Mapping California woodland-chaparral ecosystems following wildfire with diverse drone images and computer vision

David J. Russell[1*], Amritha S. Pallavoor[1], Gary M. Bucciarelli[2], Andrew M. Latimer[1,2], and Derek J. N. Young[1]

[1]Department of Plant Sciences, University of California Davis

[2]Institute of the Environment, University of California Davis

[*]Corresponding author: djrussell@ucdavis.edu

Open Research: All raw and processed data are available on Open Science Framework. The code for preparing, analyzing, and visualizing data is available via Zenodo.

Keywords— California, chaparral, woodland, wildfire, disturbance, recovery, drone, computer vision

Abstract

Fire is a key driver of vegetation dynamics in California's woodland-chaparral ecosystems, and its role has become ever more important in recent decades as wildfire extents and frequencies increase. Understanding post-fire vegetation transitions and the likelihood of type conversion is essential for effective land management. Remote sensing represents a powerful tool to map vegetation cover and study post-fire dynamics, but current approaches—generally based on satellite sensors—are limited by their spatial and temporal resolution and generally broad application extents. In contrast, uncrewed aerial vehicles, or "drones," offer great potential to yield low-cost, high-resolution, locally-tailored data on vegetation cover and its variation across time and space. With recent rapid development of technologies for translating raw drone imagery into ecologically relevant data, the power of drone-based research is increasing along with its analytical decision space. In this work, we apply modern methods in image processing and computer vision to generate vegetation maps from a large and diverse dataset of drone images collected under realistic operational constraints. Specifically, our imagery was collected at three study sites across three years, by multiple pilots flying different drone models with varying flight parameters. Our analytical approach uses an automated method to spatially co-register all overlapping datasets into a common reference frame. We then generate vegetation predictions within each raw image using a computer vision model and translate image-level predictions to a geospatial map based on the known positions of the drone camera. Finally, we unify all geospatial predictions from similar dates into a best available prediction for each location. Using this merged representation, we conduct change analysis across years for the landscape area common between years—approximately 100 ha at each of two study sites. When predicting our eight vegetation classes on unseen images, we achieved 94% overall accuracy and 88% class-balanced accuracy. Change analysis yielded surprisingly little change over 3-4 years post-fire, with key changes being shrub (re)establishment and tree resprouting. Our findings demonstrate the viability of scalable drone-based approaches for tracking vegetation change in fire-prone landscapes.

1 Introduction

Understanding how ecosystems respond to wildfires and drought is critical to gauging and ensuring the persistence and functional integrity of unique ecological communities. This is especially relevant in Mediterranean-climate ecosystems, many of which are simultaneously biodiversity hotspots and climatic change hotspots. One manifestation of the disproportionate effects of climate change has been fire regime shifts, including changes in fire frequency, severity, and behavior. For example, global instances of extreme fire events have increased 2.2-fold since 2003 Cunningham et al. (2024) and wildfire seasons have lengthened while fire intervals have shortened Moritz et al. (2012); Jolly et al. (2015); Fairman et al. (2019); Whitman et al. (2019). It is predicted that climate change will further modify wildfire behavior through the remainder of the century, with the number of days conducive to extreme wildfire events increasing between 20-50% in fire-prone regions Bowman et al. (2017).

The outsized effects of climate on fire have been particularly strong in Mediterranean-climate regions of the world, including large parts of California and the Pacific Northwest in the United States, especially the chaparral and woodland areas of California. In 2020, California lightning-ignited wildfires burned 4.2 million acres across Northern and Central California. The three largest of these wildfire events alone – the August Complex (Northwestern mountains), SCU Lightning Complex (Central Coast), and LNU Lightning Complex (Northern Coast Range) – burned nearly 2.0 million acres. A large factor driving ignition of these wildfires was a series of storms that resulted in up to 15,000 dry lightning strikes, causing numerous, small wildfires that under highly favorable wildfire conditions progressed to extreme or megawildfires.

With so much recently burned land, there is an increasingly urgent need for data on how the ecosystems are responding to fire, both for mapping and predicting outcomes and for informing management. In smaller wildfires, it can be feasible to capture representative vegetation responses using traditional field sampling, e.g. the 23,000 ha Storrie Fire Crotteau et al. (2013). In recent megafires of 10^5 - 10^6 acres, adequately sampling across vegetation types and spatial and environmental gradients will require scaling up data collection through complementing intensive field sampling with broad-scale, high-resolution remote sensing. Extensive data collection is especially needed in fire-prone Mediterranean climate vegetation types, including grasslands, shrublands, and oak woodlands, where there has been less previous work on vegetation response to fire. These ecosystem types are typically mosaics of patches with slowly-shifting boundaries, and en-

compass "climate change refugia" potentially important for buffering species from climate change effects in California Thorne et al. (2020). Oaks generally have high resilience to fire because they can usually can resprout even when entirely top-killed Holmes et al. (2008). Wildfire can, however, gradually reduce oak woodland cover and convert to grassland Callaway and Davis (1993). Key open questions about postfire response of these vegetation types include how rapidly and consistently oak cover will bounce back after fire, and to what extent smaller-scale factors including topography, tree or shrub size, and surrounding vegetation affect post-fire resprouting vs. mortality in oak woodland.

Geospatial remote sensing datasets – especially maps of spectral reflectance measured by satellite – are increasingly being used to infer variation in vegetation cover across space and through time Allred et al. (2021); Cingolani et al. (2004). For instance, the U.S. National Land Cover Database produces annual vegetation cover maps at 30 m spatial resolution using each pixel's spectral reflectance as measured by Landsat, a series of U.S. public spectral imaging satellites Wulder et al. (2022). However, a number of factors constrain the accuracy and degree of granularity–both spatial and in terms of vegetation classes–that can be achieved by satellite. One key factor is spatial resolution: public spectral imaging satellites producing freely available data relevant to vegetation mapping are not available at resolutions finer than 10 m. In the U.S. and some other countries, national programs such as the National Agriculture Imagery Program (NAIP; U.S. Department of Agriculture (2011)) use piloted aircraft to collect finer-resolution imagery over large areas (e.g., 60 cm in the case of NAIP), but at relatively low frequency (e.g., 2-3 years in the case of NAIP), limiting their utility for rapid response following disturbance.

In comparison to satellite and broad-scale aerial imaging programs, uncrewed aerial vehicles (UAVs, or "drones") with imaging sensors have several distinct advantages for inferring post-disturbance vegetation cover and change Anderson et al. (2025). First, imagery can generally be collected when it is needed, including immediately following disturbance. Second, the resulting imagery generally has a much higher spatial resolution (e.g., 1-10 cm for many popular low-cost commodity drones; Chang et al. (2025)). Imagery at this resolution allows one to resolve plant structural components such as stems and leaves and to infer stature and shape from shadows and other textural cues. This greatly increases the potential for detailed and accurate classification of vegetation cover types. Drone-based image collection generally involves capturing a dense grid of highly overlapping images, which are then "stitched" together and georectified using a process called "photogrammetry" Young et al. (2022); Mlambo et al. (2017). The set of overlapping raw images thus

contains multiple (often many) views of any given object, each from different angles, and predictions on raw images instead of orthomosaics enables much greater classification accuracy Russell et al. (2024).

