

PalmProbNet: A Probabilistic Approach to Understanding Palm Distributions in Ecuadorian Tropical Forest via Transfer Learning

Kangning Cui City University of Hong Kong Kowloon, Hong Kong kangnicui2-c@my.cityu.edu.hk

Victor Pauca Wake Forest University Winston-Salem, NC, USA paucavp@wfu.edu

João Pinheiro Wake Forest University Winston-Salem, NC, USA deolj19@wfu.edu Zishan Shao Wake Forest University Winston-Salem, NC, USA shaoz20@wfu.edu

Sarra Alqahtani Wake Forest University Winston-Salem, NC, USA sarra-alqahtani@wfu.edu

Manqi Wang Wake Forest University Winston-Salem, NC, USA wangm21@wfu.edu Gregory Larsen Wake Forest University Winston-Salem, NC, USA larseng@wfu.edu

David Segurado Wake Forest University Winston-Salem, NC, USA dsegurado@wfu.edu

David Lutz
Dartmouth College
Hanover, NH, USA
david.a.lutz@dartmouth.edu

Robert Plemmons Wake Forest University Winston-Salem, NC, USA plemmons@wfu.edu

ABSTRACT

Palms play an outsized role in tropical forests and are important resources for humans and wildlife. A central question in tropical ecosystems is understanding palm distribution and abundance. However, accurately identifying and localizing palms in geospatial imagery presents significant challenges due to dense vegetation, overlapping canopies, and variable lighting conditions in mixed-forest landscapes. Addressing this, we introduce PalmProb-Net, a probabilistic approach utilizing transfer learning to analyze high-resolution UAV-derived orthomosaic imagery, enabling the detection of palm trees within the dense canopy of the Ecuadorian Rainforest. This approach represents a substantial advancement in automated palm detection, effectively pinpointing palm presence and locality in mixed tropical rainforests. Our process begins by generating an orthomosaic image from UAV images, from which we extract and label palm and non-palm image patches in two distinct sizes. These patches are then used to train models with an identical architecture, consisting of an unaltered pre-trained ResNet-18 and a Multilayer Perceptron (MLP) with specifically trained parameters. Subsequently, PalmProbNet employs a sliding window technique on the landscape orthomosaic, using both small and large window sizes to generate a probability heatmap. This heatmap effectively

Miles Silman Wake Forest University Winston-Salem, NC, USA silmanmr@wfu.edu

visualizes the distribution of palms, showcasing the scalability and adaptability of our approach in various forest densities. Despite the challenging terrain, our method demonstrated remarkable performance, achieving an accuracy of 97.32% and a Cohen's κ of 94.59% in testing.

CCS CONCEPTS

• Computing methodologies \rightarrow Computer vision problems; Neural networks; • Applied computing \rightarrow Environmental sciences.

KEYWORDS

transfer learning, palm detection, density map, remote sensing

ACM Reference Format:

Kangning Cui, Zishan Shao, Gregory Larsen, Victor Pauca, Sarra Alqahtani, David Segurado, João Pinheiro, Manqi Wang, David Lutz, Robert Plemmons, and Miles Silman. 2024. PalmProbNet: A Probabilistic Approach to Understanding Palm Distributions in Ecuadorian Tropical Forest via Transfer Learning. In 2024 ACM Southeast Conference (ACMSE 2024), April 18–20, 2024, Marietta, GA, USA. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3603287.3651220

1 INTRODUCTION

Detecting palm trees in tropical forests is important in both biological and computational domains. Biologically, palms serve as important ecological indicators, providing insights into biodiversity, soil quality, and overall health of the forest ecosystem, support human livelihoods in indigenous and rural communities, and are keystone resources for tropical wildlife [4, 15]. Computationally, particularly within image processing and machine learning, palm detection is challenging due to issues such as noise and artifacts

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACMSE 2024, April 18–20, 2024, Marietta, GA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0237-2/24/04...\$15.00 https://doi.org/10.1145/3603287.3651220

in image orthomosaics, variable illumination within forests, lack of labeled data, and data imbalances [4, 13, 15]. Addressing these challenges is essential not only for ecological studies that rely on identification and localization of forest resources but also for advancing computer vision techniques for remote sensing in general.

