral trajectory. Still, imaging spectroscopy is likely best suited to be the inputs to a machine learning model if it is available.

In the case of detecting plant stress at the individual tree level during drought, in the absence of some other disturbance agent, it is unlikely that the spectral resolution WorldView currently offers is likely to be sufficient for ample separation. The literature indicates with multiple passes of high spatial resolution imaging spectroscopy, detection of increasing tree stress through time may be possible. Tracking the trajectory through time can help normalize some of the confounding variables, especially if the imagery is acquired annually at roughly similar times of the year. One suggested approach uses changes in two vegetation indices NDVI (a proxy for the amount of leaf area) and $CI_{red\ edge}$[35] (an indicator of subtle shifts in the red edge of plants known to be correlated with stress) as a promising tracker of stress and recovery. There is currently no satellite capable of detecting at the spectral resolution, spatial resolution, and temporal resolution needed to make this approach work currently. However, given industry trends, it is very likely that an imaging spectroscopy microsatellite will be launched shortly (< 2 years) that would be capable of this analysis[36]. For now, the most common approach is to track a broadband spectral index, generally NDVI, through time and assume negative shifts if chlorophyll levels, as indicated by the index, are indicative of stress. At a coarse level, this may be suitable, but it is likely to miss early stress signals and will not allow for a detailed understanding of what is causing negative NDVI shifts through time.

Annotated Bibliography

Immitzer, M., & Atzberger, C. (2014). Early detection of bark beetle infestation in Norway spruce (Picea abies, L.) using WorldView-2 data. *Photogramm. Fernerkund. Geoinf*, *2014*(351367), 73.

STUDY AREA: Central Austria.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Norway spruce and European spruce bark beetle.

IMAGE SOURCE AND RESOLUTION: WorldView-2, 2 m pixel size.

METHODS USED: Logistic Regression[37] and RF.

TRAINING DATA SOURCE: Field Survey.

VALIDATION DATA SOURCE: Unselected bootstrapped[38] data.

OUTPUT CLASSES: Healthy Tree, Green attack, Red/Gray Attack

SUMMARY: This study uses two WorldView-2 images, one acquired in June and the other in July, to detect green attack and red attacked spruce tree mortality. Trained on-the-ground foresters identified green attack and red attack trees. The crowns of these trees were then

---

[35] A imaging spectroscopy vegetation index with the formula of $R_{Near\ Infrared}/R_{Red\ Edge}^{-1}$.

[36] For one possible example of an upcoming space borne imaging spectroscopy project, see: https://im-mining.com/2021/06/15/orbital-sidekick-goes-above-and-beyond-to-improve-hyperspectral-imaging/

[37] One of the simplest supervised machine learning regression techniques. Given a single feature finds optimal parameters to fit a logit curve to data.

[38] A technique where the data used is repeatedly sampled with replacement from a larger dataset. Since the sampling is with replacement some data points will be left out entirely.

manually delineated in the imagery and the sunlit pixels were selected for training and validation. The samples were then bootstrapped to train an RF and logistic regression model. The methods do not make clear whether individual pixels were sampled in the bootstrapping methods or whole trees, but if some pixels within the same tree crown were allowed to be in the training set and others in the validation set this may be a major source of overfitting in the models. Regardless, using all eight WorldView-2 channels the authors were able to get moderately high separation ~0.75 overall accuracy, however, this accuracy is lifted by the near-perfect separation of the red attack. There wasn't a significant difference in accuracy between the June and July imagery. Logistic regression on individual channels was also performed to determine which channels were best for separating green attack and healthy trees. Blue, green, red, yellow, and the first near-infrared bands were found to be the most important. Overall the results were not accurate enough to operationally separate healthy and green attack trees, but it might be accurate enough to at least find hot spots for further field analysis. Furthermore, the degree the RF classifier would be robust to the shifts in the domain from year to year or in different seasons is also not clear.

Lottering, R., Mutanga, O., Peerbhay, K., & Ismail, R. (2019). Detecting and mapping Gonipterus scutellatus induced vegetation defoliation using WorldView-2 pan-sharpened image texture combinations and an artificial neural network. *Journal of Applied Remote Sensing*, *13*(1), 014513.

