



FIgLib & SmokeyNet: Dataset and Deep Learning Model for Real-Time Wildland Fire Smoke Detection

Anshuman Dewangan¹, Yash Pande¹, Hans-Werner Braun², Frank Vernon³, Ismael Perez², Ilkay Altintas², Garrison W. Cottrell¹ and Mai H. Nguyen^{2,*}

- ¹ Computer Science and Engineering Department, University of California, San Diego, CA 92093, USA; adewangan@ucsd.edu (A.D.); ypande@ucsd.edu (Y.P.); gary@ucsd.edu (G.W.C.)
- ² San Diego Supercomputer Center, University of California, San Diego, CA 92093, USA;
 - hwb@ucsd.edu (H.-W.B.); i3perez@sdsc.edu (I.P.); ialtintas@ucsd.edu (I.A.)
- ³ Scripps Institution of Oceanography, University of California, San Diego, CA 92093, USA; flvernon@ucsd.edu
 - Correspondence: mhnguyen@ucsd.edu

Abstract: The size and frequency of wildland fires in the western United States have dramatically increased in recent years. On high-fire-risk days, a small fire ignition can rapidly grow and become out of control. Early detection of fire ignitions from initial smoke can assist the response to such fires before they become difficult to manage. Past deep learning approaches for wildfire smoke detection have suffered from small or unreliable datasets that make it difficult to extrapolate performance to real-world scenarios. In this work, we present the Fire Ignition Library (FIgLib), a publicly available dataset of nearly 25,000 labeled wildfire smoke images as seen from fixed-view cameras deployed in Southern California. We also introduce SmokeyNet, a novel deep learning architecture using spatiotemporal information from camera imagery for real-time wildfire smoke detection. When trained on the FIgLib dataset, SmokeyNet outperforms comparable baselines and rivals human performance. We hope that the availability of the FIgLib dataset and the SmokeyNet architecture will inspire further research into deep learning methods for wildfire smoke detection, leading to automated notification systems that reduce the time to wildfire response.

Keywords: wildland fire mitigation; smoke detection; deep learning; computer vision; artificial intelligence; machine learning; remote sensing; HPWREN

1. Introduction

Climate change has had a devastating impact on California in the form of increased wildfire activity and intensity. In 2018 alone, 8527 fires burned an area of 1.9 million acres in California (7700 km²; nearly 2% of the state's area), with an estimated economic cost of USD 148.5 billion [1]. It is therefore imperative to detect and react to fire ignitions before they grow out of control. Currently, fire management teams rely on a combination of fire lookout experts, human-monitored camera feeds, and public reports to detect fire ignitions. However, it can take much longer than the first few crucial minutes for a fire to be reported using these existing methods, especially in areas with less human activity. Automated techniques using computer vision and deep learning hold promise in addressing this need. Deep learning-based wildfire smoke detection systems can accurately and consistently detect wildfires and provide valuable intel to reduce the time to alert authorities.

The goal of a wildfire smoke detection system can be structured as a binary image classification problem to determine the presence of smoke within a sequence of images. Priorities include quick time-to-detection, high recall to avoid missing potential fires, high precision to avoid frequent alarms that undermine trust in the system [2], and efficient performance to operate in real time on edge devices. However, the task proves challenging in real-world scenarios given the transparent and amorphous nature of smoke; faint, small, or dissipating smoke plumes; and false positives from clouds, fog, and haze. While the



Citation: Dewangan, A.; Pande, Y.; Braun, H.-W.; Vernon, F.; Perez, I.; Altintas, I.; Cottrell, G.W.; Nguyen, M.H. FIgLib & SmokeyNet: Dataset and Deep Learning Model for Real-Time Wildland Fire Smoke Detection. *Remote Sens.* 2022, *14*, 1007. https://doi.org/10.3390/rs14041007

Academic Editor: Carmen Quintano

Received: 16 December 2021 Accepted: 14 February 2022 Published: 18 February 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). idea of an automated wildfire smoke detection system has been previously explored, the difficulty of acquiring a large, labeled wildfire smoke dataset has limited researchers to using small or unbalanced datasets [3,4], manually searching for images online [4–7], or synthetically generating datasets [7–9].

In this work, we propose the following contributions to address the need for a consistent evaluation benchmark for real-world performance: (1) Fire Ignition Library (FIgLib), a publicly-available dataset of nearly 25,000 labeled wildfire smoke images captured in Southern California, and (2) SmokeyNet, a novel deep-learning-based model using image sequences for real-time wildfire smoke detection. We begin with a review of historical methods of wildfire smoke detection. We then provide details on the FIgLib dataset, the SmokeyNet architecture, our training procedure, and our experimental setup. After a discussion of the results of SmokeyNet's performance against comparable baselines, including human classification performance, we conclude with directions for future work.

2. Related Work

Before the rise in popularity of deep learning methods, computer vision algorithms leveraging hand-crafted features identified that the visual (e.g., color), spatial, and temporal (i.e., motion) qualities of smoke are essential for the machine detection of wildfires [3,10–12]. More recently, deep learning approaches use a combination of convolutional neural networks (CNNs) [5–7,13–17], background subtraction [13,16,18], and object detection methods [4,8,17,19,20] to incorporate visual and spatial features. Long short-term memory (LSTM) networks [4,16] or optical flow [14,18,21] methods have been applied to incorporate temporal context from video sequences.

Despite many papers reporting above 90% image classification accuracy for the detection of smoke, the lack of a large, labeled publicly available wildfire smoke dataset makes it difficult to compare performance between approaches. For example, Ko [3] and Jeong [4] used only 10 and 24 videos in their test sets, respectively, in which half the videos have smoke and half the videos have no smoke; there are no videos in which a fire starts in the middle of the sequence. With hundreds of frames per video but so few fire scenes evaluated, the high accuracies reported may not be representative of real-world performance across different scenarios. Li [6] and Park [7] collected 4595 (36% positive) and 6354 (22% positive) wildfire images online, respectively; however, since these images were not from video sequences, the smoke plumes in the image are more visible, and likely easier to detect, compared to wildfire smoke initially forming after ignition when seen from a continuous video sequence. Yin [5] also manually acquired images online, but the images represent smoke from a variety of indoor and outdoor scenarios beyond wildland fires. Many works synthetically generate images to overcome the lack of available data; Park [7] uses generative adversarial networks (GANs), Zhang [8] uses live smoke in front of a green screen, and Yuan [9] uses computational simulation to generate these images. Given the diversity and challenges of the datasets across these works, it is hard to identify which models are best or choose a particular dataset to use as a benchmark for wildfire smoke classification.

