Intelligent Helipad Detection from Satellite Imagery

David Specht
Rowan University
Glassboro, N.J. , USA

Ghulam Rasool
Rowan University
Glassboro, NJ, USA

Nidhal Bouaynaya
Rowan University
Glassboro, NJ, USA

Cliff Johnson
Federal Aviation Administration
Egg Harbor, NJ, USA

ABSTRACT

Location data about U.S. heliports is often inaccurate or nonexistent in the FAA's databases, which leaves pilots and air ambulance operators with inaccurate information about where to find safe landing zones. In the 2018 FAA Reauthorization Act, Congress required the FAA to collect better information from the helicopter industry under part 157, which covers the construction, alteration, activation and deactivation of airports and heliports. At the same time, there is no requirement to report private helipads to the FAA when constructed or removed, and some public heliports do not have up to date records. This paper proposes an autonomous system that can authenticate the coordinates in the FAA master database, as well as search for helipads in a designated large area. The proposed system is based on a convolutional neural network model that learns optimal helipad features from the data. We used the FAA's 5010 database and others to construct a benchmark database of rotocraft landing sites. The database consists of $9,324$ aerial images, containing helipads, helistops, helidecks, and helicopter runways in rural and urban areas, as well as negative examples, such as rooftop buildings and fields. The dataset was then used to train various state-of-the-art convolutional neural network models. The outperforming model, EfficientNet-b0, achieved nearly 95% accuracy on the validation set.

INTRODUCTION

Background and Motivation

Accurate information about the location and type of rotorcraft landing sites is an essential asset for the Federal Aviation Administration (FAA) and the Department of Transportation (DOT). However, the acquisition, verification, and regular updating of information about these landing sites is a challenging task. The lack of reliable information on helipad sites is a risk factor in several accidents and incidents involving rotorcrafts. The U.S. Helicopter Safety Team (USHST), of which the FAA is a key member, has identified and produced recommendations from their infrastructure working group to modernize and improve "the collection, dissemination, and accuracy of heliport/helipad landing sites" as a high priority to increase helicopter safety.

There are thousands of landing locations for helicopters spread across the United States. In general, rotorcraft operators can get information about helipads, heliports, and landing sites using various databases, such as the FAA's 5010 database. However, it is also well-known that the 5010 database and similar databases contain multiple inaccuracies where some helipads in the database may no longer exist or their coordinates are imprecise, and other helipads are missing from the database. The unreliability of this database is a consequence of the fact that there is no system to verify that coordinates remain accurate, nor is there a system to search for unreported helipads.

In this paper, we propose a machine learning solution to identify helipads, heliports, and other landing sites, from aerial imagery, using convolution neural networks (CNNs) (Ref. 1). We built a comprehensive database by manually checking the FAA and other databases as well as augmenting them with satellite images from Google Earth. We subsequently trained and validated different state-of-the-art CNN models to determine an appropriate neural network model for this task.

The proposed deep learning solution will allow the FAA and USDOT to automatically maintain an updated database of helipads, heliports, and landing site infrastructure for the rotorcraft community. This work presents the first step towards autonomous identification of specialized heliport infrastructure and can be optimized with minimal cost using Google Earth API. The results of this project will help the FAA and USDOT achieve the first strategic goal of "Improving durability and extending the life of infrastructure" by providing an updated record of the infrastructure without committing additional resources for data collection and recording.

Related Work

We can group the literature of identifying helipads from satellite or aerial imagery into two main approaches. The first is a

Presented at the Vertical Flight Society's 77th Annual Forum & Technology Display, Palm Beach, Florida, USA, May 11–13, 2021.

model-based approach, which relies on domain expert knowledge to extract features that can be used to identify helipads from images. A common feature used to identify helipads is the "H" marking (Ref. 2), (Ref. 3). For vision-based autonomous landing systems, an improved version of the Scale Invariant Feature Transform (SIFT), called Speeded Up Robust Features (SURF), was used in (Ref. 2) to perform feature points matching and tracking. Features points are compared to points in an "H" template to determine the similarity of the template and the aerial image. In (Ref. 4), the detection process consists of finding candidate helipads based on the following four properties: 1) a bold circle surrounding the "H", 2) presence of "H" in a bright color inside this circle against a dark background, 3) "H" is centered at the center of the circle, and 4) intersection of diagonals of "H" at the center of the circle. A Hough transforms was used to identify circles (Ref. 4). Due to a large number of false positives, the authors used three tests to eliminate these false positives. None of these tests are precise and as a consequence, error ranges were added based on experiments. After a helipad has been detected, a Median Flow tracker (Ref. 5) was used to track the region.

A vision-based helipad detection algorithm based on curvature was proposed in (Ref. 3). The method creates blobs of connected pixels, and exploits some intrinsic properties of each blob, such as the location of its center of mass, the Euler number, the eccentricity, the perimeter, and the area, to identify the blobs which represent the helipad marks, namely the character "H" and the circumscribing circles. The Euler number is an integer value defined as the number of connected components minus the whole numbers. In particular, the Euler number is equal to zero for circle blobs and one for "H" blobs. A final classification level checks the ratios between the areas and perimeters of blobs against expected values. Following detection, an identification step checks if the Euclidean distance of the centroids of the detected blobs and the ratios of related areas and perimeters are still met. (Ref. 3). Once the helipad marks have been identified, the Canny edge detector is performed in order to extract the 12 corners of "H" edge. Instead of using feature extraction operators, such as the Hough transform and line following algorithms, the authors used a radius of curvature for every 2-D point of "H" edge to detect the corners of interest. A big radius value denotes that the point is far from a corner while a small value indicates that the point might be a possible candidate to be a corner. Three checks are performed for all the possible corner candidates, based on the knowledge of "H" size and exploiting the Euclidean distances between these points and the centroid of H contour.