In response to the outlined challenges, this study introduces PalmProbNet, a deep learning approach tailored for palm tree detection in Ecuador's tropical forests, capitalizing on UAV-derived orthomosaic imagery. We employ a strategy of training two models, both with a consistent architecture, on image patches of varying sizes, resulting in probability maps that depict palm distributions. Experimental results highlight the robustness and efficacy of Palm-ProbNet in classifying palm trees. Our contributions are threefold: a data labeling process that categorizes two sets of image patches containing palm and non-palm characteristics in different sizes, feature extraction via transfer learning integrated with deep learning-based classification, and the application to the full UAV-derived orthomosaic imagery. The results emphasize the promise of integrating UAV technology with deep learning for efficient palm tree detection within dense forest canopies.

The subsequent sections of the paper are structured as follows: Section 2 provides a comprehensive review of related work in palm detection and deep learning. Section 3 details our dataset, including data collection, preprocessing, and labeling. Section 4 introduces PalmProbNet, our approach covering feature extraction, classification, and application to the landscape orthomosaic image. Section 5 showcases the experimental setup and its findings, succeeded by a comprehensive discussion on these outcomes. Finally, Section 6 concludes the paper and outlines directions for future research.

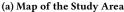
2 RELATED WORK

2.1 Object Detection

Object detection, a pivotal task in computer vision, identifies and classifies objects within images. Broadly, the methodologies can be categorized into traditional and deep learning-based approaches [5, 8, 23, 25, 26]. Traditional approaches follow a structured pipeline: selecting informative regions, extracting hand-crafted features, and classifying these features. Specifically, regions of interest (RoIs) within the image, which contains the targeted objects, are identified using methods like sliding window or selective search [14, 23]. The selected RoIs are then subjected to feature extraction approaches to describe specific features within the region, capturing object appearance, shape, and texture information [19, 23, 26]. The features are then fed into classifiers to categorize the RoIs and finalize the object detection [8, 23]. Traditional classifiers tend to be computationally efficient due to their fewer tunable parameters, in contrast to their deep learning counterparts [23, 26].

Deep learning-based techniques, on the other hand, offer an end-to-end training and prediction process by integrating feature extraction and classification into a unified framework [3, 15, 20, 25]. Notable deep learning algorithms include You Only Look Once (YOLO) and Region-Based Convolutional Neural Networks (R-CNNs) [26]. YOLO adopts a single-shot architecture that enables it to sufficiently perform object detection in real-time, while R-CNNs leverage a two-step process that first identifies regions of interest and then







(b) The Original Orthomosaic

(c) The Cleaned Orthomosaic

Figure 1: Visualizations of the Orthomosaic Image Before and After Cleaning





(a) Isolated Palm

(b) Hidden Palm

(c) Overlapped Palms

Figure 2: Palms as Captured in the Orthomosaic

classifies those regions, providing a more accurate but computationally intensive approach [17, 20, 25, 26]. Despite their capabilities, both traditional and deep learning methods encounter challenges in real-world scenarios, especially in detecting palm trees. Factors like dense canopies and diverse species in forests introduce complexities such as overlapping objects and intricate backgrounds [9, 21].

2.2 Palm Detection

Modern remote sensing technologies have significantly advanced palm tree detection by increasing the abundance, availability and resolutions of both satellite and aerial imagery. Remote sensing often leverages specialized sensors to capture distinct spectral and structural characteristics [15, 21]. For example, thermal sensors monitor surface temperature and can aid in irrigation design for palm plantations, multispectral sensors produce NDVI images which can help distinguish the presence, types, and organismal qualities of vegetation, hyperspectral sensors yield detailed profiles of spectral reflectance across large segments of the electromagnetic