Govil et al. [2] is the only work we are aware of that also uses the FIgLib dataset to evaluate wildfire smoke detection performance. They used an InceptionV3 CNN [22] trained from scratch as the primary image classification architecture. Instead of using a sigmoid threshold of 0.5 for the prediction of smoke, as is common in classification models, a dynamic threshold was implemented based on the average prediction during the same time of day over the prior three days. This was used to incorporate periodic environmental events (e.g., fog, solar reflection, smog, haze); the data from prior days is not included in the FIgLib dataset, but is available through the HPWREN Archive. An analysis of results from this work is discussed in Section 7.3.

3. Data *3.1. FIgLib Dataset*

The FIgLib dataset addresses the need for a large, labeled publicly-available dataset for wildfire smoke detection. FIgLib reflects sequences of wildland fire images as seen from fixed-view cameras, part of the High Performance Wireless Research and Education Network (HPWREN), on remote mountain tops in Southern California. As of December 2021, the dataset consists of 315 fire sequences from 101 cameras across 30 stations occurring between June 2016 and July 2021. Each sequence typically contains images from 40 min prior to and 40 min following the start of the fire, serving as binary smoke/no-smoke labels for each image, and are spaced approximately 60 s apart for a total of 81 images per fire sequence. However, 114 fires are missing an average of 6.6 images each; missing images are either at the beginning, end, or randomly dispersed throughout the sequence. In total, the dataset contains 24,800 high-resolution images that are 1536 \times 2048 or 2048×3072 pixels in size, depending on the camera model used. The ignition detection and view prior to the ignition are enabled by a cluster deployment of cameras, where four 90+ degree views stay consistent for years, covering 360 degrees around a mountaintop. Examples from FIgLib can be seen in Figure 1, and the full dataset can be accessed at the following link: http://hpwren.ucsd.edu/HPWREN-FIgLib/ (accessed on 16 December 2021).





Figure 1. Images from the FIgLib dataset: (**a**) no smoke with strong glare, (**b**) no smoke with misleading haze, (**c**) very apparent wildfire smoke, and (**d**) faint wildfire smoke.

3.2. Data Preparation

The number of fires and images in the train, validation, and test splits of the FIgLib dataset for our machine learning task of wildfire smoke detection are shown in Table 1. To avoid out-of-distribution sequences, we removed fires with black and white images (N = 10), night fires (N = 19), and fires with questionable presence of smoke (N = 16), including one fire with 180 images and no labels, from the dataset (3700 images from 45 fires removed in total). In addition to binary smoke/no-smoke labels for each image, the smoke in 144 fires has been manually annotated with bounding boxes and contour

masks. Since we divide the images into tiles (described in Section 4.1), we can use these annotations to provide smoke/no-smoke labels at a more granular level. Hence, we used images from these 144 annotated fires for training (53.3% of eligible fires, 11,300 images); the remaining 126 fires (9800 images) were split between the validation and test sets (collectively, the evaluation sets) such that the number of images in each is roughly equivalent. The exact list of fires omitted and used for the train, validation, and test splits can be accessed at the following link: https://gitlab.nrp-nautilus.io/-/snippets/63 (accessed on 16 December 2021).

| Model | # Fires | # Images | | |
|------------|---------|----------|--|--|
| Train | 144 | 11.3 K | | |
| Validation | 64 | 4.9 K | | |
| Test | 62 | 4.9 K | | |
| Omitted | 45 | 3.7 K | | |
| Total | 315 | 24.8 K | | |

Table 1. Cross-validation splits of the FIgLib dataset.

While it is uncommon to use such a low percentage of the data for the training set relative to that of the evaluation sets (an 80/10/10 split is more common), our model requires annotations for training and we only had access to 144 annotated fires. We also did not want to omit any fires from the evaluation sets, because larger evaluation sets provide a more robust representation of how the model would perform in diverse conditions. Splitting the data by fires instead of images ensures that no data related to the test set are in the training set, to obtain a more representative evaluation of the model's real-world performance. Otherwise, the model would be evaluated on images of fires for which it has already trained on using other frames of the video sequence; the performance of the model in this scenario might be overstated since the model might learn features specific to the particular video sequence that do not extrapolate to completely unseen scenarios.

We perform the following transformations during data loading to improve the performance of our model. We first resize the images to the empirically-determined size of 1392×1856 pixels to improve training and inference speed. We also crop the top 352 rows of the image, specifically determined to ensure that only rows well above the horizon are cropped to reduce false positives from clouds for additional performance gains. Resizing and further cropping the height of the images to a final size of 1040×1856 pixels enables us to evenly divide the image into overlapping 224×224 tiles (see Section 4.1). We then randomly apply data augmentations including horizontal flip, vertical crop, color jitter, brightness and contrast jitter, and blur. Finally, we normalize the images to 0.5 mean and 0.5 standard deviation, as expected by the deep learning package used (torchvision).

4. Methods

4.1. Tiling

Our goal is the binary classification of images to determine the presence of smoke as early in the sequence as possible. Training the model with standard CNN techniques by leveraging solely image labels does not provide a sufficient training signal for the model to identify small plumes of smoke within the large images. Object detection models using bounding box and contour mask annotations can better localize the target object using anchors and a regression head [23]; however, these models require precise annotations, which poses a challenge in our scenario given the amorphous and transparent nature of smoke.

Consequently, we build upon previous work by tiling the image into 224×224 tiles, overlapping by 20 pixels for a total of 45 tiles [2]. We also generate corresponding binary *tile labels*: positive if the number of pixels of smoke in the tile, determined by the filled

polygon of the contour mask, is greater than an empirically-determined smoke detection threshold of 250 (0.5% of the total pixels in the tile). Tile labels provide the entirety of our localized feedback signal; we do not otherwise use the bounding box or contour mask annotations during training.