Although quite exhaustive, These model-based detection algorithms have many restrictions. First, they were shown to work only in simple simulated environments and may fail in more complex environments. Secondly, these algorithms have limited effectiveness at further distances and angles. Some of these issues were addressed in (Ref. 6), where the authors mainly relied on flat ellipse detection as it is the most visible feature of a helipad seen from long distances. An adaption of the Hough transform was devised for the specific case of very flat ellipses. A validation step using many other properties and visual clues performs the verification of the presence of the helicopter landing platform in the research areas delimited by the obtained ellipses.

The main advantage of the model-based approach is its explainability and its relatively good performance on small datasets with no prior labeling. However while model-based methods can identify helipads that adhere to the recommended standard set in the FAA's 150/5390-2C (Ref. 7), neither the circle nor the "H" is a requirement for building a helipad. Model-based methods will need to consider all possible features of all types of helipads/heliports, including those that do not adhere to the recommended standard, to generalize their performance.

Data-driven algorithms, on the other hand, involve the collection of large amounts of labeled data, autonomously learning salient features from the raw data, and identifying helipads based on learned features. As such, data-driven systems can identify complex patterns of helipads that may be hard to model. The price paid is the large data and computational resource requirements. To the best of our knowledge, data-driven approaches to identify helipads are under-explored, despite the growing prevalence of learning systems in real-world applications. Nonetheless, there are online systems available. HelloPad (Ref. 8) is a system that uses a machine learning algorithm to identify helipads within a specified region. The system uses a sliding window and a trained neural network model (ResNet) to identify if a helipad exists at a given location. HelloPad reported 67.2% precision and 90% recall in a Los Angeles downtown area. However, HelloPad collected negative (non-helipad) examples from urban settings, and will likely not transfer well to all areas of the U.S.

CONVOLUTIONAL NEURAL NETWORKS

Learning Features with Convolutional Neural Networks

Object detection and identification requires considerable domain expertise to design features that transform the raw data (such as the pixel values of an image) into a lower-dimensional representation that is discriminatory for the input. Convolutional Neural Networks (CNNs) are designed to process multidimensional data arrays, such as images, by automatically discovering the representations needed for detection or classification. There are three types of layers in a CNN: convolutional layers, pooling layers, and fully connected layers. Each convolutional layer obtains, through convolutions followed by non-linear operators, representations that are important for the classification task. A hierarchical composition of these representations (starting with the raw input), where each representation is fed to the next convolutional layer, leads to learned features that are optimal for discrimination. The first (convolutional) layers typically learn low-level features, such as edges, and later layers extract more complex semantic features. The key aspect of CNNs is that these layers of features are not designed by human engineers or domain experts: they are learned from data (Ref. 9).

A problem with the output feature maps is that they are sensitive to the precise location of the features in the input. This

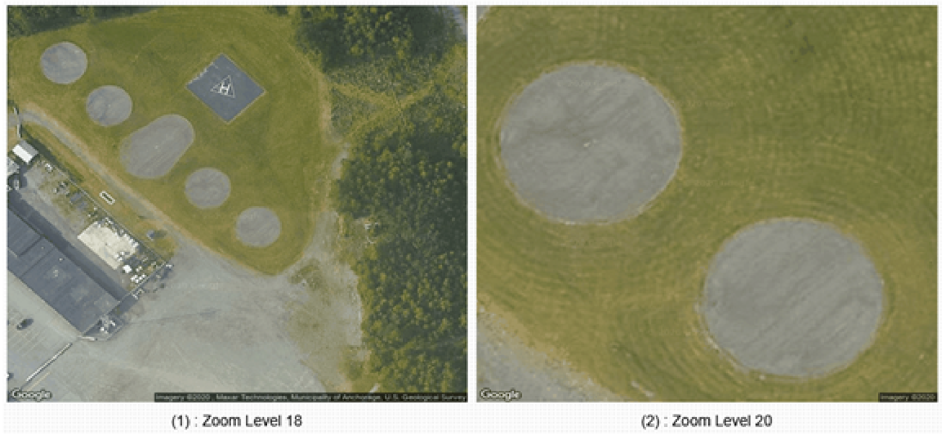

Figure 1. Zoom level in Google API. Left: zoom level of 18; Right: zoom level of 20. At a zoom of 18, numerous possible parking pads and a helipad are visible. The helipad is not visible at a zoom of 20.

means that small variations in the position of the feature in the input image will result in a different feature map. One approach to address this sensitivity is to coarse-grain the position of each feature through down-sampling, referred to as “local translation invariance”. The role of pooling layers is to summarize the feature maps by down-sampling, i.e., discarding the finer details that may not be useful to the task, creating an invariance to small shifts, while maintaining important structural elements. A typical pooling unit computes the maximum value for each patch of the feature map.