To take advantage of such detailed imagery to infer land cover, researchers are increasingly employing computer vision (CV) techniques driven by artificial neural networks Brodrick et al. (2019); Detka et al. (2023). Specifically, these methods rely on "deep learning", where multiple sequential processing steps with tunable parameters are applied to translate imagery inputs into classifications such as "tree" and "shrub". Of particular relevance to vegetation cover mapping is the image processing task known as "semantic segmentation" Csurka et al. (2023), in which each pixel is assigned a classification. Relative to conventional segmentation approaches, many popular CV-based segmentation methods are particularly skilled at incorporating the context of the pixels nearby each focal pixel, providing texture and structural cues which help differentiate pixels which may be indistinguishable by color alone.

The availability of drone-derived imagery of natural landscapes is continually growing OpenAerialMap (2025); Young (2025). Land stewardship institutions like the University of California Natural Reserve System have increasingly embraced the use of drone data across broad areas as a means to understand ecosystem dynamics, often piecing together multiple independent data collection campaigns to develop a more holistic representation of a landscape. Given the increasing availability of drone imagery across disparate platforms, dates, and collection protocols, automated vegetation mapping approaches that can accommodate diverse data sources are becoming increasingly necessary, and successful implementations may have increasingly great power for advancing ecological understanding. In this study, we evaluate the potential for using computer vision applied to raw drone images–and thus multiple views of each object–to infer vegetation cover following large wildfires that burned in a California grassland-chaparral-woodland ecosystem and to evaluate the vegetation trajectory and its drivers over a four-year period following fire. Our project leveraged imagery collected over 3 years by 2 different initiatives employing numerous pilots, aircraft, imagery collection dates, and flight parameters to develop standardized maps of predicted land cover with the goal of informing ecological understanding of post-fire vegetation dynamics. In this paper, ask how effective and accurate these novel computer vision methods are in the context of a large-scale monitoring projects. We use the resulting vegetation predictions to analyze vegetation cover change over the first 3-4 years after wildfire. Specifically, we ask what proportion of oak woodland and chaparral areas recover via resprouting after initial fire-induced dieback.

2 Materials and Methods

2.1 Study Sites

We focused on three key sites that were impacted by the dry lightning-ignited wildfires of 2020. These sites are a part of the protected network of lands and fieldstations in the Univeristy of California Natural Reserve System (UCNRS). Each is managed primarily to provide access for research and class use, with essentially no management of wildland habitat. The three sites include Blue Oaks Ranch Reserve (BORR), Hastings, and Quail Ridge and are primarily a mix of oak woodland and grassland habitats. BORR is a 3,259 acre site located in Santa Clara County, Hastings is a 2,373 acre site located in Monterey County, and Quail Ridge is a 2,500 acre site located in Napa County. BORR burned as a part of the SCU Lightning Complex Fire, Hastings burned during the River Fire, and Quail Ridge burned burned in the LNU Lightning Complex Fire. Severity of fire differed within and between these sites, as did the total area burned at each site.

2.2 Drone imagery collection and photogrammetric processing

Data was obtained from 60 drone flights conducted between February 2020 and May 2024. All the imagery used in this study was RGB (red, green, and blue channels of visible light) and was collected in a highly-overlapping manner. The following platforms were used: AgEagle eBee X, Autel Evo II v2, DJI Matrice 100, DJI Matrice 210 RTK, DJI Matrice 300, and DJI Mavic 3 Multispectral (using the M3M RGB camera). Comprehensive flight planning metadata was only provided by the pilots for 12/60 flights, and for those cases all data was collected in a nadir orientation, with 75-80% front and side overlap with nominal flight altitudes ranging from 91-122m. Most flights were conducted at a fixed altitude above sea level, meaning that the altitude above ground varied substantially across different images for terrain with variable elevation.

Once the data has been collected the individual images for a given flight need to be registered together. Structure from motion is a technique that identifies visually-distinct features across multiple images and uses these correspondences to estimate the structure of the scene as well as the locations and orientations from which each image was taken. We use the Agisoft Metashape Agisoft, LLC. (2025) software to perform structure from motion. We use the parameters identified by Young et al. (2022) that perform well in forested environments using similar flight parameters to our data collection. Each drone flight was processed

Class Name	Description	Fraction of labeled pixels
Bare Earth	Any ground where no plants are growing and only dirt is present. Includes ash, fire-killed and dead herbaceous.	0.3712
Herbaceous Live	Green fields or any areas with grass or low-growing plants.	0.1199
Artificial Object	Any artificial structure or object like buildings, vehicles, etc.	0.0017
Shrub Dead	Any shrubs that are clearly gray. Includes fire-killed shrubs.	0.0168
Shrub Live	Any green shrub, growing low and in clusters. Includes resprouting and dormant shrubs.	0.0879
Tree Dead	Any tree that has brown leaves or stripped branches. The branches and trunk may be gray or brown colored. Includes fire-killed trees.	0.1001
Tree Live	Any tree that is any shade of green.	0.2870
Water	Any river or clear body of water.	0.0153

Table 1: Description of vegetation cover classes.

independently to produce the highest-quality data and reduce spatial distortion from time-varying GPS bias.

2.3 Computer vision model training and inference

To classify vegetation cover from high-resolution drone imagery, we developed a deep learning based semantic segmentation pipeline. The model was trained on eight vegetation cover classes, as described in table 1. We began by selecting a representative set of images for model training. Each of the three reserves received roughly the same number of images. The years were split unevenly, with 50%, 25%, and 25% allocated to 2020, 2023, and 2024 respectively to account for the similarities between the later two years. Images were spatially distributed across the landscape by using K-means clustering to divide all image locations into a fixed number of clusters and then annotating a central image from each cluster. Due to a fixed annotation budget, not all clusters were labeled, but they were done in a spatially-stratified manner such that annotations were well distributed across the landscape.

Annotation was conducted by domain experts, the majority of whom had performed field work at at least one study site. They used VIAME Dawkins et al. (2017), an open-source image annotation tool, to manually label polygons delineating regions corresponding to each vegetation class. We did not attempt to label every pixel, since labeling at the boundaries of spatially-intermixed vegetation classes is challenging and laborious. Instead, we chose to label only the central region of homogeneous regions of one vegetation class with coarse polygons, since prior work Davila et al. (2022) has shown that this more efficient style of annotation can be used to train models that still produces accurate predictions even at class boundaries. The annotated polygons were rasterized into single-channel segmentation masks. The annotated masks and labels were randomly split into 80-20% as training and validation data.