One challenge of the dataset is that 1213 (approximately 20%) of the positive images are missing contour mask annotations. A total of 280 annotations are missing because the smoke is difficult to see, generally occurring at the beginning of the fire sequence or at the end, when the smoke has dissipated. A total of 486 annotations are missing contour masks but have bounding box annotations, generally because the smoke is too small to reasonably outline a fine contour mask. The remaining 447 missing annotations are randomly spread throughout the fires.

For images with bounding box annotations where contour masks are not available, we determined the tile labels by filling the bounding boxes as polygons instead of the contour masks (486 images affected). We attempted other methods to incorporate feedback from images with missing annotations, including using feedback from only image labels (as opposed to both image and tile labels) and copying contour masks from the closest available image in the sequence. However, neither of these methods improved model performance; consequently, we did not train on the remaining positive images with missing annotations (727 images total). For future work, we aim to resolve these missing labels for more robust training data.

4.2. Smokeynet Architecture

The SmokeyNet architecture (Figure 2) is a novel spatiotemporal gridded image classification approach for wildfire smoke detection combining three different types of neural networks: a CNN [24], an LSTM [25], and a vision transformer (ViT) [26]. The input to our model is the tiled raw image and its previous frame in the wildfire video sequence to incorporate the motion of the smoke. A CNN, pretrained on the ImageNet dataset [27], initially extracts representations of the raw image pixels from each tile of the two frames independently. A ResNet34, a lighter-weight version of the popular ResNet50 model, is our preferred choice of CNN backbone [28]. Then, an LSTM combines the temporal information of each tile from the current frame with its counterpart from the previous frame. Finally, all temporally-combined tiles are fed into a ViT, which incorporates spatial information *across* tiles to improve the image prediction.



Figure 2. The SmokeyNet architecture takes two frames of the tiled image sequence as input and combines a CNN, LSTM, and ViT. The yellow blocks denote "tile heads" used for intermediate supervision, while the blue block denotes the "image head" used for the final image prediction.

The outputs of the ViT are spatiotemporal embeddings for each tile, as well as a CLS token embedding that summarizes representations for the whole image [26]. The CLS token embedding is passed to an "image head", consisting of three fully-connected layers with ReLU activation with output sizes of 256, 64, and 1, respectively, and a sigmoid layer with a threshold of 0.5 to generate a single prediction for the whole image. Given the modular nature of each of the components, we can experiment with different approaches to capture spatiotemporal information while still training the model end-to-end. The full SmokeyNet

codebase, including code to run all experiments conducted in this work, can be accessed at the following link: https://gitlab.nrp-nautilus.io/anshumand/pytorch-lightning-smoke-detection (accessed on 16 December 2021).

4.3. Loss

The initial component of our loss applies standard binary cross-entropy (BCE) loss between the outputs of the image head and the ground-truth binary image labels. We can increase the weight of positive examples when calculating this BCE image loss to trade off precision for higher recall. Increasing the positive weight increases the penalty for missing positive examples; while potentially incurring more false positives, the model will also be able to detect the actual presence of smoke more quickly and accurately, which is of utmost importance. We use the empirically-determined positive weight of five to achieve more balanced precision and recall and improve the overall accuracy and F1-score.

To leverage the localized information provided by the tile labels, we also apply intermediate supervision to each of the model components [29]. Since the model's components, the CNN, LSTM, and ViT, also produce embeddings on a per-tile basis, we pass each component's embeddings through individual "tile heads", consisting of three fully-connected layers with ReLU activation with output sizes of 256, 64, and 1, respectively, and a sigmoid layer to generate predictions for each *tile*. We then apply BCE loss between the outputs of the tile heads and the binary tile labels. To address the class imbalance in which negative tiles occur more frequently than positive tiles, we weight positive examples by 40, the ratio of negative tiles to positive tiles.

If *I* is the total number of tiles, the overall training loss can be summarized as

$$loss = BCE^{image} + \sum_{i}^{I} \{BCE_{i}^{CNN} + BCE_{i}^{LSTM} + BCE_{i}^{ViT}\}$$

Since we have tile labels for only the training data, we define our validation loss as the average number of *image* prediction errors and use this validation loss for early stopping. The BCE loss equations are elaborated upon in Appendix B.

5. Experiments

5.1. Comparable Baselines

We experiment with alternate CNN backbones to the ResNet34, including a MobileNetV3Large (denoted "MobileNet") [30], MobileNet with a feature pyramid network (FPN) [31] to better incorporate spatial scales, EfficientNet-B0 [32], and Data Efficient Image Transformer (DeiT-Tiny) [33]. Using the ResNet34 as the backbone, we also try inputting three frames (i.e., two additional frames of temporal context) instead of two and conduct an ablation study by removing different parts of the model to evaluate each component's benefits. We then experiment with different architectures that can capture the temporal information from sequential frames, including replacing the LSTM with a transformer [34]; using a CNN and a ResNet18-3D CNN to replace both the LSTM + ViT [35]; and incorporating motion information using MOG2, a Gaussian-mixture-based background removal method [36,37], as an additional channel of input. Finally, we compare the model's performance to three baseline architectures: ResNet50, the standard for image classification models [28]; faster-RCNN, a standard object detection model [23]; and mask-RCNN, an image segmentation model leveraging both contour masks as well as bounding boxes for training signal [38].

For baseline models that do not use a ViT as the last architectural component (e.g., ResNet34 + LSTM, ResNet50, ResNet34 + ResNet18-3D, etc.), there is no CLS token embedding that summarizes representations for the whole image that we can use for our image prediction. Consequently, we determine the overall image prediction by passing the model's tile predictions into a single fully connected layer with sigmoid activation, outputting a single prediction for the image. We also experimented with the simple decision rule that if the prediction for any tile is positive for smoke, the full image is also classified as positive; however, this resulted in worse performance. Image predictions for object detection models (e.g., faster-RCNN, mask-RCNN) were determined as positive if the model

predicted any bounding box with a confidence score above the empirically-determined threshold of 0.5 ({0, 0.2, 0.4, 0.5, 0.6} all tested). Additional model implementation details for alternative architectures are described in Appendix C.