spectrum, and LiDAR sensors provide topographic and structural data [13, 14]. These sensors are deployed on spaceborne platforms (satellites), airborne platforms (drones or occupied aircrafts) or ground-based platforms, with trade-offs of cost, scalability, and observation scale depending on the platform and its deployment [14, 15]. In this context, unoccupied aerial vehicles (UAVs) provide a distinctly scalable, cost-effective opportunity for canopy monitoring with ultra-high image resolutions, achievable due to their typically high-resolution sensors and ability to operate safely at very low altitudes [13, 17]. UAVs enable efficient monitoring of vast extents of terrain, catering especially to needs of small-scale farmers monitoring diverse palm species [14, 23]. The integration of automated palm detection techniques, whether conventional or deep learning-based, further accentuates the potential of UAV applications for mapping and monitoring vegetation land use and land cover.

UAVs and machine learning approaches rapidly enhance vegetation analysis for agricultural and forestry management, with methods like sliding window, histogram of oriented gradients, and SVM effectively detecting palm trees in UAV imagery from Malaysia and Saudi Arabia [13, 19, 19, 23, 23]. However, while effective, machine learning classifiers depend on feature quality. Introducing a new dataset necessitates fresh feature extraction, which complicates the training phase [10]. In contrast, deep learning approaches, particularly CNNs, inherently and adaptively extract relevant features during training, remarkably simplifying the process. CNNs have been applied to canopy monitoring with UAVs, effectively addressing complex object detection challenges [9, 12, 20, 21]. R-CNNs can effectively identify Mauritia flexuosa palms in Amazonian forests and oil palm trees in Malaysian plantations, demonstrating CNNs' adaptability in both forest and agricultural settings [17, 20]. Moreover, other architectures like YOLO have been successfully utilized to detect loose fruits of oil palms from UAV images, contributing a major advance to precision agriculture [12].

Identifying palm trees from high-resolution aerial or satellite imagery presents several challenges. Conventional classifiers may not capture hierarchical feature representations adequately, whereas deep-learning methods demand extensive training datasets [14, 26]. Integrating CNN-based methods with conventional classifiers can harness the strengths of both, offering a more robust solution for palm tree detection [8, 26]. This has been demonstrated, using CNNs to detect oil palm trees in satellite images using sliding window and post-processing techniques [16].

3 DATASET

3.1 Raw and Orthomosaic Data

The raw data used in this study pertains to a natural reserve in the Northwest region of Ecuador, known as the Ecuadorian Choco forest $(00^{\circ}23'28'' \text{ N}, 79^{\circ}41'05'' \text{ W})$, see Figure 1. The reserve is a high diversity humid tropical forest at 500m elevation, receiving \sim 3000 mm precipitation per year. Precipitation is seasonal, with a 5-month minimum from July-November accompanied by persistent fog (see [18] for detailed description). The forests contain 16 species of palms that can have exposed crown [1, 18].

The data was collected in collaboration with the Fundación para la Conservación de los Andes Tropicales (FCAT). The UAV imagery



Figure 3: Comparison of Small Patches

was collected in two campaigns. The first in June 2022 covered 95 ha using a DJI Phantom 4 RTK drone with a 1" CMOS sensor and controlled using UgCS mission planning software (CITS). Missions were flown at 90 m/s with 70% sidelap and 80% endlap. A total of 387 photos were taken and an orthomosaic and 3D point cloud were generated using Structure from Motion (SfM) in Agisoft Metashape 2.0 (Agisoft 2022). Flight plans, raw imagery, and orthomosaics are available upon request.

An orthomosaic is a large photogrammetrically orthorectified image product created from an image collection of a particular geographical region. Orthomosaics generated from SfM can contain considerable distortion due to the orthorectification process, especially near the edges of the reconstruction. Typically, distorted regions are removed to improve the visual consistency of the imagery, as shown in Figure 1.

3.2 Dual-Scale Manual Labels

A ground survey provided the exact location of 2,929 palms in the study region. Upon observation, we found that many of these palms are not visible in the orthomosaic image, due to severe occlusion or overlap with surrounding trees, see e.g. Figure 2. The resulting number of sufficiently visible palms in the orthomosaic was deemed too small for supervised classification. Thus, the following strategy was adopted to generate labeled data.