STUDY AREA: Eastern South Africa.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Eucalyptus and eucalyptus weevil.

IMAGE SOURCE AND RESOLUTION: WorldView-2, 0.5 m pixel size after pansharpening.

METHODS USED: Texture analysis and NN.

TRAINING DATA SOURCE: Field data.

VALIDATION DATA SOURCE: Partitioned field data.

OUTPUT CLASSES: Percent defoliation.

SUMMARY: This study uses WorldView-2 data to predict defoliation in a eucalyptus plantation. One interesting aspect of this study is it uses texture metrics. The texture metrics are various means of describing the variability (variance, homogeneity, smoothness, etc.) of pixels over a moving window. The window was calculated over several variable sizes running from 5 X 5 pixels to 9 X 9 pixels squares. Vegetation indices were also used to predict defoliation levels. Various NN was trained over 5000 epochs with a learning rate of 0.01 and a momentum of 0.30. NN trained on texture-based parameters were more predictive of defoliation levels in the validation data than vegetation index-based methods (0.90 vs 0.80 R^2 values). NN trained with more texture bands did better than NN with fewer, demonstrating the importance of texture information. Using a NN and texture analysis is fairly similar to using a CNN, however, a CNN would be able to better learn which pattern of textures was most important to it, and therefore would likely be superior.

Murfitt, J., He, Y., Yang, J., Mui, A., & De Mille, K. (2016). Ash decline assessment in emerald ash borer infested natural forests using high spatial resolution images. *Remote Sensing*, *8*(3), 256.

STUDY AREA: Southeastern Ontario, Canada.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Ash trees and emerald ash borer.

IMAGE SOURCE AND RESOLUTION: WorldView-2 imagery, 0.5 m pixels size after pansharpening.

METHODS USED: Imaging spectroscopy was used to capture pure crown spectra, these spectra were then used to simulate the spectral impacts of chemical changes associated with vegetation decline. The best vegetation indices for tracking the vegetation decline through time were identified and then tested on a time series of imaging spectroscopy data.

TRAINING DATA SOURCE: Field observation (57 trees).

VALIDATION DATA SOURCE: Field observation (29 trees).

OUTPUT CLASSES: Percentage decline.

SUMMARY: WorldView-2 was used to first delineate tree crowns, then classify tree species, and finally predict the impact of emerald ash borer in a mixed hardwood forest. Only four bands of the multispectral data were used to simulate an aerial image. The combination of crown delineation and classification had an overall accuracy of 63.3%. Several multi-spectral indices were then calculated and correlated with field observed decline levels. Overall, NDVI was picked as the best vegetation index with a 0.38 R^2 value when compared with field measurements, although the correlation was best at lower infestation levels. In contrast, a similar study using imaging spectroscopy data to predict emerald ash borer with similar methods was able to achieve a 97% accuracy[39]. Other studies[40] using WorldView-2 were able to achieve higher correlations with field observations for stress, primarily using the red edge channel.

Navarro-Cerrillo, R. M., Varo-Martínez, M. Á., Acosta, C., Rodriguez, G. P., Sánchez-Cuesta, R., & Gomez, F. J. R. (2019). Integration of WorldView-2 and airborne laser scanning data to classify defoliation levels in Quercus ilex L. Dehesas affected by root rot mortality: Management implications. *Forest Ecology and Management*, *451*, 117564.

STUDY AREA: Southwestern Spain.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Holm oaks and root rot.

---

[39] Pontius, Jennifer, Mary Martin, Lucie Plourde, and Richard Hallett. "Ash decline assessment in emerald ash borer-infested regions: A test of tree-level, hyperspectral technologies." *Remote Sensing of Environment* 112, no. 5 (2008): 2665-2676.

[40] Other studies in the annotated bibliography show this as well, but specific to ash and WorldView-2 see: Waser, Lars T., Meinrad Küchler, Kai Jütte, and Theresia Stampfer. "Evaluating the potential of WorldView-2 data to classify tree species and different levels of ash mortality." *Remote Sensing* 6, no. 5 (2014): 4515-4545.

IMAGE SOURCE AND RESOLUTION: WorldView-2 0.5 m pixel size after pansharpening and airborne LIDAR (12 points per square meter).