We selected SegFormer Xie et al. (2021) as our deep learning model. SegFormer is a state-of-the-art semantic segmentation model that combines the powerful feature extraction capabilities of transformers with a lightweight and efficient decoder. Also, the authors demonstrated this this model is robust to common image degradation (ex. blur or noise) that might dramatically decrease performance of other approaches.

We used the implementation from the MMSegmentation Contributors (2020) framework, which is built on PyTorch and optimized for training segmentation models at scale. The default parameters were used unless otherwise stated. The SegFormer-B5 variant was initialized with Imagenet-pretrained weights and fine-tuned on our annotated dataset. The training was done on images of size 1024 x 1024. Data augmentations included random resize, random crop, random flip and controlled photometric distortion. Augmentation techniques that significantly altered the color of the imagery were avoided to preserve the color distinctions between different types of vegetation cover. Mean and standard deviation values were computed from the dataset to normalize the input images. For improved generalization, we also implemented test-time augmentation (TTA) in the pipeline, which included multi-scale inference and horizontal flipping.

The AdamW optimizer, configured with a learning rate of 6e-5 was used to train the model. Training began with a linear warm-up phase for the first 1500 iterations, with a start factor of 1e-6, followed by a polynomial learning rate decay. Cross-entropy loss (without sigmoid activation) was used as the loss function. A total of 10000 iterations were run, and validation was performed every 1000 iterations. The SegFormer model was trained on a partial NVIDIA A100 GPU with 20 GB of VRAM and 16 CPU cores. Both training and validation were performed with a batch size of 1. The end-to-end training process took approximately 1.5 hours.

2.4 Projecting image-based cover predictions into geographic space

The majority of methods for predicting land use/land cover from drone imagery generate predictions from the top-down orthomosaic derived from photogrammetry Lu et al. (2023); Bellia and Lanfranco (2019); von Nonn et al. (2024); Detka et al. (2023). Recent research has demonstrated that higher accuracy can be achieved by instead generating predictions on individual drone images and projecting these predictions into geospatial coordinates Liu and Abd-Elrahman (2018); Russell et al. (2024). We use an open source software Geograypher Russell et al. (2025) for this task. This approach uses the 3D mesh model and precise camera locations and orientations derived from photogrammetry to determine correspondences between each location in the image and the 3D structure of the scene. Importantly, this approach is robust to oblique views, terrain relief, and inter-object occlusion. We use Geograypher to project the segmentations from the model described in the previous section onto the mesh, and then take the most commonly predicted class for each location on the mesh, across all images that observed that location. Since multiple views observe each location, the agreement of predictions across views gives us a metric for prediction confidence. We exclude locations on the mesh where more than 20% of views disagree with the most commonly-predicted class from subsequent ecological analysis because these predictions are less likely to be correct.

The 3D representation of the data must be converted to a more traditional 2D geospatial format to conduct ecological analysis. The 3D geometry is represented by a mesh data structure, consisting of triangular faces which are each annotated with a single class. Since the mesh is geospatially referenced, the 2D, top-down geospatial projection of each of these faces can be easily computed. Then, all faces with the same class are merged together into multi-polygon data structures. Since the mesh has a high spatial resolution, these vector data products have correspondingly-detailed boundaries. To speed up subsequent operations, the detail of the boundary is reduced using morphological operations and simplification. In rare cases, two faces with different class labels have a 2D projection that overlaps, such as under an overhanging tree. In these cases, the class which is rarest across the entire scene is retained and the overlapping region is removed from the representation of the dominant class. Finally, the data is cropped to the boundaries of the flight polygon, to exclude low-quality predictions at the boundary of the observed region.

2.5 Data registration across drone flights

Our data consists of 60 distinct drone flights which must be spatially registered together. Notably, the majority of flights were conducted without RTK corrections, and there were no fixed ground control points present year-to-year. Using GPS data alone to georeference photogrammetry products often results in several meters of registration error Sohl and Mahmood (2024). Precise registration between different datasets is especially important for environments like ours where multiple classes are spatially intermixed at a relatively fine scale.

To automatically register the flight datasets together, we first identify pairs of flights that overlap more than a threshold area of 2,500 m^2. Then we attempt to find the best translation-only registration between the two datasets using geometric features. Specifically, we compute a canopy height model (CHM) using the digital surface model (DSM; top of canopy) and digital terrain model (DTM; ground estimate) produced by the photogrammetry software. Then we use an iterative registration algorithm implemented in the ITK medical image registration software Ibanez et al. (2025) to find the translation that minimizes the mean squared difference between the CHM models for the two datasets.

This procedure generates relative shifts for multiple pairwise dataset correspondences. Since these pairwise shifts may not suggest a consistent absolute shift for each dataset, we use a least-squares formulation to find the optimal shift for each dataset. This has a tunable weight that balances respecting the pairwise shifts determined by the CHM registration with maintaining the initial location of each dataset. The optimized shift for each dataset is applied to register all datasets into a common reference frame.

2.6 Summarizing predicted vegetation cover

To develop unified maps of "best-available" inferred vegetation cover across our drone-surveyed areas, we generated a "flattened" layer for two time points: 2020 and 2023/2024 (combined), which we hereafter refer to as "early" and "late", respectively. Within a given time point, we had multiple potentially partially overlapping drone footprints, each with a subset of regions predicted with high enough confidence to be included in analysis. For each location, we computed how many drone missions predicted each class. If the number of predictions for a given class achieved a plurality of votes, that class was assigned to the corresponding location in the flattened dataset. In the case of a tie (e.g. two datasets predicted class Tree Dead and two predicted Shrub Dead) the class which was least common across all datasets was chosen, to

ensure that rare classes were well represented in the final product.

After preliminary evaluations, we determined that it was too challenging to differentiate live vs. dead trees during leaf-off conditions. As such, we restricted our analysis to drone flights collected during the period we could conservatively assume leaves would be present in this ecosystem, between April 15th and November 1st. This left us with 27 out of 60 datasets.

Depending on phenology, herbaceous vegetation cover can appear live or dead, and it can also appear as bare ground in low-biomass areas when it senesces. Due to the high reproductive output and effective dispersal of many of the herbaceous species common in our study sites (primarily invasive grasses), bare ground almost always supports herbaceous vegetation during the growing season. For these reasons, we lumped together the herbaceous live, herbaceous dead, and bare ground classes. We additionally included artificial objects in this combined cover class to account for the possibility that dirt or gravel roads were sometimes classified as artificial and/or confused with bare ground.

2.7 Quantifying temporal vegetation change

Although our land cover predictions were produced as vector data derived from sub-meter resolution inputs, we quantified vegetation change between the early and late periods at the 30 m scale to minimize the influence of fine-scale prediction noise and/or error and focus on consistent patterns at a scale relevant to landscape ecology and management. Starting with the summarized vegetation cover layers created as described in the previous section, we superimposed a 30 m grid. For each pixel, if the pixel was at least 25% covered by predictions, and at least 60% of the area of the predictions was of a single cover type, that cover type was assigned to the pixel. For assessment of vegetation change, we retained only those pixels which were assigned a cover class in both the early and late periods. Within the burned areas of each reserve, we quantified the change (or stasis) as both the absolute and relative area of the retained pixels that contained different (or the same) cover classes between the two periods.