Fine-scale labeled data. We used the location of visible palms in the orthomosaic to extract 6000 small image patches of 40 by 40 pixels, capturing as much palm leaf feature diversity as possible. We endeavored to ensure that the extracted patches contain at least 90% of palm data within them. Similarly, we extracted 6000 40-by-40-pixel patches of non-palm features, such as other types of tree crowns, ground soil, roads, tree trunks, and other land cover types.

Sample images for the resulting palm and non-palm classes are shown in Figure 3. The images on the left are patches containing palm features, capturing the distinctive characteristics of a palm tree. The images on the right, on the other hand, showcase patches without palm features, highlighting variations that may help the model to differentiate between palms and non-palms.

Coarse-scale labeled data. The second part of our labeling process involved larger 100x100 pixel patches labeled as having high/low probability of containing palms. This scheme not only captures intricate leaf details but also broader tree crown contexts, crucial for distinguishing partially occluded palms from other trees.

A semi-manual approach was employed for this purpose. We used a portion of our fine-scale labeled data to train a network and identify small palm patches throughout the entire orthomosaic image. The locations of these patches reveals the presence of a palm tree, even in areas where the palm is mostly occluded. We then



Figure 4: Comparison of Large Patchs

explore 100x100 pixel areas around these patches, labeling them as either palms or non-palms. This approach resulted in about 3367 palm and 4091 non-palm samples, as shown in Figure 4, with the left four images showing palm patches and the right four showing non-palm patches, including roads and other tree crowns. Notably, some palms emerged within the low probability patches and some non-palms within the high probability ones, underscoring the effectiveness of the large patch datasets in tree crown identification.

4 METHODOLOGY

In this section, we introduce the proposed PalmProbNet that focuses on generating a heatmap to show the probability distribution of palm presence, aiding in biological applications. PalmProbNet encompasses three pivotal stages: feature extraction, classification, and application to the landscape orthomosaic image. The subsequent sections provide a detailed introduction of these stages and Figure 5 shows the workflow of the proposed PalmProbNet.

4.1 Feature Extraction

Transfer learning is a potent technique that leverages pre-trained models, reducing the need for extensive data collection and computational resources [24]. Customization of these models is often achieved by adding specialized layers tailored to the specific task. Multi-Layer Perceptrons (MLPs), feedforward neural networks with multiple layers, are commonly used for this purpose due to their capability to learn intricate patterns [22]. Residual Networks (ResNet) address challenges in training deep neural networks by introducing shortcut or "skip" connections, making them suitable for complex tasks like palm tree detection [11]. In our study, given the limited labeled data, we utilized the pre-trained ResNet-18 model, a variant with 18 layers. Originally trained on the ImageNet dataset [7], ResNet-18 offers a robust feature hierarchy. We fine-tuned its final layers to our dataset, ensuring a robust foundation for the PalmProb-Net. Additionally, we incorporated data augmentation techniques to enhance the diversity of our training data, addressing issues such as variable illumination and occlusions [17].

4.2 Classification

To classify the palm images based on the extracted features from ResNet layers, PalmProbNet employs an MLP that can capture complex patterns in data by introducing non-linearity [22]. Our MLP architecture, shown in Figure 5, incorporates a batch normalization layer, enhancing the model's ability to learn from ResNet's feature vectors and subsequently map them to the appropriate class labels. The overall workflow is illustrated in Figure 5.

4.3 Application to the Orthomosaic Image

Applying the model trained by small/large patches to large landscape orthomosaic images in PalmProbNet involves several steps. Depending on the training set, the model uses either 40×40 or 100×100 pixel patches. The process begins by loading the image,

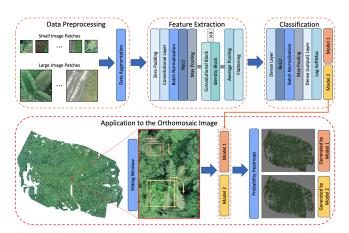


Figure 5: Workflow of The Proposed PalmProbNet

preprocessing, and capturing patches using a sliding window with a stride of 10. Patches with 25% or more missing pixels are immediately categorized as 'non-palm'. The ResNet layers then extracts features from each patch, which are predicted by the trained MLP layers. Probabilities from these predictions are aggregated into an accumulator array, with overlapping regions being averaged to ensure accurate probabilities are resized to generate a heatmap, which is then overlaid onto the orthomosaic image to emphasize the areas containing palms. The final results, including the highlighted image indicating the palm regions and the heatmap of the landscape orthomosaic, are saved for further analysis, capturing all computed probabilities.