METHODS USED: RF.

TRAINING DATA SOURCE: Field sampling of trees.

VALIDATION DATA SOURCE: Both cross-validation and partitioned field data.

OUTPUT CLASSES: Low Defoliation and high defoliation.

SUMMARY: A combination of WorldView-2 and LIDAR data was used to identify defoliation as a result of root rot. The LIDAR data were used both to segment tree crowns as well as to generate canopy structure metrics. Vegetation indices were generated from the WorldView-2 images. Both sets of information were fed into an RF classifier to predict defoliation levels of trees (as determined by field sampling). Separation of the low and high defoliation classes was fairly successful (around 85% overall accuracy). The model performed best when given predictive parameters both from WorldView-2 and the LIDAR. From WorldView-2, spectral indices related to leaf area levels, leaf senescence, and leaf chlorophyll content were most important for predicting defoliation. LIDAR metrics that were important directly measured canopy structure. The relative location of the vegetation within the canopy, especially whether there was a high percentage of the canopy near the top of the crown, was particularly important.

Niemann, K. O., Quinn, G., Stephen, R., Visintini, F., & Parton, D. (2015). Hyperspectral remote sensing of mountain pine beetle with an emphasis on previsual assessment. ***Canadian Journal of Remote Sensing*****,** ***41*****(3), 191-202.**

STUDY AREA: Central British Columbia.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Lodgepole pine and mountain pine beetle.

IMAGE SOURCE AND RESOLUTION: Airborne imaging spectrometer (AISA Dual imaging spectrometer) 2m pixel size.

METHODS USED: Treetops were isolated using a LIDAR-derived digital canopy model. Continuum removal and the first derivative of spectra were extracted from healthy and green attack trees. A nonparametric ANOVA was then run on both spectra sets.

TRAINING DATA SOURCE: Visual interpretation of multi-year imagery.

VALIDATION DATA SOURCE: N/A

OUTPUT CLASSES: Healthy, green attack, red attack.

SUMMARY: An airborne imaging spectrometer was flown over a lodgepole pine plantation in September 2008 and July 2009. A total of 81 spectra were extracted from the 2008 imagery representing healthy, green attack, and red attack trees. Green attack trees were selected by identifying trees that are visually healthy in 2008 but had transitioned to red attack in 2009. To normalize brightness differences, two analysis techniques were separately run on the spectra: continuum removal and derivative analysis. Both techniques were then examined for statistical

significance. Although the imaging spectrometer sampled between 400 nm and 2400 nm only the visible to near-infrared (400-1300 nm) showed statistical differences between healthy vegetation and green attack vegetation. In particular, a blue shift of the red edge (around 680 nm) was noted in green attack trees versus healthy trees. Several wavelengths noted to be associated with water absorption in vegetation (around 940 nm and 1250 nm) were also statistically significant in their ability to separate green attack and healthy vegetation. These changes were detectable regardless of the normalization technique used. These features are all too spectrally narrow to be detectable by a multispectral sensor such as WorldView.

Oumar, Z., & Mutanga, O. (2013). Using WorldView-2 bands and indices to predict bronze bug (Thaumastocoris peregrinus) damage in plantation forests. *International Journal of Remote Sensing*, *34*(6), 2236-2249.

STUDY AREA: Eastern South Africa.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Eucalyptus plantations, and bronze bug.

IMAGE SOURCE AND RESOLUTION: WorldView-2, 0.5 m pixel size after pansharpening.

METHODS USED: Partial least squares regression (PLS)[41] on bands and vegetation indices.

TRAINING DATA SOURCE: Field sampling of the tree damage amount.

VALIDATION DATA SOURCE: Same as the train (70% of trees randomly assigned to train 30% to validation).

