



Article

Identifying Critical Infrastructure in Imagery Data Using Explainable Convolutional Neural Networks

Shiloh N. Elliott ^{1,*}, Ashley J. B. Shields ¹, Elizabeth M. Klaehn ¹ and Iris Tien ²¹ Idaho National Laboratory, 1955 N Fremont Avenue, Idaho Falls, ID 83415, USA² Georgia Institute of Technology, 790 Atlantic Drive, Atlanta, GA 30332, USA

* Correspondence: shiloh.elliott@inl.gov

Abstract: To date, no method utilizing satellite imagery exists for detailing the locations and functions of critical infrastructure across the United States, making response to natural disasters and other events challenging due to complex infrastructural interdependencies. This paper presents a repeatable, transferable, and explainable method for critical infrastructure analysis and implementation of a robust model for critical infrastructure detection in satellite imagery. This model consists of a DenseNet-161 convolutional neural network, pretrained with the ImageNet database. The model was provided additional training with a custom dataset, containing nine infrastructure classes. The resultant analysis achieved an overall accuracy of 90%, with the highest accuracy for airports (97%), hydroelectric dams (96%), solar farms (94%), substations (91%), potable water tanks (93%), and hospitals (93%). Critical infrastructure types with relatively low accuracy are likely influenced by data commonality between similar infrastructure components for petroleum terminals (86%), water treatment plants (78%), and natural gas generation (78%). Local interpretable model-agnostic explanations (LIME) was integrated into the overall modeling pipeline to establish trust for users in critical infrastructure applications. The results demonstrate the effectiveness of a convolutional neural network approach for critical infrastructure identification, with higher than 90% accuracy in identifying six of the critical infrastructure facility types.

Keywords: remote sensing; critical infrastructure detection; convolutional neural networks; explainability; machine learning



Citation: Elliott, S.N.; Shields, A.J.B.; Klaehn, E.M.; Tien, I. Identifying Critical Infrastructure in Imagery Data Using Explainable Convolutional Neural Networks. *Remote Sens.* **2022**, *14*, 5331. <https://doi.org/10.3390/rs14215331>

Academic Editors: Jun Zhou, Danfeng Hong and Jing Yao

Received: 19 August 2022

Accepted: 17 October 2022

Published: 25 October 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/>).

1. Introduction

Critical infrastructure (CI) systems in the United States contain a diverse array of facilities, functions, and dependencies [1]. Failure of a facility in one sector can lead to cascading events that negatively impact CI across multiple sectors [2]. An example of this is infrastructure damaged during Hurricane Harvey, where a series of power failures led to a chemical plant explosion and chemical spills impacting the surrounding area [3]. The complexity of these systems, in addition to the siloed nature of CI in the United States, makes it difficult to identify and analyze these systems and their relationships using existing methods and information. The lack of a comprehensive understanding of the locations, types, and dependencies of CI across an area inhibits both the anticipation and response time of state and federal agencies when reacting to natural disasters or other events [4].

To date, no methods exist that enable the detailing of locations and functions of CI across the United States in a computationally efficient and repeatable manner. Instead, CI data exists in silos at the federal, state, municipal, and private levels [4]. Without understanding CI and their dependencies, it is difficult to provide accurate information in implementing risk mitigation measures and in emergency response. The ability to identify assets across multiple CI sectors is particularly important. Doing so enables identification of not only individual assets but also inferences to be made of functional relationships between assets in different sectors across both service provision and geographic infrastructural

interdependencies [5]. An additional challenge associated with the study of CI is the evolving nature of CI systems. Thus, it is important to facilitate the timely updating of information regarding the construction and decommissioning of different CI facilities.

In this paper, we present a new approach to address these challenges in CI analysis. The main contributions are to: (1) provide a novel repeatable machine-learning method for the cross-sector identification of multiple CI facilities in satellite imagery with a high degree of accuracy, an approach currently absent in the existing body of literature, and (2) provide explanations for the model's conclusions. To achieve this, we use a combination of unique data generation practices, a DenseNet161 convolutional neural network (CNN) architecture, and two explainability frameworks: local interpretable model-agnostic explanations (LIME) and Shapley additive explanations (SHAP). The method enables the model to be easily transferred for use with new data as updated aerial imagery becomes available.

The rest of the paper is organized as follows:

- Section 2 describes related work in this area and the advancements of this work compared to prior studies.
- Section 3 describes our methodology, including data generation and model selection.
- Section 4 presents our results and provides discussion on the accuracy of the outcomes. To provide further insights into the results, this section also describes our explainability analysis, including the explainability frameworks implemented and the subsequent analysis of the results.
- Finally, Section 5 provides conclusions on this work and describes future research directions based on the outcomes of this study.

2. Background and Related Work

The identification of objects in satellite imagery data has been an extensively studied field. In the domain of CI, there has been limited application of current machine-learning techniques to the identification process. Recent work within the related domain of damage detection in urban images focused on utilizing modified weakly supervised attention networks to detect destruction within an urban image [6]. Additionally, several works have identified components of CI, such as airports [7–9], ships [10–12], and roads [13]. These efforts include a range of approaches, including AlexNet [7] or VGG16 [14] architectures. The most common approach for these problems is using variations of R-CNN architectures [8–12]. However, these previous works have focused on detection of a single type of facility. Few, if any, efforts have included the domain diversity (i.e., range of CI sectors and facility types) provided in this work, and none have been developed with CI analysts as a target customer for the final results.