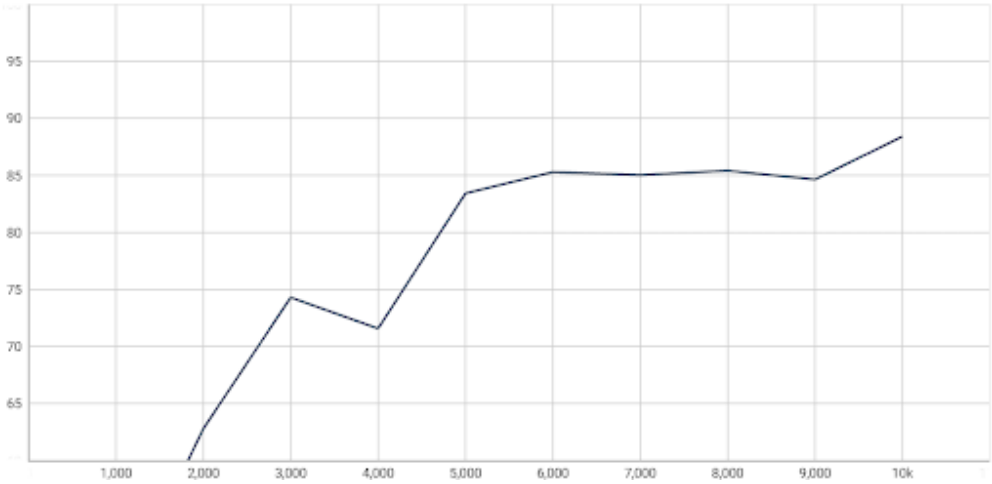

Figure 1: Validation mAcc across 10k training iterations.

3 Results

3.1 Computer vision model accuracy

The fine-tuned SegFormer-B5 model demonstrated strong performance in classifying vegetation cover types from the drone imagery. We evaluated the performance of the semantic segmentation model by analyzing the Mean Accuracy (mAcc), Mean Intersection over Union (mIoU) and pixel accuracy (aAcc) metrics. mAcc calculates the accuracy for each individual class and then averages the results, which highlights performance on rare classes that represent a small fraction of all pixels. mIoU evaluates the overlap between predicted and ground-truth segmentation masks divided by the union of the two masks, averaged across all classes. aAcc represents the overall pixel-wise accuracy across all classes. After 10k iterations, the model achieved an overall mAcc of 88.40%, mIoU of 81.69%, and aAcc of 93.96% on the validation set. Figure fig. 1 illustrates the progression of mAcc on the validation set as training iterations increase. The model demonstrated rapid gain in accuracy early in training, followed by a gradual plateau, converging around iteration 10000. The per-class IoU and accuracy on the validation set is summarized in table 2.

In fig. 2 we show the normalized confusion matrix for the model's predictions on the validation set. Most vegetation classes were classified with high accuracy and minimal misclassifications. For instance, Tree Live and Water had near-perfect results, with 96.46% and 100% of their respective pixels correctly predicted. In contrast, Shrub Dead showed notable misclassification, particularly with Shrub Live (39.33%) and Tree Dead (14.48%). This reflects the lower IoU for the class and is likely due to the visual similarities

Class	IoU (%)	Precision (%)	Recall (%)
Bare Earth	92.91	96.54	96.11
Herbaceous Live	85.34	91.68	92.51
Artificial Object	81.23	86.2	93.38
Shrub Dead	38.77	85.4	41.52
Shrub Live	83.09	86.49	95.48
Tree Dead	78.63	84.65	91.71
Tree Live	94.09	97.45	96.46
Water	99.43	99.43	100.0

Table 2: SegFormer-B5 class-wise performance on validation data.

Reserve	Type	Surveyed area (ha)	Predicted area (ha)	Predicted (%)
BORR	early	1277.6	1033.1	80.9
BORR	late	236.4	185.6	78.5
BORR	overlap	236.3	167.7	71.0
Hastings	early	275.9	233.4	84.6
Hastings	late	398.9	343.0	86.0
Hastings	overlap	189.1	138.1	73.0
Quail	late	171.0	132.5	77.5

Table 3: Summary of surveyed and predicted areas per reserve. The “early” category refers to 2020 and the “late” category refers to data combined between 2023 and 2024. The ”overlap” represents the area that overlaps between early and late time points.

between dead shrubs, live shrubs, and dead trees. The matrix also shows some confusion between Artificial Object and Bare Earth (5.08%), possibly due to overlapping spatial or color characteristics in certain areas. Overall, the matrix highlights that while most classes are clearly distinguishable, classes like Shrub Dead may improve from additional labeled data (see table 1).

3.2 Geospatial vegetation predictions

Using the approaches described above, predicted vegetation maps were produced using all the available leaf-on data. The surveyed area for each reserve is described in table 3. The Quail Ridge study site only had leaf-on data for 2023, so it cannot be used for change analysis. Both BORR and Hastings had data from both the early and late period, with an overlapping region of 236.3 ha for BORR and 189.1 ha for Hastings that can be used for change analysis. To conduct this analysis, we need high confidence predictions of vegetation types for both early and late time periods. The faction of confidently predicted area is lower

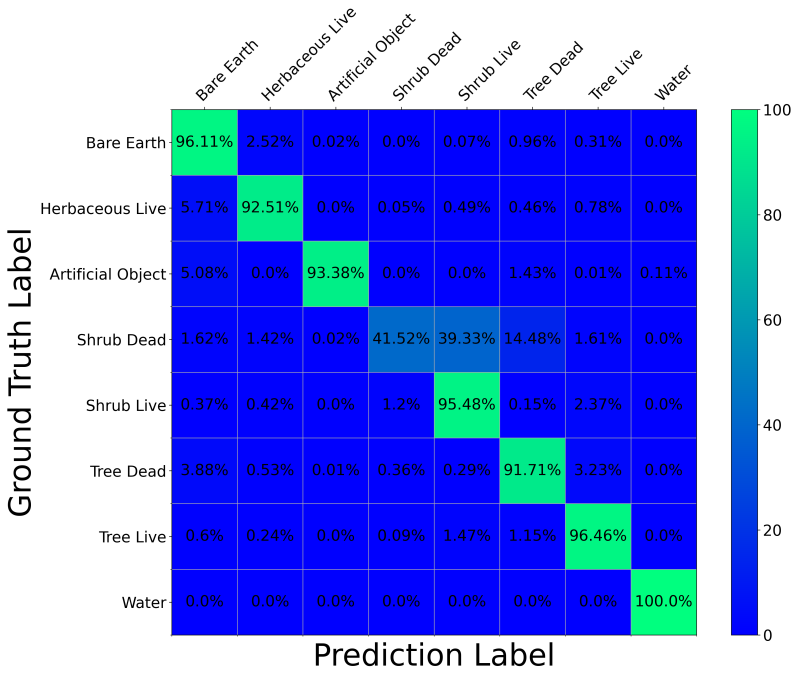

Figure 2: Normalized confusion matrix of SegFormer model's predictions on validation images.

in the overlapping categories than either the early or late categories because a region must be confidently predicted in both categories to be used for analysis. Despite this requirement, 71.0% of the total area surveyed at BORR and 73.0% of the area surveyed at Hastings met this criteria, resulting in 167.7 ha and 138.1 ha respectively of area available for change analysis at the two reserves.