Layers of convolutions, non-linearities, and pooling are stacked to learn robust optimal features for the data, followed by fully-connected layers that form the classifier for the extracted features. Backpropagating gradients through a CNN is as simple as through a regular neural network, allowing all the weights in all the filters to be trained.

How Convolutional Neural Networks perceive the World

While CNNs have achieved higher-than-human accuracy in many computer vision tasks, they provide little insight into their decision-making process (Ref. 10). With the composition of convolutions, non-linearires, pooling and fully-connected layers, very complex functions can be learned, making deep learning models black boxes. This poor interpretability significantly hinders the robustness evaluation of the network, its further optimization, as well as understanding the network adaptability and transferability to different datasets. In the case of helipad detection, this question becomes “Does the network detect salient features of helipads in the image, or does it detect other features that typically correlate with the presence of a helipad?”. An understanding of the learning process will allow for the identification of cases where the algorithm might fail, and also build trust in learning systems to allow for their safe deployment.

An intuitive approach to understand the inner workings of deep learning models (such as CNNs) is the *gradient saliency map*. This approach computes the gradient of the class score with respect to the input image; thus, highlighting the areas of the input image that are discriminative with respect to the predicted class (Ref. 11). A popular gradient saliency method is the Gradient-weighted Class Activation Mapping (Grad-CAM). Grad-CAM uses the gradient information flowing into the last convolutional layer of the CNN to assign importance values to each neuron for a particular decision of interest (Ref. 12).

In this project, Grad-CAM provides a multifaceted advantage. First, the saliency map will be able to verify that the network classifies imagery as helipads because of the presence of helipads and not supporting facilities. Second, it can help with understanding and mitigating false positives, i.e., the non-helipad samples classified as a helipad. Lastly, the saliency map can help locate the helipad, which will allow for larger regions to be searched for helipads.

ROTORCRAFT LANDING DATASET

We acquired three datasets through the FAA, one dataset from the IOWA DOT website, and another dataset on arcGIS. These 5 datasets provide longitude and latitude of potential helipad landing locations. We used Google Earth’s API to extract the corresponding images as well as to sample negative helipad locations. We noticed some discrepancy in the FAA datasets and had to manually curate the coordinates to ensure accuracy for our use cases. In the sequel of this section we will elaborate on each dataset, our curation approach, and our collection method for negative samples.

Google static maps API

We used Google static maps API to collect, satellite imagery of positive (helipad) and negative (non-helipad) locations. The service is accessed by sending an *HTTP* request with a query containing the desired parameters. The Google server responds with an image based on the provided parameters. The parameters used here are: center, zoom, size, and maptype. The center provides the coordinates of the center of the

Figure 2. Sampling Google Earth for helipad labeling. The sampled area is enclosed in the black box. There is no helipad inside the sampled area. However, a helipad is present just outside the sampled area.

image. Zoom determines the zoom level, which defines the resolution of the current view. Size determines the number of pixels in the image. Maptype determines the type of image to be retrieved (as Google maps contains road maps). For the purposes of this project, size was set to the maximum value of 640×640, and the maptype was always satellite. The center was set to the desired coordinates to be sampled for the image. Zoom was set to 18. The highest resolution images are at a zoom of 20; however, a zoom of 18 was used instead. The difference between the two zooms can be seen in Fig. 1. A lower zoom results in a larger area that will allow for sampling helipads using fewer API calls. There is a cost associated with making API calls beyond a certain limit, so efficiency of calls becomes important when scaling up.

One major issue with the initial databases is the incorrect reporting of landing site coordinates. In our experiments, a coordinate was considered correct if there was a landing area present in the Google imagery taken of the area. This is to allow for a margin of error in the reported coordinates. The margin is considered acceptable as it is believed to be reasonable for a pilot to identify a helipad within the given area. However, there are few cases where helipads would be within a reasonable range of the coordinates, yet not present within the imagery being sampled. Figure 2, shows a case where there is a helipad near the coordinates, however the helipad is outside the range that was annotated. A lower zoom could be used to sample a larger area, however while this may still be within an acceptable margin of error, the markings on helipads become less noticeable.

Another known issue is the recency of Google's satellite imagery. The images used in Google maps are not real-time images, but rather imagery taken during an area survey. This means that the overhead view that was sampled does not actually reflect the current state of the area. Google attempts to keep the images up to date such that the available imagery should be less than 3 years old; yet this may still lead to inaccuracies in landing site locations

Lastly, there is an issue of some missing imagery. Google Maps does maintain a database that covers most of the world, however it does not contain high resolution imagery for every coordinate in the world. Typically, at higher levels of zoom, there are fewer coordinates with available imagery. Even when using a zoom of 18, there are a few coordinates that simply did not have imagery available. If a zoom of 20 were used, there would likely be fewer locations where data could be collected from.