5.2. Training Details

Hyperparameter tuning was performed sequentially, with the best result from one set of experiments used in subsequent experiments, in the following order: learning rate ({1e-2, **1e-3**, 1e-4}), weight decay ({1e-4, **1e-3**}), image resizing ({100%, **90%**, 80%, 50%} of 1536 \times 2048), smoke detection threshold to determine tile labels ({0, 10, 100, **250**} pixels per tile), dropout ({**0**, 0.1}), and image BCE loss positive weight ({1, 2, 5, 10}). Final models were trained using an SGD optimizer with learning rate 0.001, weight decay 0.001, and no dropout. The batch size used was the larger of 2 or 4, depending on which would fit into GPU memory, and gradient batches were accumulated such that the effective batch size was 32. Models were trained for 25 epochs using a single NVIDIA 2080Ti GPU; the model with the lowest validation loss was used for evaluation on the test set.

5.3. Evaluation Metrics

For each experiment, we report the following evaluation metrics typical for binary classification problems:

Accuracy =
$$\frac{TP+TN}{TP+TN+FP+FN}$$
 Precision = $\frac{TP}{TP+FP}$ Recall = $\frac{TP}{TP+FN}$ F1 = $\frac{2*Precision*Recall}{Precision+Recall}$

derived from true positives (*TP*), false positives (*FP*), true negatives (*TN*), and false negatives (*FN*) calculated between the model predictions and the ground truth labels. We also report the average time-to-detection, calculated as the number of minutes until the model correctly predicts the first positive frame of a wildfire sequence, averaged over all fires. We include the number of parameters (in millions) and inference time (ms/image) of each model, which should be minimized for deployment to edge devices.

6. Results

6.1. Experimental Results

Table 2 reports test evaluation performance for each of the experimental architectures with two frames of input (unless otherwise stated) as an average over five runs. The SmokeyNet architecture with a ResNet34 backbone and two frames of input achieves an image accuracy of 83.49% and F1-score of 82.59% while delivering on the objectives of high precision (89.84%), high recall (76.45%), fast performance (51.6ms/image), and low average time to detection (3.12 min). One additional frame of input (ResNet34 + LSTM + ViT (3 frames)) only marginally improves performance at the cost of a 55.6% increase in inference time. Large backbones, such as the ResNet34 or EfficientNet-B0, trade off model size and inference time for better accuracy compared to smaller backbones, such as the MobileNet or MobileNetFPN.

From the ablation study, we observe that the standalone CNN or CNN + LSTM models perform poorly at the task. Adding the ViT to the CNN significantly improves performance with little impact to inference speed. All three alternate architectures to incorporate temporal information perform slightly worse than SmokeyNet; however, the ResNet34 + ResNet18-3D architecture provides another viable alternative if prioritizing model size. Comparing MobileNet + LSTM + ViT to MobileNet + LSTM + ViT(MOG2), we see that the addition of background subtraction improves performance by almost 2%, mainly by improving recall by over 3%, while maintaining high precision. This comes at the cost of model size and inference time, but it is worth exploring adding MOG2 as another SmokeyNet variant in the future.

| Model | Params (M) | Time (ms/it) | Α | F1 | Р | R | TTD (mins) |
|----------------------------------|---------------|-----------------|-------|-------|-------|-------|---------------|
| Variants of SmokeyNet: | | | | | | | |
| ResNet34 + LSTM + ViT | 56.9 | 51.6 | 83.49 | 82.59 | 89.84 | 76.45 | 3.12 |
| ResNet34 + LSTM + ViT (3 frames) | 56.9 | 80.3 | 83.62 | 82.83 | 90.85 | 76.11 | 2.94 |
| MobileNet + LSTM + ViT | 36.6 | 28.3 | 81.79 | 80.71 | 88.34 | 74.31 | 3.92 |
| MobileNetFPN + LSTM + ViT | 40.4 | 32.5 | 80.58 | 80.68 | 82.36 | 79.12 | 2.43 |
| EfficientNetB0 + LSTM + ViT | 52.3 | 67.9 | 82.55 | 81.68 | 88.45 | 75.89 | 3.56 |
| TinyDeiT + LSTM + ViT | 22.9 | 45.6 | 79.74 | 79.01 | 84.25 | 74.44 | 3.61 |
| ResNet34 (1 frame) | 22.3 | 29.7 | 79.40 | 78.90 | 81.62 | 76.58 | 2.81 |
| ResNet34 + LSTM | 38.9 | 53.3 | 79.35 | 79.21 | 82.00 | 76.74 | 2.64 |
| ResNet34 + ViT (1 frame) | 40.3 | 30.8 | 82.53 | 81.30 | 88.58 | 75.19 | 2.95 |
| ResNet34 + Transformer + ViT | 58.9 | 50.6 | 82.74 | 81.39 | 90.69 | 73.95 | 3.18 |
| ResNet34 + ResNet18-3D | 38.0 | 57.5 | 83.10 | 82.26 | 88.91 | 76.65 | 2.87 |
| MobileNet + LSTM + ViT (MOG2) | 57.3 | 55.5 | 83.12 | 82.65 | 88.79 | 77.43 | 3.72 |
| ResNet50 (1 frame) | 26.1 | 50.4 | 68.51 | 74.30 | 63.35 | 89.89 | 1.01 |
| FasterRCNN (1 frame) | 41.3 | 55.6 | 71.56 | 66.92 | 81.34 | 56.88 | 5.01 |
| MaskRCNN (1 frame) | 43.9 | 56.9 | 73.24 | 69.94 | 81.08 | 61.51 | 4.18 |

Table 2. Accuracy (A), F1, precision (P), recall (R), and average time-to-detection (TTD) evaluation metrics on the test set. Best results are **bolded**.

The SmokeyNet architecture clearly outperforms standard image classification, object detection, and image segmentation baselines. All models have acceptable time to detection—most within 3–4 min. While the ResNet50 is clearly the fastest to detect a fire, it only accomplishes this by generating an unacceptable amount of false positives.