5 EXPERIMENT

5.1 Data Division and Augmentation

To ensure a fair evaluation of PalmProbNet, we divided our dataset (for both small and large patches datasets) into two parts: 80% for training and 20% for testing. The training set was further subjected to 5-fold cross-validation, providing a more reliable assessment of the model's performance.

Data augmentation enhances model robustness and performance, especially in image-based tasks with limited labeled datasets. For the training set, we implemented several preprocessing and augmentation techniques. The images were resized to 224×224 pixels for compatibility with the pre-trained ResNet layers. Augmentations included random horizontal and vertical flips, which mirror the image in different orientations, and color jittering to simulate various lighting conditions. Lastly, normalization was applied to scale pixel values to a standardized range, facilitating faster convergence during training. For the test set, we avoided augmentations that might alter the intrinsic characteristics of the images, preserving consistency with the model's expectations.

5.2 Model Training and Evaluation

In the training of PalmProbNet, we fine-tuned the MLP layer parameters for task-specific adaptations, while keeping most ResNet parameters frozen to guarantee robust feature extraction. The final layer of ResNet was specifically fine-tuned to better align the network's output with the unique characteristics of our dataset.

Table 1: Classification Results for Different Configurations

Patch Size	nnodes	Acc	ROC AUC	AA	κ	Precision	Recall
Small	64	0.9101	0.9642	0.9105	0.8203	0.9126	0.9101
	128	0.9041	0.9653	0.9047	0.8084	0.9099	0.9041
	256	0.9110	0.9674	0.9114	0.8221	0.9132	0.9110
Large	64	0.9698	0.9963	0.9694	0.9391	0.9698	0.9698
	128	0.9712	0.9949	0.9705	0.9418	0.9712	0.9712
	256	0.9732	0.9959	0.9735	0.9459	0.9733	0.9733

During each fold of the cross-validation, the training data was augmented using the described techniques, while the validation data remained unchanged. We submitted SLURM jobs on Wake Forest University's cluster to accelerate the computation. The training employed the Negative Log-Likelihood loss function to evaluate prediction accuracy.

The core training phase consisted of 500 epochs for small patches and 200 epoches for large patches, employing a batch size of 64. We monitored the validation loss to select the best model, characterized by the lowest loss across epochs and folds. This model was then retrained on the entire training dataset with the same augmentation techniques and evaluated on a non-augmented test set. This systematic training approach, integrating cross-validation and validation loss monitoring, provided a robust evaluation framework, ensuring optimal model performance and reliability.

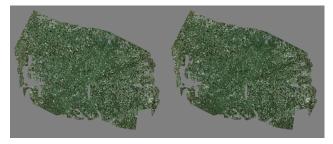
We utilized several metrics to evaluate the performance of Palm-ProbNet. Accuracy, given by $Acc = \frac{Number \ of \ Correct \ Predictions}{Total\ Predictions}$, measures the proportion of correct classifications. Average Accuracy (AA) is the mean of individual class accuracies, offering a balanced view, especially for imbalanced datasets. Cohen's κ coefficient gauges prediction reliability by assessing agreement between predicted and actual classes beyond chance. For binary classification, Precision, defined as Precision = $\frac{True\ Positives}{True\ Positives\ +\ False\ Positives},$ evaluates the model's exactness, while Recall, given by Recall = $\frac{True\ Positives\ +\ False\ Negatives\ +\ Palse\ Negatives\ +\ Negatives$