Building the dataset

Positive set Areas with helipads are needed to create the positive set for our database. While areas can be randomly sampled and helipads in those areas labeled, this would be an incredibly inefficient process. There is an extremely low probability that a randomly sampled location would contain a helipad. We used the initial FAA, IOWA DOT, and ArcGIS helipad datasets to sample positive areas, The FAA's 5010 was the largest database. To ensure accuracy, all coordinates were manually annotated so that only coordinates where a helipad would be visible in the collected image was added to the training set. From an initial $6,333$ coordinates in the dataset, only $3,887$ were manually annotated to be helipads. An additional 157 positive coordinates were added from other databases provided by the FAA, including the Lifeflight of Maine dataset

Two publicly available datasets were used. The first is a dataset found on ArcGIS containing the coordinates of hospital helipads found in California. This dataset contained 170 coordinates, and after annotation, 169 of these coordinates were used. The second is Iowa DOT's dataset, which listed 126 locations, and 111 of these coordinates were considered to contain helipads.

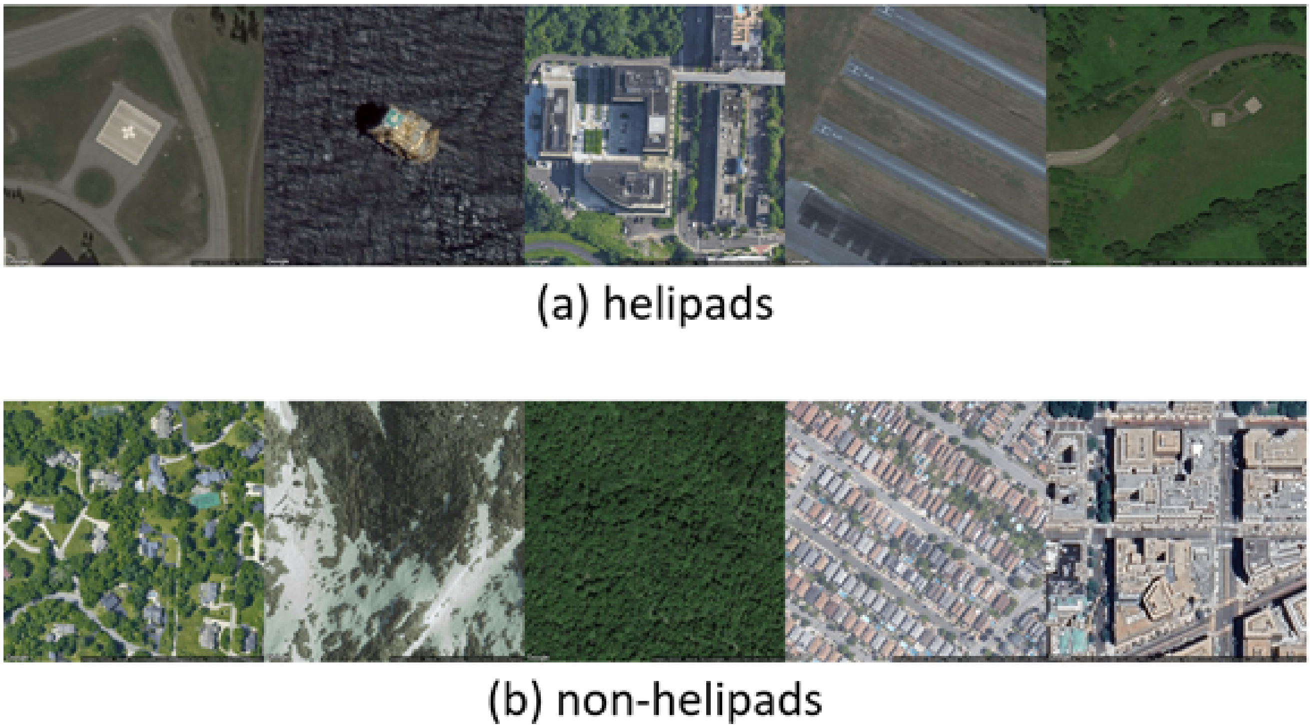

Figure 3. Sample aerial images from our Benchmark dataset. While the term helipad is used for the positive set, the positive set also contains areas that helicopters are intended to land at, e.g., helicopter runways

Negative set A negative (non-helipad) set of images is also needed to train the machine learning model. The negative set was collected using random sampling of Google Maps. These random samples were manually checked to not contain any landing site. As the current goal is to identify helipads in the U.S., the sampling was limited to an area such that the sampling region includes most of the mainland U.S. However, most of these samples were of forested areas and farmlands and contained very few urban areas. This could bias the network to predict helipads mainly in urban areas It is therefore important to sample negative locations from urban areas as well. It is noted that urban areas will likely have a higher helipad density, and thus a helipad will be more likely to be found there. To lessen this risk, locations like Washington D.C. and New York City were chosen due to the lower density of helipads. In New York City, ownership of rooftop helipads became more restricted after the 1977 crash at the Pan Am building (Ref. 13), along with noise complaints continuing to restrict helicopter flights. Washington D.C. is in restricted airspace and allows only a few helipads to operate.

Final Benchmark Dataset

After careful and meticulous collection, labeling, and curation steps, a helipad identification benchmark dataset was created. The positive set contains $4,324$ samples. Some areas are more represented than others, as some of the datasets used were specific to certain regions. However, the largest dataset making up over 80% of the final dataset is the FAA's dataset spread over the United States and its territories covering different types of landing areas including helicopter parking pads, helidecks, EHLFs(Emergency Helicopter Landing Facilities), and heliports.

The negative set was created by randomly sampling 5,000 coordinates. 2,000 of these coordinates were from sampling the mainland United States and contains mostly woodland and other rural areas. The remaining 3,000 negative images came from urban areas, such as San Jose, Washington D.C., New York City, and San Antonio.