OUTPUT CLASSES: Predicted percentage damage

SUMMARY: In this paper, the impact of the bronze bug, a sap-sucking insect was examined. Bronze bug infestations generally weaken trees and can cause leaf discolorations. Intensive field surveys were done to identify the percentage of necrotic tissue on 50 mature leaves for each tree. PLS regression was then done against various combinations of bands and spectral indices to predict the average percentage damage of leaves. Variable importance was extracted from the PLS regression to identify which parameters were most important in the prediction. When PLS was run on the feature set of the sensor bands, relatively high correlation scores were achieved ($R^2 = 0.63$), whereas just the vegetation indices got relatively weak correlations ($R^2 = 0.42$). Combining the vegetation indices and bands with the highest variable importance score performed the best ($R^2 = 0.65$). Overall channels near the red edge and near-infrared were most predictive of damage. Vegetation indices correlated with pigment change especially carotenoids

[41] A supervised regression machine learning technique. Partial least squares reduces the dimensionality of the feature space focusing on the collinearity between features and then performs ordinary least squares regression on the reduced dimensionality space.

and anthocyanins[42], and the reddening of vegetation (RGI) were most predictive of bronze bug impact.

Zarco-Tejada, P. J., Hornero, A., Hernández-Clemente, R., & Beck, P. S. A. (2018). Understanding the temporal dimension of the red-edge spectral region for forest decline detection using high-resolution hyperspectral and Sentinel-2a imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, *137*, 134-148.

STUDY AREA: Western Spain.

TARGET VEGETATION TYPE AND DISTURBANCE AGENT: Maritime pine plantation and drought stress.

IMAGE SOURCE AND RESOLUTION: Airborne imaging spectroscopy, 40 cm pixel size.

METHODS USED: Imaging spectroscopy was used to capture pure crown spectra; these spectra were then used to simulate the spectral impacts of chemical changes associated with vegetation decline. The best vegetation indices for tracking the vegetation decline through time were identified and then tested on a time series of imaging spectroscopy data.

TRAINING DATA SOURCE: Simulation.

VALIDATION DATA SOURCE: Field sampling.

OUTPUT CLASSES: Correlation with field-identified decline classes.

SUMMARY: High spatial resolution imaging spectroscopy was flown three times (January 2015, February 2015, and July 2016) over a pine plantation that was experiencing declining health. Concurrent with this, four extensive field surveys were conducted to establish the trajectory of the pine health in individual trees. The high-resolution imagery allowed for pure vegetation pixels to be extracted for various levels of stress and decline. Simulations showed temporal trends in plant decline over time particularly associated with the NDVI (a stand-in for leaf area index) and the $CI_{red\ edge}$. Both vegetation indices declined in value concurrent with tree health decline. These trends were then verified to be identifiable and predictive of field observed vegetation decline both for the airborne imaging spectroscopy and medium spatial resolution satellite imagery (Sentinel-2). The spectral resolution to calculate $CI_{red\ edge}$ is not currently available in commercial fine resolution satellite imagery.

[42] Carotenoid reflectance index (CRI) is a vegetation index associated with carotenoids, a leaf pigment. It was proposed as an imaging spectroscopy vegetation index as: $(1/R_{510}) - (1/R_{700})$. A related vegetation index, the anthocyanin reflectance index, has the equation of $(1/R_{550}) - (1/R_{700})$. In this study, WorldView band 3 was used for the 510 and 550 nm wavelength and band 7 was used for the 700 nm wavelength.

4
EVALUATION OF COMMERCIAL VENDORS

This section details some commercial vendors active in the space of using high resolution remote sensing for utility vegetation management purposes. This is not a comprehensive summary of the vendors in this space but a representative sample. Each vendor was contacted and a company representative interviewed. Where possible additional sourcing from the company web site or other relevant sources was added. Companies are listed in alphabetical order.

AiDash

AiDash "is an AI-First vertical SaaS company enabling satellite-powered operations and maintenance in core industries with distributed assets."[43] From examples highlighted on its website, a core focus is on electric utility vegetation management, along with storm, wildfire, and extreme weather event risk mitigation. AiDash maintains an web and mobile app to deliver vegetation management solution for utility companies and third party vegetation management companies. AiDash claims 25 existing utility customers in the United States, Canada, the United Kingdom, and Australia (including National Grid, Avista Utilities, and an unnamed Fortune 500 power utility).[44,45,46] Several AiDash webpages claim customers can reduce vegetation management costs by 20% and improve reliability by 15%.[47] [48] [49] Their technology can identify dead and dying vegetation in the strike zone of power lines. They can also identify healthy trees within striking distance, which, based on pattern analysis of tree fall in the area, are at risk of falling.