The predicted regions are shown in fig. 3. Each represents a single high-resolution prediction synthesized from all leaf-on flights at that site in the specified time period. Note that the largest dataset, BORR early, represents 1033.1 ha of predicted area (table 3). There are limited abrupt changes in classification at the boundaries of input datasets, which shows that predictions were effectively merged across multiple input datasets. Regions of low confidence predictions, denoted in white, occur preliminarily at the boundary between different classes.

3.3 Temporal change in vegetation cover

Across both Hastings and Blue Oak Reserves, within the repeat-surveyed burned area with vegetation predictions at the 30 m scale, land cover largely remained constant (86% of landscape) (fig. 4). The primary transitions were of herbaceous/bare to live shrub (14% of the herbaceous/bare cover made this transition)

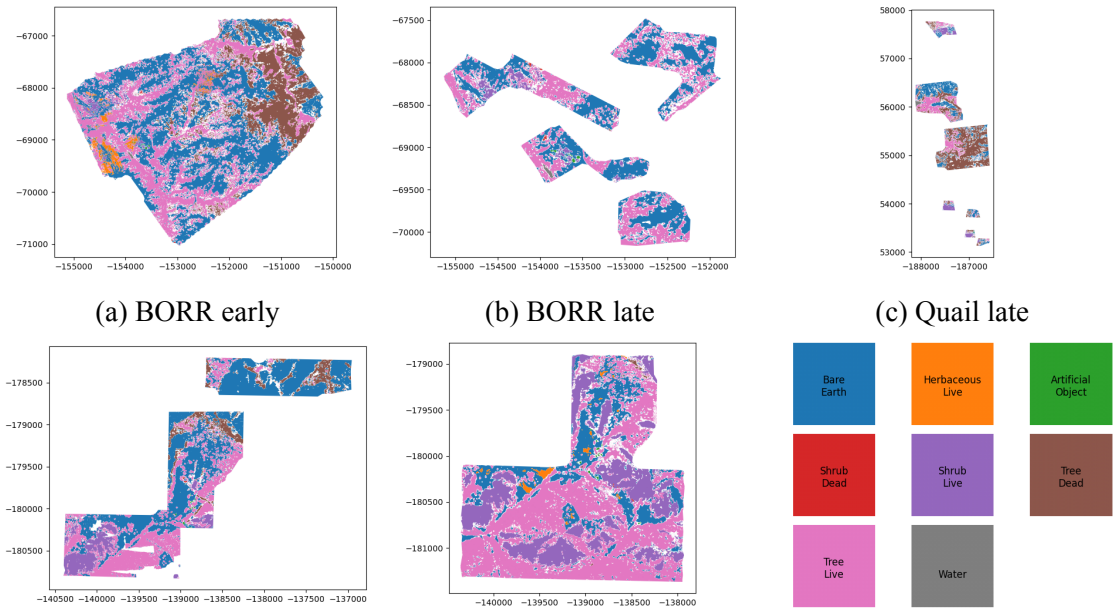

(a) BORR early (b) BORR late (c) Quail late

(d) Hastings early (e) Hastings late (f) Colormap

Figure 3: Predicted vegetation type maps for each site from the early/late periods containing leaf-on data. The axis are in the California Albers coordinate reference frame (EPSG:3310) with units of meters.

Figure 4: Sankey diagram of transitions in predicted land cover from early (left) to late (right) periods, where the vertical width of each band corresponds to the proportion of the landscape undergoing each transition. For purposes of this figure, "landscape" refers to burned 30 m pixels with confident predictions.

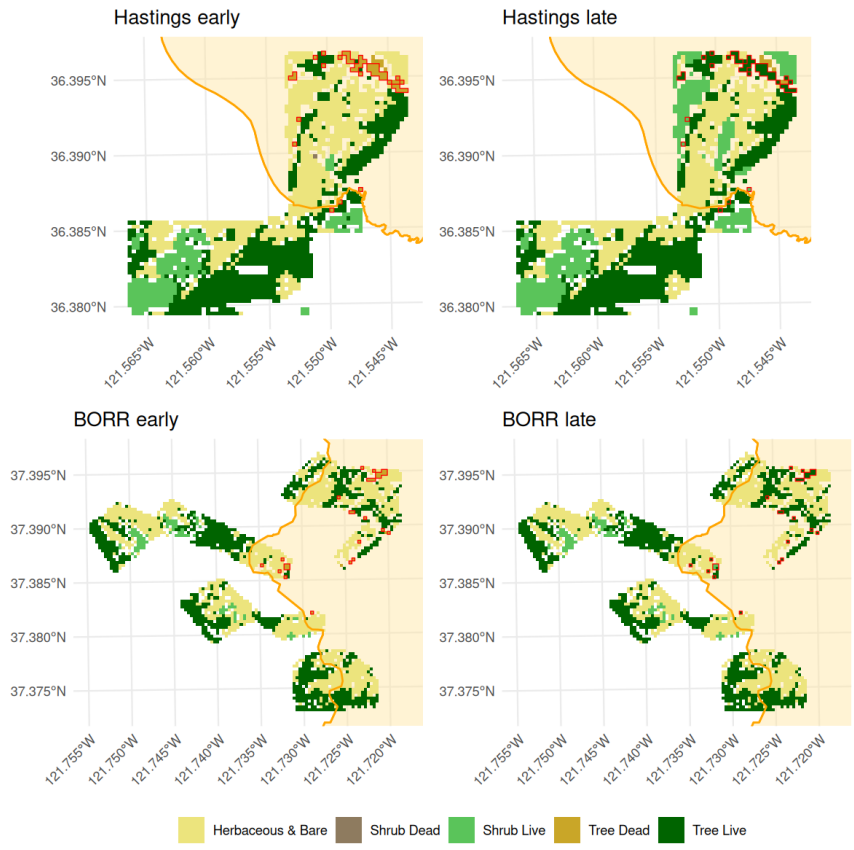

Figure 5: Predicted landcover class at the 30 m scale for all pixels with confident predictions in both the early and late time periods. Pixels outlined in red are those that transitioned from "dead tree" to "live tree" or "live shrub" (the latter presumably a misclassification of "live tree"). The fire footprint is depicted by orange shading.

and of dead tree to live tree or live shrub (87% of the dead tree cover made this transition) (fig. 4). Of the 84 dead tree pixels, only 10 remained dead in the late predictions (late predicted class "dead tree" (8 pixels) or "bare/herbaceous" (2 pixels)). All 8 pixels with late predictions "dead" were relatively closely clustered in the northeast section of the Hastings reserve study area, while the two "bare/herbaceous" pixels were relatively close to each other in the eastern section of the Blue Oaks Reserve study area (fig. 5).