The final benchmark dataset has $9,324$ satellite images labeled as either helipad or non-helipad. Figure 3 shows some of the images in the dataset. Figure 3 (a) shows some landing locations, including helistops, helidecks, and helicopter runways. Figure 3 (b) shows some of the randomly sampled imagery, with the three on the left being samples from the rural areas, and the two on the right coming from urban areas. It is also noteworthy to mention the variety of landing sites shown in Fig. 3(a). In particular helipads have different sizes, as their minimum required lengths are decided by the rotor diameter of helicopters intended to land. This causes the areas they represent in squared meters to be different. Other factors, such as the zoom level which takes into account the distance from the satellite, the elevation, and the latitude add to the complexity of the landing sites imagery.

EXPLAINABLE IDENTIFICATION OF HELIPADS

Convolutional Neural Network Models

Four different CNN models were evaluated for this classification task: ResNet101 (Ref. 14), Inception (Ref. 15), Xception (Ref. 16), and efficienetnet-b0 (Ref. 17). These models where chosen as they represent a variety of possible architecture families (see Table 1)

ResNet101 ResNet101, proposed in (Ref. 14), is a residual network that contains skip connections. Residual networks were designed to mitigate the problem of *accuracy degradation with network depth*. With the network depth increasing, accuracy gets saturated and then degrades rapidly. Skip connections simply perform identity mapping, and their outputs are added to the outputs of the stacked layers. It was experimentally shown that deep residual networks: 1) exhibit lower training error when the depth increases compared to their counterpart "plain" networks, and 2) enjoy accuracy gains from considerably increased depth, producing significantly better results than their counterpart "plain" networks.

Inception-V3 Inception-V3, proposed in (Ref. 15), is part the family of *Inception* networks. This family of networks uses the Inception module, which leverages multiple types of filter sizes in a convolutional layer instead of being restricted to a single filter size. The motivation of Inception stems from the human visual cortex, which identifies patterns at different scales and then accumulates them to form larger perceptions of objects. Therefore, Inception modules have the potential to improve optimal feature extraction, and hence improve the learning.

Xception Xception, proposed in (Ref. 16), improves upon the Inception family of architectures by replacing Inception modules with depthwise separable convolutions. A depthwise separable convolution is a spatial convolution performed independently over each channel of an input, followed by a pointwise convolution, i.e., a 1×1 convolution, projecting the channels output onto a new channel space. Xception is build by stacking depthwise separable convolutions. This model also uses skip connections.

EfficientNet-b0 The EfficientNet family of architectures was proposed in (Ref. 17) to address the issue of scaling CNN models for better accuracy. Based on a compound scaling method that balances network width, depth, and resolution, eight different models (b0 - b7) were proposed, where higher values correspond to larger networks. The models are made up of sections of repeating layers which can be efficiently scaled to create a deeper model.

EXPERIMENTAL RESULTS

We performed 10-fold validation for each network. The results from each of the 10 runs were averaged to produce the average performance of the model as shown in Fig. 4. As can be seen from Fig. 4, EfficientNet-b0 performed the best on our final benchmark helipad dataset.

Figure 5 shows the results from the grad-CAM implementation. This overlay shows a heatmap of the most salient pixels in the model's prediction (EfficientNet-b0). Observe that the network relied on features of the helipad to make its prediction, as opposed to background features, such as nearby buildings.

HELIPAD SEARCH IN LARGE AREAS

Using the CNN model validated in the previous section, it is now possible to identify imagery with helipads, which can be used to verify the accuracy of coordinates in helipad databases. This system can then be extended to be able to detect helipads within a designated region. In computer vision, the distinction between identification and detection is that identification can determine the presence of an object, while detection determines where in the image an object is. In this section, we extend the problem of helipad identification from aerial images to detection of helipads from a larger area, e.g., downtown Los Angeles. To solve this new problem, without requiring new labeling, we use a sliding window approach to determine where in a larger image a helipad is.

Collages

Sampling a larger Google Earth area can be done by using a lower value zoom, dividing it into sections, then upsampling the images. This approach would minimize the number of API calls; However, the images retrieved will be of lower resolution. The second approach would be to sample using a higher zoom for higher resolution imagery, then combine the samples to form a larger image referred to as a *collage*. This collage can then be searched for helipads with an overlapping sliding window. A mapping between the lat/lon coordinates and pixel values must be derived. Google has provided the following relationship:

$$\frac{\text{meter}}{\text{pixel}} = 156543.03392 \times \frac{\cos(\text{latitude} \times \frac{\pi}{180})}{2^{\text{zoom}}} \tag{1}$$

The distance represented by a pixel decreases as we sample further from the equator. Equation (1) does not factor in elevation, and may cause issues at different elevations.

Assuming that the circumference of earth is 40.075 million meters and taking elevation into account, we can derive the following mapping from pixels to change in lat/lon.

$$\frac{\Delta\ \text{latitude}}{\text{pixel}} = 156543.03392 \times \frac{cos(\text{latitude} \times \frac{\pi}{180})}{2^{\text{zoom}} \times 111320} \tag{2}$$

$$\frac{\Delta\ \text{longitude}}{\text{pixel}} = 156543.03392 \times \frac{1}{2^{\text{zoom}} \times 111320} \tag{3}$$

An example of the created collage can be seen in Fig. 6. This collage is created from a 5×5 sliding window, and shows an area about 25 times larger than the initial aerial images, while still keeping the level of detail in a higher zoom. Sub-images can then be extracted from this area to search for helipads.