6.2. Performance Visualization

Figure 3 (on the following page) further visualizes SmokeyNet's performance on images from the test set. Additionally, a video of the model's performance per image can be viewed at the following link: https://youtu.be/cvXQJao3m1k (accessed on 16 December 2021). The model performs well in a variety of real-world scenarios, correctly identifying apparent smoke plumes while avoiding clouds and haze. However, the model still makes systematic misclassifications of low-altitude clouds as false positives (20200806_BorderFire_om-e-mobo-c, 20200828_BorderFire_sm-s-mobo-c, 20200806_BorderFire_lp-s-mobo-c). The model completely missed the 20200930_DeLuzFire_rm-w-mobo-c fire (top row of Figure 3) because the smoke occurs directly behind a transmission tower.

6.3. Human Performance Baseline

Due to the lack of suitable benchmarks for performance, we additionally measured human performance of smoke classification on the FIgLib dataset. Participants were three lab members experienced in classifying images for the presence of wildfire smoke. For the experimental setup, one image from each of the 62 fires from the test set was randomly selected for prediction. Participants were presented with the images for prediction, each preceded by the previous frame of the image sequence, to replicate the temporal context our machine learning model receives and a real-time inference scenario. The participants then recorded if they believed wildfire smoke was present in the image.

The experimental setup proves challenging due to differences from how the ground truth labels were generated. The ground truth labels were created by a human expert who had full temporal context, including access to all past and *future* frames of a video sequence, to precisely determine the first frame in which smoke is visible. Consequently, it is difficult to correctly identify positive images early in the fire sequence when the smoke is faint or small with only a single preceding frame and no additional temporal context to see how the smoke grows over time.



Figure 3. SmokeyNet's performance per fire on both negative and positive images. Green denotes a correct prediction; red denotes an incorrect prediction; white denotes images missing from the sequence. Hence, red on the left are false positives, red on the right are misses. Common misses include faint smoke occurring at the start of the fire or dissipating smoke at the end of the fire sequence. Common false positives include low-altitude clouds and haze.

The three participants achieved an average accuracy of 78.5% ($\sigma = 1.52\%$), F1-score of 82.8% ($\sigma = 0.73\%$), precision of 93.5% ($\sigma = 4.66\%$), and recall of 74.4% ($\sigma = 1.90\%$). The low accuracy and recall signify the false negatives from missing positive images early in the fire sequence. SmokeyNet achieved a higher accuracy and similar F1-score compared to human performance. The model also had higher recall at the expense of lower precision. If the false positives derived from low-altitude clouds could be corrected, SmokeyNet would achieve 85.8% accuracy and 94.8% precision, matching the precision and further surpassing the accuracy of human performance. Future work should focus on this issue.

7. Discussion

7.1. Innovations in the SmokeyNet Architecture

SmokeyNet establishes a strong baseline for performance on the FIgLib dataset for the task of wildfire smoke detection. With a finely-tuned custom architecture, SmokeyNet improves upon works that leverage standard out-of-the-box models for image classification or smoke detection [2,17,20]. As the first transformer-based architecture for wildfire smoke detection that we are aware of, SmokeyNet should see improved performance over CNNs when training on even larger datasets [34]. SmokeyNet also includes features that enable quick experimentation for further improvement. Modularity allows the user to easily swap out components of the SmokeyNet architecture to test alternate backbones or temporal aggregators. Intermediate supervision mitigates the likelihood of unstable training even with very large models [29]. Lastly, the tiling of input images enables SmokeyNet to adapt to input images of any size.

7.2. Comparison of FIgLib to Previous Datasets

The structure of the FIgLib dataset enables a more robust evaluation of machine learning models for wildfire smoke detection in real-world scenarios than datasets in previous work. Since FIgLib depicts fires from only wildland areas, models trained on FIgLib for wildfire smoke detection avoid a data distribution shift that would occur when training models on datasets with both indoor and outdoor scenes [5]. FIgLib's use of video sequences, as opposed to the static images used in Li et al. [6] and Park et al. [7], allows models to leverage temporal data to improve predictions. FIgLib includes frames in sequences before and after the initial fire ignition, encouraging models to learn to detect the ignition itself, when smoke plumes are small. This is of great practical importance, as fast time-to-detection can prevent fires from getting out of control. These features contrast with the datasets used in Ko et al. [3] and Jeong et al. [4], in which the videos contain either all smoke or no smoke, making it difficult for a model to learn to detect fire ignition. Furthermore, FIgLib represents an order of magnitude more data, 315 fire sequences from 101 cameras, versus the 10 and 24 videos used in previous work [3,4]. This variety ensures models are trained and evaluated in a number of different scenarios to avoid overfitting. Going forward, we plan to continue adding more recent fires to the FIgLib dataset, expanding the HPWREN camera network, and, upon further validation, releasing the human-annotated bounding boxes and contour masks to assist model training.

It is important to note that the FIgLib dataset contains roughly 50% positive examples and 50% negative examples representing frames 40 min before and after the start of each fire. While a balanced dataset makes it easy to train a machine learning model for binary classification, one limitation is that this ratio is not representative of real-world scenarios in which positive examples of visible smoke are much more rare than negative, no-smoke examples. This may induce models to incur a higher false positive rate on real-world data. To mitigate this, models can leverage unlabeled data from the HPWREN Archive (http://cl.hpwren.ucsd.edu/archive/, accessed on 16 December 2021) in combination with FIgLib to obtain a more representative sample of positive and negative examples to fine-tune the model for real-world deployment.