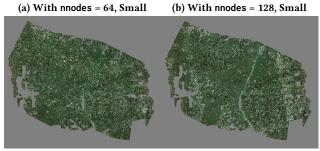
5.3 Numerical Results and Discussions

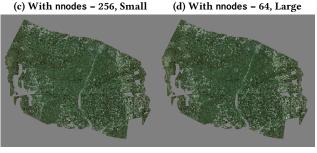
In evaluating the numerical results, the trained models tailored for both large and small patch sizes were tested on the testing patches and deployed on the landscape orthomosaic image. A stride of 10 was employed during this process, which allowed for a comprehensive and efficient scan of the entire image.

Our comparative analysis of patch classification, encompassing both small and large patches, involved three configurations of hidden layer nodes (nnodes): 64, 128, and 256. The performances, shown in Table 1, reveal a uniform level of effectiveness across the configurations for both patch sizes. For small patches, the 256-node configuration showed a slight edge, although the overall stable metrics across all configurations indicate that additional fine-tuning may be necessary for significant performance improvements. In the realm of large patches, the 256-node setup similarly displayed a marginal but noticeable enhancement in performance measures, such as accuracy, average accuracy, and Cohen's κ coefficient. These parallel findings across different patch sizes underscore the robustness of the classification approach and suggest that the 256-node configuration provides a subtle yet consistent advantage.

Figure 6(a-c) presents the probability heatmaps generated from the landscape orthomosaic image using the PalmProbNet trained on







(e) With nnodes = 128, Large

(f) With nnodes = 256, Large

Figure 6: Probability Heatmaps Produced by PalmProbNet Trained on Small & Large Patches

small patches. These heatmaps visually encode the model's assessed probability of palm presence, with brighter areas corresponding to higher probabilities of palm features. The intensity within each 10×10 segment of the heatmap is indicative of the model's confidence in identifying palm elements within that specific area.

Considering three different hidden layer configurations (nnodes = 64, 128, 256), the results, while broadly similar, exhibit subtle differences upon closer inspection. Each configuration efficiently detects a majority of the palm trees, though there are occasions where non-palm tree crowns are erroneously classified as palms. Such instances emphasize the necessity for more refined model adjustments to reduce false positives. A visual evaluation suggests that configurations with nnodes = 64 and nnodes = 256 marginally surpass the performance of nnodes = 128. The consistent detection patterns across diverse node configurations underline the reliability of the feature extraction and classification phases, simultaneously indicating areas where further tuning could elevate model accuracy.

Figure 6(d-f) displays the probability heatmaps derived from the landscape orthomosaic image, utilizing PalmProbNet trained on large patches. Each map is the result of a different hidden layer node configuration (nnodes = 64, 128, 256), with the intensity of the highlights indicating the likelihood of palm presence. The three

configurations yield very similar results, making it challenging to determine a clear superior performer. However, a common issue across all configurations is the misclassification of certain areas. Notably, patches adjacent to blank areas—those that have been either deleted due to poor image quality or are missing—are prone to be inaccurately labeled as containing palms. Addressing these misclassifications may require incorporating training samples that include these edge cases or applying post-processing steps to exclude them, thereby refining the model's ability to discern true palm features from artifacts introduced by image processing.

6 CONCLUSION AND FUTURE WORK

Identification of ecologically and economically important species within tropical forests canopies expands the frontiers of research and conservation. The paper details a thorough investigation into identifying palm trees in Neotropical rain forest using PalmProb-Net on UAV-captured images. In particular, it tackles challenges commonly encountered in detection tropical palm canopies like noise and uneven lighting, providing precise palm localization and landscape spatial distributions. PalmProbNet employed two models with the same architecture but trained on different-sized image patches. The models, particularly with a 256-node configuration, yielding robust results: for small patches, an accuracy of 0.9110, ROC AUC of 0.9674, and a Cohen's κ coefficient of 0.8221; for large patches, an accuracy of 0.9732, ROC AUC of 0.9959, and a Cohen's κ coefficient of 0.9459. Probability heatmaps were introduced as a new way to illustrate palm distributions, useful for both environmental and economic research and computational tasks. The integration of UAVs and deep learning has been validated for palm detection in dense forests through high classification accuracy across various model configurations.