Table 1. Comparison of selected CNN models

Model	Skip connections	Inception Modules	Trainable Parameters (millions)	Image-net Top 5-accuracy
ResNet101	yes	no	44.71	92.8%
Inception-V3	no	yes	23.85	93.7%
Xception	yes	no	22.91	94.5%
EfficientNet-b0	yes	no	5.33	97.1%

Figure 4. Training (blue) and validation (orange) accuracy curves for ResNet101 (top left), Inception-V3 (top right), Xception (bottom left), and EfficientNet-b0 (bottom right).

Searching for Helipads in Los Angeles

We applied this collage technique to create a Los Angeles (LA) region. LA makes for an interesting testing area, as it has a high helipad density, so there will be many helipads to detect in a small region. However, the LA region is notably different than the other cityscapes in the dataset. To fix this data imbalance, some of the data from LA was used to supplement the benchmark dataset. Figure 7 shows the LA area under study. It was formed via an 11×11 collage, and will be broken up using a 20×20 sliding window producing 400 smaller images. The 200 images making up the top half of the image will be part of the supplemental training set, and the 200 images in the bottom half of the image will make up the testing case. The test achieved the following performance measures: Accuracy = 76.0%, Precision = 61.6%, and Recall = 97.4%.

CONCLUSION

We developed a deep learning model for helipad identification and detection from aerial Google Earth imagery. We also devised a framework to begin searching for helipads in designated areas. We achieved good performance in detecting helipads in the LA region; Nonetheless, more experimentation is needed before this approach is ready for more widespread testing. Notably a larger variety of data should be considered, and a wide variety of locations should be incorporated for testing to ensure that the algorithm will perform in these different locations.

Although the study was limited to the US, this approach is readily extendable to helipad identification and detection across the globe. Future work includes leveraging Grad-CAM interpretability maps to estimate the location of the helipads after their identification.