The exact nature of the machine learning algorithm or the remote sensing technology being used was not publicly available. The remote sensing used was entirely spaceborne, and the highest resolution of their product was 30 cm pixel size. Their product could be updated with new remote sensing rapidly, if needed. The company notes that their product has been validated extensively with previous projects with partners and customers.

---

[43] https://www.aidash.com/our-company

[44] https://www.aidash.com/resource/aidash-secures-6m-series-A-funding-led-benhamou-global-ventures-and-national-grid-partners, and https://www.aidash.com/resource/aidash-transforming-vegetation-management-national-grid-using-satellite-analytics-and-ai

[45] https://www.aidash.com/resource/webinar-wildfire-and-risk-mitigation-using-satellite-analytics-avistas-journey

[46] https://www.aidash.com/resource/how-aidash-transformed-vegetation-management-fortune-500-company

[47] https://www.aidash.com/resource/aidash-launches-intelligent-vegetation-management-system-distributech-2020

[48] https://www.aidash.com/resource/aidash-launches-2021-vegetation-management-best-practices-guide-electric-utilities

[49] https://www.aidash.com/resource/decoding-intelligent-aidash-intelligent-vegetation-management-system

LiveEO

LiveEO[50] focuses on vegetation management of linear infrastructure using satellite data and artificial intelligence. The company has partnered with major utility companies from across the world including Germany's E.DIS, and other European and U.S. utilities.[51] The company has also worked with Deutsche Bahn, Europe's largest railway company.[52] LiveEO was also selected to participate in EPRI's Incubatenergy initiative that supports technology demonstrations with U.S. electric utilities.[53] LiveEO maintains an app that provides an interface point between remotely sensed alerts, utility companies, and third-party vegetation management services. LiveEO tracks species, height, and tree health among other metrics. Tree health is monitored via tracking NDVI through time. This combined with species information, gives an idea of chlorophyll content and leaf area through time.

Beyond the use of vegetation indices, LiveEO also uses neural networks to perform semantic segmentation of species and generates warnings. LiveEO's models have been run over extremely large scales, including an analysis of vegetation risk across the entire US's electric grid.[54]

LiveEO uses a constellation of space-borne satellites and archival imagery of various resolutions to perform their analysis. The exact makeup of the constellation was not given, but the finest resolution they use is 30 cm. The return interval for analysis can be as rapid as weekly, but the exact revisit time interval should be balanced with project needs and cost-efficiency.

Previous validation has been done in collaboration with customers who collected on the ground data. Previous results are iteratively fed back into the algorithm for continuous improvement. LiveEO claims 95% accuracy in tree detection for individual trees and above 99% accuracy for groups of trees.

Salo Sciences

Salo Sciences' mission is "to accelerate and expand the impact of conservation with data & analytics."[55] Although Salo Sciences'[56] focus is on using remote sensing for ecosystem conservation and restoration, their products can be useful for utility vegetation management. As part of a project funded by the U.S. Department of Agriculture, the company mapped tree mortality across the state of California.[57] They are also currently collaborating with a large California electric utility to assist in its wildfire mitigation plan. Salo Sciences outputs tree mortality maps at 10 x 10 m resolution and 3 x 3 m resolution.

Salo Sciences' mortality mapping is based on convolutional neural networks (CNNs). They have customized CNNs to take in all channels of the multispectral data they consume. CNNs are

---

[50] https://live-eo.com/

[51] https://live-eo.com/liveeo-launches-grid-wide-vegetation-monitoring-for-e-dis/

[52] https://www.deutschebahn.com/en/Digitalization/startups/db_startups/LiveEO-3520728

[53] https://live-eo.com/epris-incubatenergy-labs-selects-liveeo-as-a-2021-cohort-awardee/

[54] https://live-eo.com/case-study-entire-us-power-transmission-grid-analyzed/

[55] https://salo.ai/about

[56] https://salo.ai/

[57] https://salo.ai/projects/identifying-tree-mortality

trained from scratch on remotely sensed images which should make the CNNs highly adept at processing remotely sensed textures. Salo Sciences has evaluated its mortality product against self-collected field measurements and found it to have high accuracy.[58] Currently, the company does not have a vegetation stress product, but they have collaborated with researchers on mapping plant moisture stress in the past and are looking for collaborators to develop partnerships on a moisture stress product.