4 Discussion

4.1 Accuracy considerations

Our analysis of per-image predictions on data not seen during model training shows that most classes are predicted with high accuracy. The overall accuracy of our model predictions within raw drone images (94%) exceeds accuracies achieved in similar systems by other drone-based mapping efforts, which tend to be applied over substantially smaller areas. A study over 41 ha in a California shrubland-woodland system Detka et al. (2023) achieved overall accuracy across 8 cover classes of roughly 83-87%, and a study over 10 ha in California chaparral achieved overall accuracy across 4 cover classes of 88% von Nonn et al. (2024). Our approach achieved accuracy at the raw image pixel level similar to that achieved by the NLCD at the level of 30 m geospatial pixels Wickham et al. (2023), with the additional advantages that it can distinguish between live and dead vegetation and is tailored to the species present in our study system. In our study, the class with lowest prediction accuracy was dead shrub, which had the fewest annotations of any vegetation class (only greater than water and artificial object)–and thus the least data available for model training–and is challenging to distinguish from dormant live shrubs in this ecosystem.

Notably, we do not have quantitative metrics for performance at a geospatial pixel level, which would be the most informative for assessing the utility of these predictions for ecological modeling. We obtained field reference data to support this analysis, but spatial inaccuracies in the locations of observations made the data unusable in this ecosystem where vegetation types are highly intermixed. Given the strong performance of per-image predictions, we expect that the geospatial predictions have comparable or higher accuracy, as the geospatial predictions are produced by voting across multiple viewpoints and therefore should be more robust than a single view alone. Indeed, our multi-view approach to vegetation classification has been shown in a different system to roughly halve error rates relative to a single-view orthomosaic approach Russell et al. (2024). Furthermore, qualitative analysis of our geospatial predictions against the underlying imagery suggests high accuracy, as does the high degree of agreement across years.

Figure 6: Early (2020) and late (2023) post-fire imagery of the same area of Hastings Reserve classified primarily as "dead tree" in the early imagery. Brown ellipses indicate areas predicted to contain dead trees at the time of imagery collection; some areas remained classified as "dead tree" and no greening was observed.

4.2 Post-fire vegetation change

Within the burned areas, the primary landcover change was from herbaceous/bare to live shrub. The rapid establishment of shrubs post-fire likely reflects rapid resprouting of shrubs that had most aboveground biomass consumed by fire. Shrubland ecosystems in California usually burn at high severity, consuming most foliage, branches, and often even larger stems Grupenhoff and Safford (2024); Keeley et al. (2008). At the time the early post-fire imagery was collected, most pre-fire aboveground shrub biomass had likely been removed by fire, exposing bare ground and/or herbaceous cover. Even severely burned shrubs often rapidly resprout in the years following fire Cowan and Ackerly (2010); Keeley et al. (2005), likely manifesting in our assessment as the large proportion of herbaceous/bare pixels transitioning to live shrub. Without pre-fire imagery, though, we cannot differentiate between resprouting of existing severely-burned shrubs vs. post-fire establishment of new individuals. Likewise, early post-fire bare/herbaceous pixels that remained in this class may have contained shrubs pre-fire that were burned so severely they did not resprout, but the existence and prevalence of this dynamic cannot be determined without pre-fire imagery.

Nearly all "dead tree" pixels in the early post-fire imagery transitioned to "live tree" pixels in the later imagery, implying that nearly all trees that appeared dead in the early post-fire imagery were in fact still alive (but simply missing live foliage) and, like shrubs, rapidly sprouted. In the case of trees, much of the resprouting appeared to be epicormic as opposed to basal: in many areas, the early post-fire "dead trees" were covered in dead, brown leaves but 3-4 years later appeared as full-size trees covered in green foliage

(fig. 6). Only a few small areas containing dead trees in the early post-fire imagery remained dead in the later imagery, apparently not even exhibiting basal resprouting; in many of these areas, manual imagery inspection revealed that the trees appeared to have had their foliage fully consumed by fire (fig. 6). This observation suggests that greater fire intensity (as represented by foliage consumption) is associated with a greater chance of tree mortality, a pattern that is corroborated by existing studies Ackerly et al. (2019). Alternatively, it is possible that the trees classified as dead at both time points were in fact dead pre-fire. However, the persistence of the finer branches through fire (apparent in the imagery) suggests that the trees were alive at the time of fire, as live/wet wood is less likely to be consumed by fire Prichard et al. (2017); Ottmar (2014); Goodwin et al. (2021). Finally, the small fraction of early dead tree pixels that transitioned to live shrub pixels likely reflects model confusion of basally resprouting trees with shrubs due to their similar appearance.

4.3 Future directions

Our approaches to unifying a large collection of drone imagery datasets from a diverse set of imagery collection platforms, dates, locations, and methods may help to serve as a model for leveraging the increasing availability of existing imagery, potentially opening doors to ecological insights not easily achievable otherwise. For example, given disturbances are often unpredictable, studies of disturbance effects may be enabled by leveraging preexisting data and combining it with purpose-collected post-disturbance data–and the confidence of resulting insights could be increased by performing similar comparisons across numerous spatial footprints. However, robust change analysis would require co-registering the disparate datasets and unifying them in an efficient processing pipeline such as the one we demonstrate.

This work highlights the ability to manually annotate high-resolution individual drone images as an alternative to field reference information when that is unavailable. However, field reference data is still essential in some circumstances, especially when distinguishing classes more granular than functional types. We believe that an effective strategy to mitigating spatial error in field measurements is to first produce an orthomosiac from drone imagery and then annotate it in field, as done in Detka et al. (2023). These annotations can then be appropriately shifted along with the reference drone imagery, using our proposed registration pipeline. While computer vision approaches are adept at leveraging texture cues, we found that strong appearance changes driven by seasonality are still a challenge, even after restricting to leaf-on data.

These changes were especially prominent for the shrub and herbaceous classes, which can change appearance dramatically when they dry out. If data cannot be collected only in one season, it may be helpful to train independent ML models per season or provide the season as additional context for the prediction task. A major challenge in this work was the lack of pre-fire drone imagery or land cover predictions matching our classification scheme. One approach to extend this work would be to predict pre-fire vegetation from existing remote sensing data such as NAIP, either by directly annotating NAIP data or training a model using predictions from drone data for years it was available. Finally, the high spatial resolution of these data (provided as vector data, with an effective resolution of single-digit centimeters) provides an opportunity to study granular ecological questions, such as the fate of individual trees or shrubs or the fine-grained drivers of resilience.