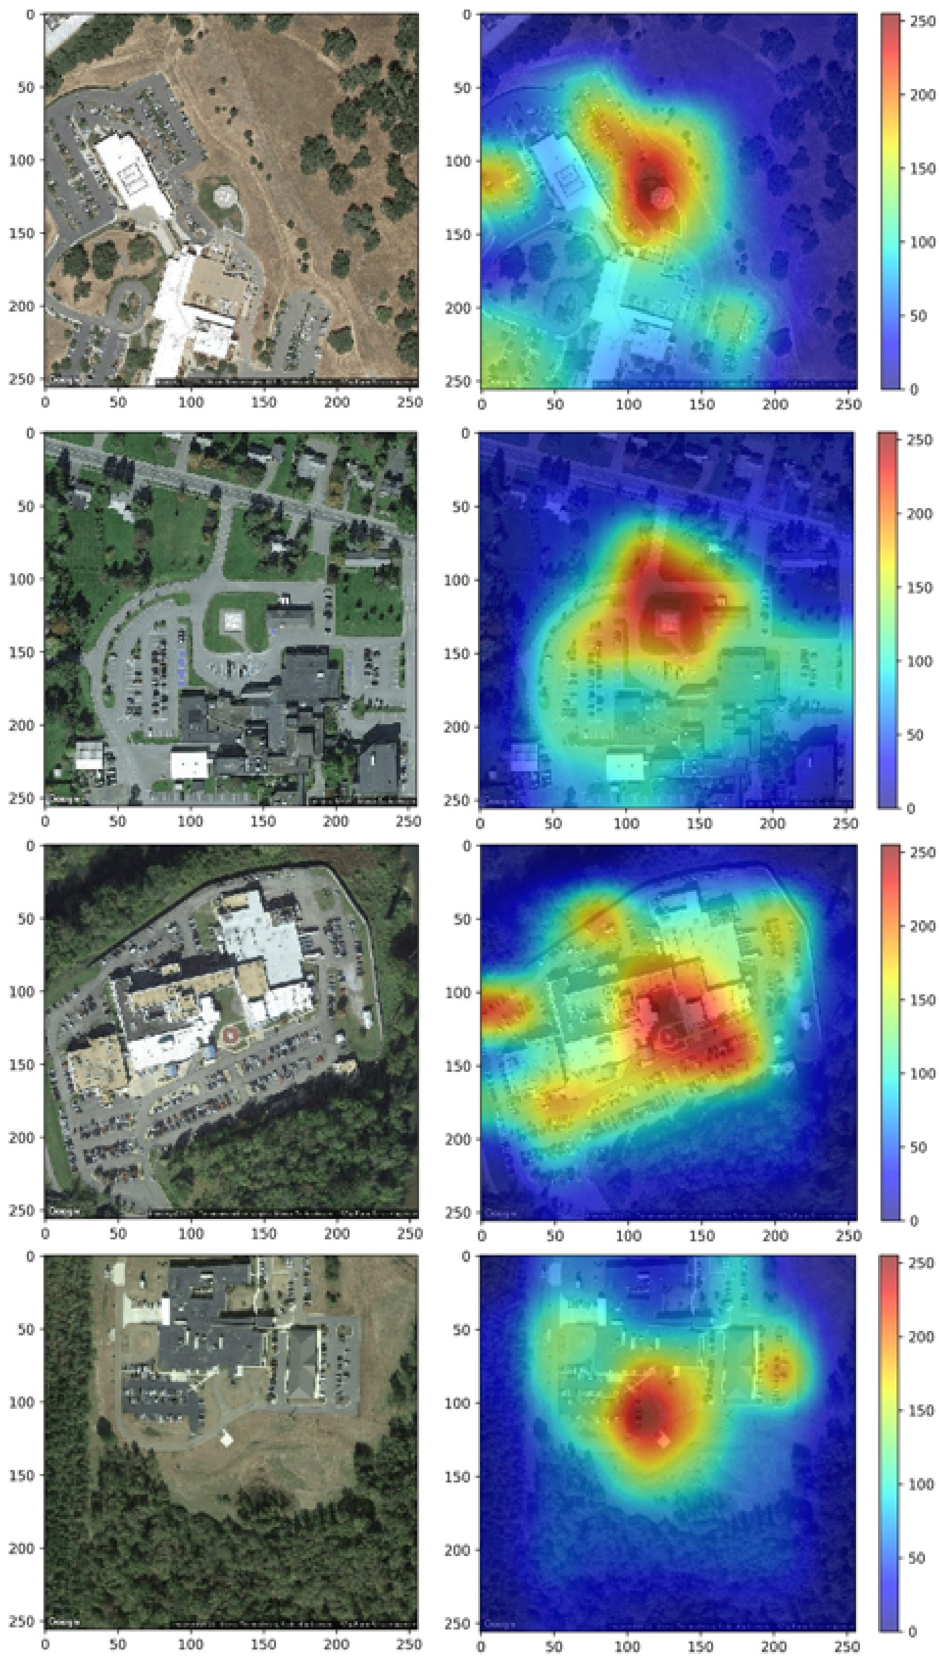

Figure 5. Grad-CAM heatmaps showing the importance of pixels in the CNN model (EfficientNet-b0) prediction. Aerial images (left column) and their corresponding Grad-CAM heatmaps (right column). The red area refers to the part of the model where the network attention is strong, and the blue part refers to the part that does not influence the prediction.

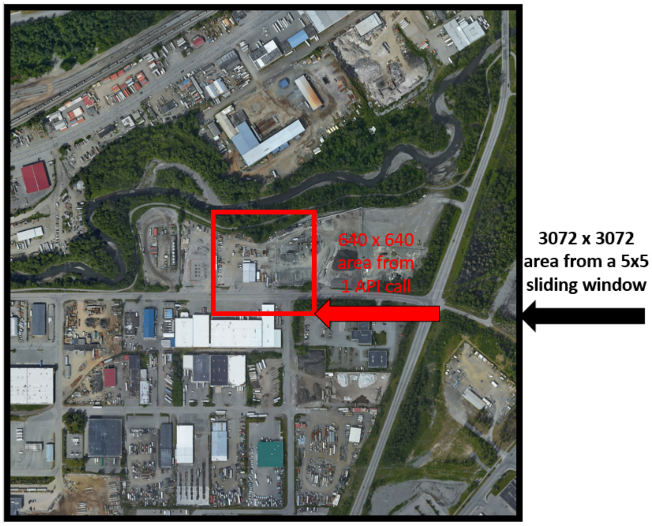

Figure 6. Area collected in a 5×5 collage vs. single API call.

Figure 7. Los Angeles (LA) region to be sampled. As the network was not trained on a similar cityscape, the top half of the area was used to supplement the training dataset and the model was tested on the lower half.

ACKNOWLEDGMENTS

This work was supported by the Federal Aviation Administration (FAA) Cooperative Agreement Number 16- G-015 and NSF Award DUE-1610911. This publication was also supported by a subaward from Rutgers University, Center for Advanced Infrastructure & Transportation, under Grant no. 69A3551847102 from the U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology (OST-R). We would also like to thank LZControl for their guidance and assistance with this effort. Via a Cooperative Research and Development Agreement with the FAA, LZControl provided a set of data from their system and subject matter expertise which provides landing zones for helicopters across the U.S., often complementing the FAA's 5010 database and including locations/sites not present in the FAA's 5010 system..

REFERENCES

1. Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, Vol. 60, (6), May 2017.

2. Prakash, R. O., and Saravanan, C., "Autonomous Robust Helipad Detection Algorithm using Computer Vision," International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT), 2016. DOI: 10.1109/ICEEOT.2016.7755163

3. Patruno, C., Nitti, M., Stella, E., and D'Orazio, T., "Helipad detection for accurate UAV pose estimation by means of a visual sensor," *International Journal of Advanced Robotic Systems*, Vol. 14, (5), 2017. DOI: 10.1177/1729881417731083

4. Rungta, A., Soni, Y., Agarwal, P., Ghosh, B., and Kumar, S., "Real-time and Autonomous Detection of Helipad for Landing Quad-Rotors by Visual Servoing," , 2020.