7.3. Comparison of SmokeyNet Performance to Previous Work

It is impossible to make a head-to-head performance comparison with previous work due to differences in the datasets used and evaluation metrics reported, as discussed in Section 2. Govil et al. [2], the only work to also use the FIgLib dataset, reported a test accuracy of 0.91 and F1-score of 0.89, higher than SmokeyNet's accuracy of 0.83 and F1-score of 0.82. However, their test set consisted of only 250 hand-selected images relative to the training set of about 8500 images. The test set also contained images from the same video sequences of fires used in the training set, with only 10 min (i.e., 10 frames) of separation between them. This procedure suggests that the test data may only be 10 min earlier or later than the training data, and therefore may overstate performance. Completely independent test data is an important methodological requirement in machine learning to measure true generalization performance. The authors also reported results from field-testing the model on 65 HPWREN cameras over a period of nine days in October 2019. After suppressing repeat detections in a one-hour timespan, only 21% of notifications showed smoke from real fires (i.e., a 79% false positive rate) and there was no report of how many actual fires were missed.

The standard in machine learning when comparing two models is that they use the same training, validation, and test sets. Consequently, despite SmokeyNet's numerically lower performance on the evaluation set, it is difficult to make a direct comparison with the performance of the model in Govil et al. [2]. Given the much larger training and test sets used, we believe SmokeyNet is more robust to unseen scenarios and more representative of real-world performance. Further experimentation should compare SmokeyNet to other architectures [2,5,6,14,19,21] using consistent training and evaluation splits from FIgLib

as the basis for fair comparison. Additionally, it is important to follow the example of Govil et al. [2] and test the model on unseen, real-world data.

7.4. Planned Future Work

For future work, we will continue improving the performance of SmokeyNet by reducing false positives in difficult scenarios such as low-altitude clouds and haze, a current limitation of the model. One method would be to take false positives from the holdout set and incorporate them into the training set as in Govil et al. [2]; however, in order to maintain the integrity of the holdout set, we would have to split our holdout set in two and leave one untouched when following this procedure. Advanced preprocessing techniques can extract candidate smoke patches before inputting the image into the model by using ViBE background removal [39], YUV color space segmentation [6], or a dark channel prior [13,21] for potential performance gains. Since fires frequently occur in similar locations and conditions during specific months of the year, incorporating fire location, date, and weather data from historical fire records can further improve predictions.

We also plan to utilize the large amount of unlabeled data from the HPWREN Archive. We can leverage SmokeyNet's performance to make predictions on unlabeled data, which can then be validated by human inspection. These validated predictions can then be used to quickly expand the amount of labeled training data. The unlabeled data can also be used in the self-supervised representation learning setting, using models such as DINO [40]. This approach allows the model to learn hidden representations from large amounts of unlabeled data to better perform in a classification task with limited labeled data. Since the vast majority of the unlabeled images will not contain fires, we can alternatively use generative adversarial networks (GANs) to synthetically generate smoke in these images to produce positive examples for additional training data as in Park et al. [7]. By effectively increasing the size of our training data and extracting better representations from the data, these approaches enable our model to make better predictions of the presence of smoke.

Lastly, we plan future work to reduce the model size for better compatibility with edge devices, another limitation of SmokeyNet, without sacrificing prediction performance. We can use pruning [41] to eliminate sparsely-used model weights to decrease model size. We can also use distillation to train a smaller deep network to perform similarly to our current model; this strategy will result in a model more suitable for edge computing [4,42].

8. Conclusions

FIgLib addresses the need for a large, labeled dataset that incorporates a variety of scenarios for the early detection of wildfire smoke. SmokeyNet provides a strong baseline for automated wildfire smoke detection. Since SmokeyNet is trained and evaluated on the publicly-available FIgLib dataset, its performance can be easily compared to subsequent approaches. We ultimately hope our contributions enable further research and development of automated notifications for wildfire smoke detection.

Author Contributions: Conceptualization, M.H.N., Y.P., A.D. and G.W.C.; methodology, A.D., Y.P., M.H.N. and G.W.C.; software, A.D. and Y.P.; validation, A.D., Y.P., M.H.N. and G.W.C.; formal analysis, A.D. and Y.P.; investigation, A.D. and Y.P.; resources, H.-W.B., F.V. and I.A.; data curation, H.-W.B., F.V. and I.P.; writing—original draft preparation, A.D.; writing—review and editing, A.D., M.H.N., G.W.C., I.A., H.-W.B., F.V. and Y.P.; visualization, A.D.; supervision, M.H.N. and G.W.C.; project administration, M.H.N.; funding acquisition, I.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded in part by NSF grant numbers 1730158, 2100237, 2120019 for Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI) and 1331615, 2040676 and 1935984 for WIFIRE, WIFIRE Commons, and SAGE.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: Publicly available datasets were analyzed in this study. This data can be found here: http://hpwren.ucsd.edu/HPWREN-FIgLib/, (accessed on 16 December 2021).

Acknowledgments: The authors would like to thank Brian Norton for sharing his invaluable expertise on wildfire management; Stephen Jarrell, Duolan Quyang, Atman Patel, and Ulyana Tkachenko for their collaboration and insights; Scott Ehling and his team of labellers at Argonne National Laboratory for contributing to dataset annotations; and the UC San Diego CHASE-CI team for computer support.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Summary of Hyperlinks

The following hyperlinks were referenced in the body of the work and are listed again here for convenience (in order of appearance):

- FIgLib Dataset: http://hpwren.ucsd.edu/HPWREN-FIgLib/, (accessed on 16 December 2021)
- Cross-Validation Splits of FIgLib Dataset: https://gitlab.nrp-nautilus.io/-/snippets/ 63, (accessed on 16 December 2021)
- Codebase for SmokeyNet & Related Experiments: https://gitlab.nrp-nautilus.io/ anshumand/pytorch-lightning-smoke-detection, (accessed on 16 December 2021)
- Visualization of SmokeyNet Performance: https://youtu.be/cvXQJao3m1k, (accessed on 16 December 2021)
- HPWREN Archive: http://c1.hpwren.ucsd.edu/archive/, (accessed on 16 December 2021)

Appendix B. Binary Cross-Entropy Loss Equations

Standard binary cross-entropy (BCE) loss for the two-class case is summarized in Equation (A1), in which N is the number of examples, y is the ground truth label, and p is the model prediction:

BCE =
$$-\frac{1}{N} \sum_{n=1}^{N} (y_n \log(p_n) + (1 - y_n) \log(1 - p_n))$$
 (A1)