Future enhancements to PalmProbNet will focus on refining the heatmap by including training samples with edge cases and employing post-processing to filter out patches near blank areas. We plan to include a fairer comparison of PalmProbNet with other classifiers, and explore segmentation networks using the produced heatmap. We also intend to involve superpixel-based methods for the fast and precise localization of individual palms in orthomosaic images, enhancing our understanding of their distribution [6]. Additionally, we aim to augment our approach by integrating UAV imagery with satellite data from platforms like Planet or WorldView, capitalizing on the complementary strengths of local detail and broad-scale context to improve detection accuracy [2].

REFERENCES

- Luke Browne and Jordan Karubian. 2016. Diversity of Palm Communities at Different Spatial Scales in a Recently Fragmented Tropical Landscape. Botanical Journal of the Linnean Society 182, 2 (2016), 451–464.
- [2] Seda Camalan, Kangning Cui, Victor Pauca, Sarra Alqahtani, Miles Silman, Raymond Chan, Evan Dethier, Luis Fernandez, David Lutz, and Robert Plemmons. 2022. Detecting Change Due to Alluvial Gold Mining in Peruvian Rainforest Using Recursive Convolutional Neural Networks and Contrastive Learning. In AGU Fall Meeting Abstracts, Vol. 2022. AGU, Chicago, USA, B52G-0905.
- [3] Seda Camalan, Kangning Cui, Victor Pauca, Sarra Alqahtani, Miles Silman, Raymond Chan, Robert Plemmons, Evan Dethier, Luis Fernandez, and David Lutz. 2022. Change Detection of Amazonian Alluvial Gold Mining Using Deep Learning and Sentinel-2 Imagery. Remote Sensing 14, 7 (2022), 1746.
- [4] Khai Loong Chong, Kasturi Devi Kanniah, Christine Pohl, and Kian Pang Tan. 2017. A Review of Remote Sensing Applications for Oil Palm Studies. Geo-spatial Information Science 20, 2 (2017), 184–200.