5 Conclusion

Our work demonstrates the potential to use multi-view drone imagery, photogrammetry, and computer vision to efficiently obtain reliable, locally tuned predictions of land cover and land cover change at fine spatial scales across large areas. The substantial consistency in predicted land cover across time periods – particularly in the unburned areas of each reserve – demonstrates robustness in model predictions, particularly considering that an entirely different set of imagery was used for each time period. The assessment of post-fire change and stasis revealed changes that align well with existing understanding of post-fire vegetation dynamics in woodland and chaparral systems, but it also reveals the limitations in assessing post-fire vegetation recovery without pre-fire vegetation data. The methods we apply in this study may help to form the basis of broader vegetation monitoring efforts.

6 Acknowledgments

Guidance on the study system was provided by Shane Waddell and Ross Brennan. Field work and imagery annotation were performed by Ellyn Paris, Justin Liu, Keith Kolenut, April Harlamoff, Anika Levy-Groth, and Emily Moore. Drone imagery was collected by Becca Fenwick, Justin Cummings, Jacob Flannagan, Clancy McConnell and Sean Hogan. This work was funded by CITRIS and the Banatao Institute at the

University of California (#2022-61), the University of California Climate Innovation and Entrepreneurship Fund (#R02CE6851), and the National Science Foundation Division of Biological Infrastructure (#DBI-2152671).

7 Conflict of Interest Statement

The authors declare no conflicts of interest.

References

Ackerly, D. D., M. Kozanitas, M. Oldfather, P. Papper, and M. Clark (2019). Mortality and resprouting in california oak woodlands following mixed-severity wildfire. *International oaks 30*.

Agisoft, LLC. (2025). Agisoft metashape user manual, professional edition, version 2.2. https://www.agisoft.com/pdf/metashape-pro_2_2_en.pdf.

Allred, B. W., B. T. Bestelmeyer, C. S. Boyd, C. Brown, K. W. Davies, M. C. Duniway, L. M. Ellsworth, T. A. Erickson, S. D. Fuhlendorf, T. V. Griffiths, V. Jansen, M. O. Jones, J. Karl, A. Knight, J. D. Maestas, J. J. Maynard, S. E. McCord, D. E. Naugle, H. D. Starns, D. Twidwell, and D. R. Uden (2021). Improving landsat predictions of rangeland fractional cover with multitask learning and uncertainty. *Methods in Ecology and Evolution 12*(5), 841–849.

Anderson, K., F. Gonzalez, and K. J. Gaston (2025, 06). Drones in ecology: ten years back and forth. *BioScience 75*(8), 664–680.

Bellia, A. F. and S. Lanfranco (2019). A preliminary assessment of the efficiency of using drones in land cover mapping.

Bowman, D. M., G. J. Williamson, J. T. Abatzoglou, C. A. Kolden, M. A. Cochrane, and A. M. Smith (2017). Human exposure and sensitivity to globally extreme wildfire events. *Nature ecology & evolution 1*(3), 0058.

Brodrick, P. G., A. B. Davies, and G. P. Asner (2019). Uncovering ecological patterns with convolutional neural networks. *Trends in ecology & evolution 34*(8), 734–745.

Callaway, R. M. and F. W. Davis (1993). Vegetation dynamics, fire, and the physical environment in coastal central california. *Ecology 74*(5), 1567–1578.

Chang, B., F. Li, Y. Hu, H. Yin, Z. Feng, and L. Zhao (2025). Application of uav remote sensing for vegetation identification: a review and meta-analysis. *Frontiers in Plant Science 16*, 1452053.

Cingolani, A. M., D. Renison, M. R. Zak, and M. R. Cabido (2004). Mapping vegetation in a heterogeneous mountain rangeland using landsat data: an alternative method to define and classify land-cover units. *Remote Sensing of Environment 92*(1), 84–97.

Contributors, M. (2020). MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. https://github.com/open-mmlab/mmsegmentation.

Cowan, P. D. and D. D. Ackerly (2010). Post-fire regeneration strategies and flammability traits of california chaparral shrubs. *International Journal of Wildland Fire 19*(7), 984–989.

Crotteau, J. S., J. M. Varner III, and M. W. Ritchie (2013). Post-fire regeneration across a fire severity gradient in the southern cascades. *Forest Ecology and Management 287*, 103–112.

Csurka, G., R. Volpi, and B. Chidlovskii (2023). Semantic image segmentation: Two decades of research.

Cunningham, C. X., G. J. Williamson, and D. M. Bowman (2024). Increasing frequency and intensity of the most extreme wildfires on earth. *Nature ecology & evolution 8*(8), 1420–1425.

Davila, D., J. VanPelt, A. Lynch, A. Romlein, P. Webley, and M. S. Brown (2022). Adapt: An open-source suas payload for real-time disaster prediction and response with ai.

Dawkins, M., L. Sherrill, K. Fieldhouse, A. Hoogs, B. Richards, D. Zhang, L. Prasad, K. Williams, N. Lauffenburger, and G. Wang (2017). An open-source platform for underwater image and video analytics. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 898–906.

Detka, J., H. Coyle, M. Gomez, and G. S. Gilbert (2023). A drone-powered deep learning methodology for high precision remote sensing in california's coastal shrubs. *Drones 7*(7).

Fairman, T. A., L. T. Bennett, and C. R. Nitschke (2019). Short-interval wildfires increase likelihood of resprouting failure in fire-tolerant trees. *Journal of Environmental Management 231*, 59–65.

Goodwin, M. J., H. S. J. Zald, M. P. North, and M. D. Hurteau (2021). Climate-driven tree mortality and fuel aridity increase wildfire's potential heat flux. *Geophysical Research Letters 48*(24), e2021GL094954. e2021GL094954 2021GL094954.

Grupenhoff, A. R. and H. D. Safford (2024). High fire frequency in california chaparral reduces postfire shrub regeneration and native plant diversity. *Ecosphere 15*(12), e70128.

Holmes, K. A., K. E. Veblen, T. P. Young, and A. M. Berry (2008). California oaks and fire: A review and case study. In *Proceedings of the sixth California oak symposium: today's challenges, tomorrow's opportunities. USDA Forest Service, Pacific Southwest Research Station, Redding, CA (PSW-GTR-217)*, pp. 551–565.