The total loss used in SmokeyNet is summarized in Equation (A2), in which *I* is the total number of tiles:

$$loss = BCE^{image} + \sum_{i}^{I} \{BCE_{i}^{CNN} + BCE_{i}^{LSTM} + BCE_{i}^{ViT}\}$$
(A2)

Modified from the standard BCE loss equation, Equations (A3)–(A5) elaborate upon the tile losses of the CNN, LSTM, and ViT, in which y^{tile} are the tile labels, p^{CNN} are the outputs of the CNN, p^{LSTM} are the outputs of the LSTM, and p^{ViT} are the outputs of the ViT. Note that the weight of the positive examples is 40 to address class imbalance:

$$BCE^{CNN} = -\frac{1}{N} \sum_{n=1}^{N} \left(40 * y_n^{tile} \log(p_n^{CNN}) + (1 - y_n^{tile}) \log(1 - p_n^{CNN}) \right)$$
(A3)

$$BCE^{LSTM} = -\frac{1}{N} \sum_{n=1}^{N} \left(40 * y_n^{tile} \log(p_n^{LSTM}) + (1 - y_n^{tile}) \log(1 - p_n^{LSTM}) \right)$$
(A4)

$$BCE^{ViT} = -\frac{1}{N} \sum_{n=1}^{N} \left(40 * y_n^{tile} \log(p_n^{ViT}) + (1 - y_n^{tile}) \log(1 - p_n^{ViT}) \right)$$
(A5)

Finally, Equation (A6) expands upon the overall image loss, in which y^{image} are the image labels and p^{image} are the outputs of the image head corresponding to the final outputs of the model. Note that the positive weight is only 5 in this case:

$$BCE^{image} = -\frac{1}{N} \sum_{n=1}^{N} (5 * y_n^{image} \log(p_n^{image}) + (1 - y_n^{image}) \log(1 - p_n^{image}))$$
(A6)

Appendix C. Experimental Architecture Details

A feature pyramid network (FPN) is a computer vision architecture that better recognizes spatial scales by incorporating information from multiple layers of the CNN backbone [31]. Instead of producing a single 960-dimensional embedding like a standard MobileNetV3Large [30], the MobileNetFPN outputs 3 layers of 256-channel feature maps per tile sized 7×7 , 7×7 , and 4×4 respectively. We downsample each feature map through two convolutional layers of kernel size 1 and flatten the feature maps such that the total number of features per map is 784. We then concatenate all the feature maps and further downsample the concatenated features to 960 to match the embedding size of the standard MobileNetV3Large. These are the final embeddings that are passed onto the next component of the SmokeyNet architecture, the LSTM.

To incorporate MOG2 background removal as an additional input channel, we first take two sequential frames of the raw wildfire smoke video sequence and apply MOG2 background removal; this generates a single channel of dimensions equivalent to the raw image inputs. While two sequential frames of the tiled raw images are input into the CNN and LSTM of the SmokeyNet architecture, two sequential frames of the MOG2 channel are passed through a separate CNN and a separate LSTM. The embeddings resulting from both LSTMs, one for the raw image and one for the MOG2 channel, are then concatenated and passed through a single linear layer downsampling the feature maps by half before being passed onto the next component of the standard SmokeyNet architecture, the ViT.

References

- 1. Wang, D.; Guan, D.; Zhu, S.; Kinnon, M.M.; Geng, G.; Zhang, Q.; Zheng, H.; Lei, T.; Shao, S.; Gong, P.; et al. Economic footprint of California wildfires in 2018. *Nat. Sustain.* **2021**, *4*, 252–260. [CrossRef]
- Govil, K.; Welch, M.L.; Ball, J.T.; Pennypacker, C.R. Preliminary Results from a Wildfire Detection System Using Deep Learning on Remote Camera Images. *Remote Sens.* 2020, 12, 166. [CrossRef]
- Ko, B.C.; Kwak, J.Y.; Nam, J.Y. Wildfire smoke detection using temporospatial features and random forest classifiers. *Opt. Eng.* 2012, 51, 017208. [CrossRef]
- Jeong, M.; Park, M.; Nam, J.; Ko, B.C. Light-Weight Student LSTM for Real-Time Wildfire Smoke Detection. Sensors 2020, 20, 5508. [CrossRef] [PubMed]
- Yin, M.; Lang, C.; Li, Z.; Feng, S.; Wang, T. Recurrent convolutional network for video-based smoke detection. *Multimed. Tools Appl.* 2019, 78, 237–256. [CrossRef]
- 6. Li, T.; Zhao, E.; Zhang, J.; Hu, C. Detection of Wildfire Smoke Images Based on a Densely Dilated Convolutional Network. *Electronics* **2019**, *8*, 1131. [CrossRef]
- Park, M.; Tran, D.Q.; Jung, D.; Park, S. Wildfire-Detection Method Using DenseNet and CycleGAN Data Augmentation-Based Remote Camera Imagery. *Remote Sens.* 2020, 12, 3715. [CrossRef]
- Zhang, Q.X.; Lin, G.H.; Zhang, Y.M.; Xu, G.; Wang, J.J. Wildland Forest Fire Smoke Detection Based on Faster R-CNN using Synthetic Smoke Images. *Proceedia Eng.* 2018, 211, 441–446. [CrossRef]
- 9. Yuan, F.; Zhang, L.; Xia, X.; Wan, B.; Huang, Q.; Li, X. Deep smoke segmentation. Neurocomputing 2019, 357, 248–260. [CrossRef]
- 10. Ho, C.C. Machine vision-based real-time early flame and smoke detection. *Meas. Sci. Technol.* 2009, 20, 045502. [CrossRef]
- 11. Toreyin, B.U.; Cetin, A.E. Wildfire detection using LMS based active learning. In Proceedings of the 2009 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASP-09), Taipei, Taiwan, 19–24 April 2009; pp. 1461–1464.
- Genovese, A.; Labati, R.D.; Piuri, V.; Scotti, F. Wildfire smoke detection using computational intelligence techniques. In Proceedings of the 2011 IEEE International Conference on Computational Intelligence for Measurement Systems and Applications (CIMSA) Proceedings, Ottawa, ON, Canada, 19–21 September 2011; pp. 1–6.
- Luo, Y.; Zhao, L.; Liu, P.; Huang, D. Fire smoke detection algorithm based on motion characteristic and convolutional neural networks. *Multimed. Tools Appl.* 2018, 77, 15075–15092. [CrossRef]
- 14. Pundir, A.S.; Raman, B. Dual Deep Learning Model for Image Based Smoke Detection. *Fire Technol.* **2019**, *55*, 2419–2442. [CrossRef]