- [5] Kangning Cui, Seda Camalan, Ruoning Li, Victor Pauca, Sarra Alqahtani, Robert Plemmons, Miles Silman, Evan Nylen Dethier, David Lutz, and Raymond Chan. 2022. Semi-Supervised Change Detection of Small Water Bodies Using RGB and Multispectral Images in Peruvian Rainforests. In 2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS). IEEE, Rome, Italy, 1–5.
- [6] Kangning Cui, Ruoning Li, Sam Polk, Yinyi Lin, Hongsheng Zhang, James Murphy, Robert Plemmons, and Raymond Chan. 2023. Superpixel-based and Spatiallyregularized Diffusion Learning for Unsupervised Hyperspectral Image Clustering. arXiv preprint arXiv:2312.15447 (2023).
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition. IEEE, Miami, USA, 248–255.
- [8] Ms RS Sandhya Devi, VR Vijay Kumar, and P Sivakumar. 2021. A Review of Image Classification and Object Detection on Machine Learning and Deep Learning Techniques. In 2021 5th International Conference on Electronics, Communication and Aerospace Technology (ICECA). IEEE, Coimbatore, India, 1–8.
- [9] Yago Diez, Sarah Kentsch, Motohisa Fukuda, Maximo Larry Lopez Caceres, Koma Moritake, and Mariano Cabezas. 2021. Deep Learning in Forestry Using UAV-Acquired RGB Data: A Practical Review. Remote Sensing 13, 14 (2021), 2837.
- [10] Fabian Ewald Fassnacht, Hooman Latifi, Krzysztof Stereńczak, Aneta Modzelewska, Michael Lefsky, Lars Waser, Christoph Straub, and Aniruddha Ghosh. 2016. Review of Studies on Tree Species Classification from Remotely Sensed Data. Remote Sensing of Environment 186 (2016), 64–87.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. IEEE, Las Vegas, USA, 770–778.
- [12] Mohamad Haniff Junos, Anis Salwa Mohd Khairuddin, Subbiah Thannirmalai, and Mahidzal Dahari. 2022. Automatic Detection of Oil Palm Fruits from UAV Images Using an Improved YOLO Model. The Visual Computer 38 (2022), 2341–2355.
- [13] Nuzhat Khan, Mohamad Anuar Kamaruddin, Usman Ullah Sheikh, Yusri Yusup, and Muhammad Paend Bakht. 2021. Oil Palm and Machine Learning: Reviewing One Decade of Ideas, Innovations, Applications, and Gaps. Agriculture 11, 9 (2021), 832.
- [14] Zailani Khuzaimah, Nazmi Mat Nawi, Siti Nooradzah Adam, Bahareh Kalantar, Okoli Jude Emeka, and Naonori Ueda. 2022. Application and Potential of Drone Technology in Oil Palm Plantation: Potential and Limitations. *Journal of Sensors* 2022 (2022), 1–18.
- [15] Kuryati Kipli, Salleh Osman, Annie Joseph, Hushairi Zen, Dayang Nur Salmi Dharmiza Awang Salleh, Asrani Lit, and Kho Lee Chin. 2023. Deep Learning Applications for Oil Palm Tree Detection and Counting. Smart Agricultural Technology 5 (2023), 100241.
- [16] Weijia Li, Haohuan Fu, Le Yu, and Arthur Cracknell. 2016. Deep Learning Based Oil Palm Tree Detection and Counting for High-Resolution Remote Sensing Images. Remote Sensing 9, 1 (2016), 22.
- [17] Xinni Liu, Kamarul Hawari Ghazali, Fengrong Han, and Izzeldin Ibrahim Mohamed. 2021. Automatic Detection of Oil Palm Tree from UAV Images Based on the Deep Learning Method. Applied Artificial Intelligence 35, 1 (2021), 13–24.
- [18] Sarah Lueder, Kaushik Narasimhan, Jorge Olivo, Domingo Cabrera, Juana Jurado, Lewis Greenstein, and Jordan Karubian. 2022. Functional Traits, Species Diversity, and Species Composition of a Neotropical Palm Community Vary in Relation to Forest Age. Frontiers in Ecology and Evolution 10 (2022), 678125.
- [19] Salim Malek, Yakoub Bazi, Naif Alajlan, Haikel AlHichri, and Farid Melgani. 2014. Efficient Framework for Palm Tree Detection in UAV Images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 7, 12 (2014), 4692–4703
- [20] Willintong Marin, Ivan Mondragon, and Julian Colorado. 2022. Aerial Identification of Amazonian Palms in High-Density Forest Using Deep Learning. Forests 13. 5 (2022), 655.
- [21] Lucas Prado Osco, José Marcato Junior, Ana Paula Marques Ramos, Lúcio André de Castro Jorge, Sarah Narges Fatholahi, Jonathan de Andrade Silva, Edson Takashi Matsubara, Hemerson Pistori, Wesley Nunes Gonçalves, and Jonathan Li. 2021. A Review on Deep Learning in UAV Remote Sensing. International Journal of Applied Earth Observation and Geoinformation 102 (2021), 102456.
- [22] David Rumelhart, Geoffrey Hinton, Ronald Williams, et al. 1985. Learning Internal Representations by Error Propagation.
- [23] Yiran Wang, Xiaolin Zhu, and Bo Wu. 2019. Automatic Detection of Individual Oil Palm Trees from UAV Images Using HOG Features and an SVM Classifier. International Journal of Remote Sensing 40, 19 (2019), 7356–7370.
- [24] Karl Weiss, Taghi Khoshgoftaar, and DingDing Wang. 2016. A Survey of Transfer Learning. Journal of Big Data 3, 1 (2016), 1–40.
- [25] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. 2019. Object Detection with Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems 30, 11 (2019), 3212–3232.
- [26] Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. 2023. Object Detection in 20 Years: A Survey. Proc. IEEE 111, 3 (2023), 257–276.