Salo Science's product consumes images from Sentinel-1, Sentinel-2, and images from Planet's microsatellites in a custom multi-sensor approach. Project areas could be re-scanned as rapidly as around a week, but the desired return interval should be evaluated on an individual basis considering cost and needs.

Satelytics

Satelytics provides "geospatial analytics for early detection, location, and — in many instances — quantification of" risks associated with company assets.[59] The company websites notes solutions for multiple sectors, including providing electric utilities vegetation management and corridor encroachment services. Aspects of vegetation that are monitored include whether the vegetation is live or dead, vegetation health (primarily determined by differences in chlorophyll reflectance), species, and tree height. Satelytics has partnered with a "major western power utility company," Duke Energy, and ACRT (a utility vegetation management consulting company), for monitoring vegetation along power lines. [60,61]

Satelytics' system can send actionable alerts to consumers via app or API to whatever system the customer chooses. Their system also provides an interface between remotely sensed data, alerts, utility companies, and third-party contractors. The bulk of their analysis is done via CNN. Analysis results can be displayed via the app. The app also allows for viewing changes through time and communicating directly to contractors the exact location and nature of vegetation that may need removal.

A variety of inputs are used for monitoring purposes including Satellogic's satellite constellation, and Maxar's WorldView-3 satellite; however, a lion's share of the vegetation monitoring is done by the Pleiades constellation of satellites (distributed by Spot Imagery; 4 channel RGB and infrared 2 m imagery pan-sharpened to 30 cm). Imagery can be updated as rapidly as a daily basis, but exact revisit time is determined by customer needs and cost-efficiency.

Satelytics' training and validation data have been assembled via previous projects. The company notes that they have sampled around 200 different species and their various results have been validated through previous client-confidential work. Project validation errors are iteratively fed back into the algorithm for a continuously improving analysis.

---

[58] https://salo.ai/blog/2018/09/completed-snf-sampling, and personal communication, May 2021

[59] https://www.satelytics.com/

[60] https://www.satelytics.com/why-satelytics#testimonials

[61] https://acrt.com/articles/acrt-services-partners-with-satelytics-to-bring-geospatial-analytics-to-utilities/

UP42

Unliked the other companies analyzed in this section, which offer an analysis product or service, UP42 provides "a developer platform and marketplace, offering all the building blocks for powerful, scalable geospatial products."[62] The "building blocks" include a marketplace for remotely sensed imagery, and machine learning and other developer tools. The marketplace provides access to satellite and aerial imagery at a variety of different spatial resolutions, wavelengths, and aquisition intervals. Satellite-based imagery can be as fine in resolution as the Pleiades sensor (50 cm pixel size). After the desired imagery is selected, UP42 offers tools such as APIs and software development kits (SDKs), and on-demand cloud-based infrastructure to process the raw imagery from different sources into analysis-ready products. UP42 also distributes various algorithms that extract meaningful insights from the imagery.

From the perspective of utility vegetation management, relevant algorithms include those related to tree location, species, height, and plant health. UP42 does conduct a prior quality analysis of the algorithms. At the same time, UP42 encourages a rating system for users, who evaluate each processing block in terms of performance (high, medium, low). Additionally, it is possible to upload custom-made algorithms created by individual users to build the current collection of algorithms. UP42 provides a great deal of flexibility and potential cost efficiencies. However, the customer should have a familiarity with the geospatial industry and solid technical understanding of their requirements and how to accomplish them.

[62] https://up42.com/

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI members represent 90% of the electricity generated and delivered in the United States with international participation extending to nearly 40 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; Dallas, Texas; Lenox, Mass.; and Washington, D.C.

Together…Shaping the Future of Energy™

© 2021 Electric Power Research Institute (EPRI), Inc. All rights reserved.
Electric Power Research Institute and EPRI are registered marks of the Electric Power Research Institute, Inc. in the U.S. and worldwide.

3002022770

EPRI
3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 • USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com