- 15. Ba, R.; Chen, C.; Yuan, J.; Song, W.; Lo, S. SmokeNet: Satellite Smoke Scene Detection Using Convolutional Neural Network with Spatial and Channel-Wise Attention. *Remote Sens.* **2019**, *11*, 1702. [CrossRef]
- 16. Cao, Y.; Yang, F.; Tang, Q.; Lu, X. An Attention Enhanced Bidirectional LSTM for Early Forest Fire Smoke Recognition. *IEEE Access* 2019, *7*, 154732–154742. [CrossRef]
- Khan, S.; Muhammad, K.; Hussain, T.; Del Ser, J.; Cuzzolin, F.; Bhattacharyya, S.; Akhtar, Z.; de Albuquerque, V.H.C. DeepSmoke: Deep learning model for smoke detection and segmentation in outdoor environments. *Expert Syst. Appl.* 2021, 182, 115125. [CrossRef]
- Yuan, J.; Wang, L.; Wu, P.; Gao, C.; Sun, L. Detection of Wildfires along Transmission Lines Using Deep Time and Space Features. Pattern Recognit. Image Anal. 2018, 28, 805–812. [CrossRef]
- Li, X.; Chen, Z.; Wu, Q.J.; Liu, C. 3D Parallel Fully Convolutional Networks for Real-Time Video Wildfire Smoke Detection. *IEEE Trans. Circuits Syst. Video Technol.* 2018, 30, 89–103. [CrossRef]
- Jindal, P.; Gupta, H.; Pachauri, N.; Sharma, V.; Verma, O.P. Real-Time Wildfire Detection via Image-Based Deep Learning Algorithm. In Soft Computing: Theories and Applications; Springer: Berlin/Heidelberg, Germany, 2021; pp. 539–550.
- Gupta, T.; Liu, H.; Bhanu, B. Early Wildfire Smoke Detection in Videos. In Proceedings of the 2020 25th International Conference on Pattern Recognition (ICPR-20), Milan, Italy, 10–15 January 2021; pp. 8523–8530.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; Wojna, Z. Rethinking the Inception Architecture for Computer Vision. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-16), Las Vegas, NV, USA, 27–30 June 2016; pp. 2818–2826.
- Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS-15), Montreal, QC, Canada, 7–12 December 2015; Volume 28, pp. 91–99.
- Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet Classification with Deep Convolutional Neural Networks. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS-12), Stateline, NV, USA, 3–8 December 2012; pp. 1097–1105.
- 25. Hochreiter, S.; Schmidhuber, J. Long Short-term Memory. Neural Comput. 1997, 9, 1735–1780. [CrossRef] [PubMed]
- 26. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An Image is Worth 16 × 16 Words: Transformers for Image Recognition at Scale. In Proceedings of the International Conference on Learning Representations (ICLR-21), Virtual Event, 3–7 May 2021.
- Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR-09), Miami, FL, USA, 20–25 June 2009; pp. 248–255.
- He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-16), Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
- Wei, S.E.; Ramakrishna, V.; Kanade, T.; Sheikh, Y. Convolutional Pose Machines. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 4724–4732.
- Howard, A.; Sandler, M.; Chu, G.; Chen, L.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. Searching for MobileNetV3. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV-19), Seoul, Korea, 27 October–2 November 2019; pp. 1314–1324.
- Lin, T.Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; Belongie, S. Feature Pyramid Networks for Object Detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-17), Honolulu, HI, USA, 21–26 July 2017; pp. 2117–2125.
- 32. Tan, M.; Le, Q. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In Proceedings of the International Conference on Machine Learning (ICML-19), PMLR, Long Beach, CA, USA, 9–15 June 2019; pp. 6105–6114.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; Jégou, H. Training data-efficient image transformers & distillation through attention. In Proceedings of the International Conference on Machine Learning (ICML-21), PMLR, Virtual Event, 18–24 July 2021; pp. 10347–10357.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention Is All You Need. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 5998–6008.
- Tran, D.; Wang, H.; Torresani, L.; Ray, J.; LeCun, Y.; Paluri, M. A Closer Look at Spatiotemporal Convolutions for Action Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-18), Salt Lake City, UT, USA, 18–23 June 2018; pp. 6450–6459.
- Zivkovic, Z. Improved adaptive Gaussian mixture model for background subtraction. In Proceedings of the 17th International Conference on Pattern Recognition (ICPR 2004), Cambridge, UK, 26 August 2004; Volume 2, pp. 28–31.
- 37. Zivkovic, Z.; Van Der Heijden, F. Efficient adaptive density estimation per image pixel for the task of background subtraction. *Pattern Recognit. Lett.* **2006**, *27*, 773–780. [CrossRef]
- He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV-17), Venice, Italy, 22–29 October 2017; pp. 2961–2969.
- Barnich, O.; Van Droogenbroeck, M. ViBe: A Universal Background Subtraction Algorithm for Video Sequences. *IEEE Trans. Image Process.* 2010, 20, 1709–1724. [CrossRef]

- 40. Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; Joulin, A. Emerging Properties in Self-Supervised Vision Transformers. *arXiv* 2021, arXiv:2104.14294.
- Pan, H.; Badawi, D.; Cetin, A.E. Fourier Domain Pruning of MobileNet-V2 with Application to Video Based Wildfire Detection. In Proceedings of the 2020 25th International Conference on Pattern Recognition (ICPR-20), Milan, Italy, 10–15 January 2021; pp. 1015–1022.
- 42. Hinton, G.; Vinyals, O.; Dean, J. Distilling the knowledge in a neural network. *arXiv* 2015, arXiv:1503.02531.