5. Zdenek Kalal, K. M., and Matas, J., "Forward-Backward Error: Automatic Detection of Tracking Failures," International Conference on Pattern Recognition, 2010. DOI: 10.1109/ICPR.2010.675

6. Pierre, Z., Mavromatis, S., Sequeira, J., Anoufa, G., Bélanger, N., and Fillias, F.-X., "Embedding Intelligent Image Processing Algorithms: the new safety enhancer for helicopters missions," , 2018.

7. "150/5390-2c - heliport design," , Apr 2012.

8. Walker, G., "No one knows where America's helipads are, except this neural network," .

9. LeCun, Y., Bengio, Y., and Hinton, G., "Deep learning," *Nature*, Vol. 521, May 2015, pp. 436–444.

10. Gunning, D., and Aha, D., "DARPA's Explainable Artificial Intelligence (XAI) Program," Vol. 40, June 2019. DOI: 10.1609/aimag.v40i2.2850

11. Simonyan, K., Vedaldi, A., and Zisserman, A., "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps," *arXiv*, Vol. 1312.6034, 2014.

(a) LA Test Region (b) Ground Truth (Helipads) (c) Model Prediction (Helipads)

Figure 8. Large Area Helipad Search with the Sliding Window Approach. (a) The test region consists of the lower half of the LA area in Fig. 7. (b) Ground Truth helipad locations. (c) Model prediction of helipad locations.

12. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," *International Journal of Computer Vision*, Vol. 128, (2), Oct 2019, pp. 336–359. DOI: 10.1007/s11263-019-01228-7

13. Grinberg, E., "Deadly helicopter crash renews calls to ban flights over Manhattan," *CNN*.

14. He, K., Zhang, X., Ren, S., and Sun, J., "Deep residual learning for image recognition," Proceedings of the IEEE conference on computer vision and pattern recognition, 2016. DOI: 10.1109/CVPR.2016.90

15. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z., "Rethinking the inception architecture for computer vision," Proceedings of the IEEE conference on computer vision and pattern recognition, 2016.

16. Chollet, F., "Xception: Deep Learning with Depthwise Separable Convolutions," 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017. DOI: 10.1109/CVPR.2017.195

17. Tan, M., and Le, Q., "Efficientnet: Rethinking model scaling for convolutional neural networks," International Conference on Machine Learning, 2019.

AUTHORS BIOGRAPHY

David Specht is currently an MS. student with the Department of Electrical and Computer Engineering at Rowan University. He is a research fellow supported by the Federal Aviation Administration (FAA) via a research grant/cooperative agreement. His research interests focus on machine learning learning theory and applications.

Nidhal C. Bouaynaya received her Ph.D. degree in Electrical and Computer Engineering and M.S. degree in Pure Mathematics from The University of Illinois at Chicago, in 2007. From 2007-2013, she was an Assistant then Associate Professor with the Department of Systems Engineering at the University of Arkansas at Little Rock. In Fall 2013, she joined the Department of Electrical and Computer Engineering at Rowan University, where she is currently a Professor and the Associate Dean for Research and Graduate Studies. Dr. Bouaynaya co-authored more than 100 refereed journal articles, book chapters and conference proceedings. She won numerous Best Paper Awards, the most recent was at the 2019 *IEEE International Workshop on Machine Learning for Signal Processing*. She is also the winner of the Top algorithm at the 2016 Multinomial Brain Tumor Segmentation Challenge (BRATS). Her research interests are in Big Data Analytics, Machine Learning, Artificial Intelligence and Mathematical Optimization. In 2017, she Co-founded and is Chief Executive Officer (CEO) of MRIMATH, LLC, a start-up company that uses artificial intelligence to improve patient oncology outcome and treatment response.

Ghulam Rasool is an Assistant Professor of Electrical and Computer Engineering at Rowan University. He received a BS. in Mechanical Engineering from the National University of Sciences and Technology (NUST), Pakistan, in 2000, an M.S. in Computer Engineering from the Center for Advanced Studies in Engineering (CASE), Pakistan, in 2010, and the Ph.D. in Systems Engineering from the University of Arkansas at Little Rock in 2014. He was a postdoctoral fellow with the Rehabilitation Institute of Chicago and Northwestern University from 2014 to 2016. He joined Rowan University as an adjunct professor and later as a lecturer in the year 2018. Currently, he is the co-director of the Rowan AI Lab. His current research focuses on machine learning, artificial intelligence, data analytics, signal, image, and video processing. His research is funded by National Science Foundation (NSF), U.S. Department of Education, U.S. Department of Transportation (through the University Transportation Center (UTC), Rutgers University), Federal Aviation Administration (FAA), New Jersey Health Foundation (NJHF), and Lockheed Martin, Inc. His recent work on Bayesian machine learning won the Best Student Award at the 2019 IEEE Machine Learning for Signal Processing Workshop.

Charles C. Johnson works as a research engineer, program manager, and technical expert for the Aviation Research Division at the FAA William J. Hughes Technical Center in Atlantic City, NJ. During his 10+ year-career with the FAA, he has led several rotorcraft and unmanned aircraft systems (UAS) research and development simulation/flight test activities that seek to improve aviation safety. Cliff is qualified on the ScanEagle UAS platform. He holds a Bachelor's degree in Mechanical Engineering from Rowan University. He also holds a private pilot's license (single engine land/fixed-wing) with instrument rating and is pursuing his instrument, commercial, and helicopter add-on ratings.