Ibanez, L., B. Lorensen, M. McCormick, B. King, H. Johnson, D. Blezek, B. Lowekamp, J. Jomier, J. Miller, G. Lehmann, J. Cates, N. Dekker, J. H. Legarreta, D. Zukić, L. Ng, J. Kim, A. Gelas, M. Malaterre, K. Krishnan, B. Hoffman, K. Williams, S. R. Aylward, F. Budin, J. H. Legarreta, W. Schroeder, X. Liu, B. Avants, M. McCormick, S. McBride, A. Noe, M. Popoff, G. Hart, D. Zukić, T. Sundaram, H. Johnson, A. Gouaillard, M. Stauffer, N. Tustison, A. Enquobahrie, S. Pathak, A. Cedilnik, T. Chen, D. Shelton, B. Helba, C. Quammen, P. Hernandez-Cerdan, Y. Jin, D. Padfield, T. Vercauteren, H. Jae Kang, M. Turek, R. Tamburo, M. Audette, M. Foskey, N. Tustison, P. Hughett, D. Doria, J.-C. Fillion-Robin, G. Kindlmann, L. Newberg, D. Cole, V. S. FONOV, W. Turner, S. Chen, T. Tasdizen, J. Duda, J. Galeotti, S. BARRE, S. Jaume, K. Mosaliganti, M. Isakov, P. Chandra, A. Ghayoor, A. Mackelfresh, C. Mullins, Y. Zhuge, S. Rit, D. Vigneault, K. Martin, X. Xue, M. Straing, R. Estepar, A. Squillacote, B. Wyman, W. Chang, J.-P. Guyon, C. Botha, L. Baghdadi, maekclena, S. P. Awate, P. Reynolds, Z. Pincus, J. Finet, R. Venkatram, Z. Yaniv, K. Zygmunt, P. Rondot, B. Davis, L. Antiga, M. Coursolle, M. Roden, S. Park, H. Cheung, H. Greer, L. Gandel, C. Aaron Cois, R. Kaucic, M. D. Hanwell, V. Magnotta, F. Bertel, J. Hipwell, R. Chandrashekara, R. Beare, S. Gerber, D. C. Bigler, M. Straing, M. Styner, S. Robbins, V. Chalana, Y. Le Poul, A. Neundorf, P. Lamb, Z. Williamson, T. Braun-Jones, A. Wasem, N. Rannou, and I. C. Members (2025, June). Insightsoftwareconsortium/itk: Itk 5.4.4.

Jolly, W. M., M. A. Cochrane, P. H. Freeborn, Z. A. Holden, T. J. Brown, G. J. Williamson, and D. M. Bowman (2015). Climate-induced variations in global wildfire danger from 1979 to 2013. *Nature communications 6*(1), 7537.

Keeley, J. E., T. Brennan, and A. H. Pfaff (2008). Fire severity and ecosytem responses following crown fires in california shrublands. *Ecological Applications 18*(6), 1530–1546.

Keeley, J. E., C. J. Fotheringham, and M. Baer-Keeley (2005). Determinants of postfire recovery and succession in mediterranean-climate shrublands of california. *Ecological Applications 15*(5), 1515–1534.

Liu, T. and A. Abd-Elrahman (2018). Multi-view object-based classification of wetland land covers using unmanned aircraft system images. *Remote Sensing of Environment 216*, 122–138.

Lu, T., L. Wan, S. Qi, and M. Gao (2023). Land cover classification of uav remote sensing based on transformer–cnn hybrid architecture. *Sensors 23*(11).

Mlambo, R., I. H. Woodhouse, F. Gerard, and K. Anderson (2017). Structure from motion (sfm) photogrammetry with drone data: A low cost method for monitoring greenhouse gas emissions from forests in developing countries. *Forests 8*(3).

Moritz, M. A., M.-A. Parisien, E. Batllori, M. A. Krawchuk, J. Van Dorn, D. J. Ganz, and K. Hayhoe (2012). Climate change and disruptions to global fire activity. *Ecosphere 3*(6), 1–22.

OpenAerialMap (2025). Openaerialmap: The open collection of aerial imagery. https://openaerialmap.org/.

Ottmar, R. D. (2014). Wildland fire emissions, carbon, and climate: Modeling fuel consumption. *Forest Ecology and Management 317*, 41–50.

Prichard, S., M. Kennedy, C. Wright, J. Cronan, and R. Ottmar (2017). Predicting forest floor and woody fuel consumption from prescribed burns in southern and western pine ecosystems of the united states. *Forest Ecology and Management 405*, 328–338.

Russell, D., A. Sidhu, R. Prabhu, D. Young, and T. Addisherla (2025, May). open-forest-observatory/geograypher: v0.3.0.

Russell, D., B. Weinstein, D. Wettergreen, and D. Young (2024). Classifying geospatial objects from multi-view aerial imagery using semantic meshes. *arXiv preprint arXiv:2405.09544*.

Sohl, M. A. and S. A. Mahmood (2024). Low-cost uav in photogrammetric engineering and remote sensing: Georeferencing, dem accuracy, and geospatial analysis. *Journal of Geovisualization and Spatial Analysis 8*(1), 14.

Thorne, J. H., M. Gogol-Prokurat, S. Hill, D. Walsh, R. M. Boynton, and H. Choe (2020). Vegetation refugia can inform climate-adaptive land management under global warming. *Frontiers in Ecology and the Environment 18*(5), 281–287.

U.S. Department of Agriculture (2011). National Agriculture Imagery Program (NAIP) Information Sheet. (March).

von Nonn, J., M. L. Villarreal, L. Blesius, J. Davis, and S. Corbett (2024). An open-source workflow for scaling burn severity metrics from drone to satellite to support post-fire watershed management. *Environmental Modelling Software 172*, 105903.

Whitman, E., M.-A. Parisien, D. K. Thompson, and M. D. Flannigan (2019). Short-interval wildfire and drought overwhelm boreal forest resilience. *Scientific Reports 9*(1), 18796.

Wickham, J., S. V. Stehman, D. G. Sorenson, L. Gass, and J. A. Dewitz (2023). Thematic accuracy assessment of the nlcd 2019 land cover for the conterminous united states. *GIScience & Remote Sensing 60*(1), 2181143.

Wulder, M. A., D. P. Roy, V. C. Radeloff, T. R. Loveland, M. C. Anderson, D. M. Johnson, S. Healey, Z. Zhu, T. A. Scambos, N. Pahlevan, M. Hansen, N. Gorelick, C. J. Crawford, J. G. Masek, T. Hermosilla, J. C. White, A. S. Belward, C. Schaaf, C. E. Woodcock, J. L. Huntington, L. Lymburner, P. Hostert, F. Gao, A. Lyapustin, J.-F. Pekel, P. Strobl, and B. D. Cook (2022). Fifty years of landsat science and impacts. *Remote Sensing of Environment 280*, 113195.

Xie, E., W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo (2021). Segformer: simple and efficient design for semantic segmentation with transformers. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA. Curran Associates Inc.

Young, D. (2025). Open forest observatory: Drone-based mapping for forest ecology and management. https://openforestobservatory.org/.

Young, D. J. N., M. J. Koontz, and J. Weeks (2022). Optimizing aerial imagery collection and processing parameters for drone-based individual tree mapping in structurally complex conifer forests. *Methods in Ecology and Evolution 13*(7), 1447–1463.