This work utilizes a DenseNet 161 architecture. Convolutional neural networks such as DenseNet161 will convert an image into a mathematical representation (matrix). From this point on, the derived matrix will go through a series of convolution, pooling, flattening layer, dropout layer, densely connected layer and activation functions that compose the bases of CNN architectures [15]. It is through this process that the model can “learn” the common features of a class of images (e.g., airports). When classifying an image, a trained CNN will utilize the feature learning obtained from the training process to classify an unknown image. Specific architectural features vary depending on the architecture in use. For example, the DenseNet architectures contain DenseBlocks, sections of architecture that are fully connected to other layers within the block. These blocks reduce accuracy decline, which is attributed to distance between input and output layers [16]. To date, no significant work has been published regarding the use of explainable CNNs in the identification of specific CI facilities in diverse CI sectors. There has been previous work regarding the detection of infrastructure expansion [17] and infrastructure quality [18]. However, the majority of this work focuses on change detection at the country or city level. These previous works do not touch on the identification of individual facilities. In the prior work on infrastructure expansion, infrastructure quality, or work at the facility

scale focusing on detection of only single facility types, none of this work has incorporated elements of explainability for model conclusions in this domain.

This paper describes the creation of a new approach that combines machine-learning methods with subject matter expertise to result in a domain aware CI dependency analysis tool. Ultimately this paper aims to introduce the use of a deep learning technique in the identification of CI and to establish baseline performance metrics for lifeline CI sectors using a DenseNet161 CNN architecture. The result of this work is a repeatable, transferable, and explainable method for CI detection. It is applied to nine different CI asset types, highlighting the need and ability to identify and distinguish between facilities in different CI sectors. Results have applications in a range of end-use scenarios for CI including emergency response, dependency analysis, and identification of vulnerabilities that can be bolstered to ensure the safety, security, and resilience of CI systems.

3. Materials & Methods

This section details the data generation and model selection processes to build and train the proposed machine-learning model for CI identification. The detailed description of the workflow in its entirety is presented in the following text, but to provide an overview here, a schematic of the modeling and data analysis pipeline is shown in Figure 1. The figure describes the flow of information from image data to final explainable predictions. This process includes a train/test split, development of a DenseNet161 model, implementation of the model on unknown data, explainable assessments via LIME—described in more detail in Section 4.4—and a final output including the top three most probable predictions, along with an image communicating which superpixels in the image were the most influential in the classification’s top prediction.

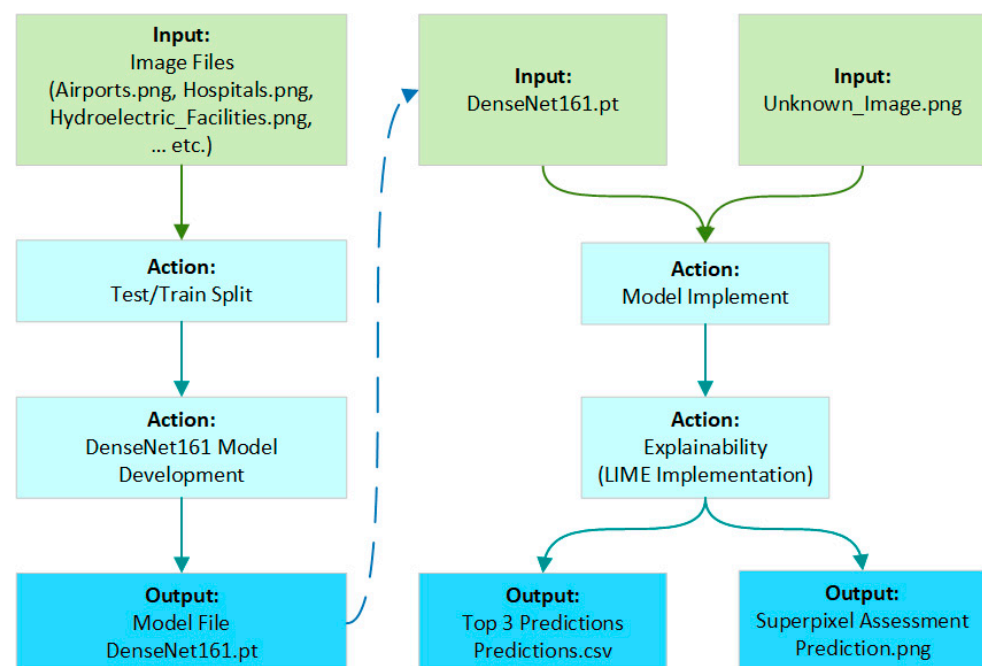


Figure 1. Modeling and analysis pipeline showing the flow of information from image data inputs to explainable predictions, with example input and output data files.

3.1. Data Generation

For our analysis, we identified nine CI facilities within five CI sectors for study. These are shown in Table 1. To qualify for consideration in this work, a facility had to be considered critical to sector operations, possess a facility footprint detectable from a standard RGB satellite image, and contribute to demonstrating the model’s ability to correctly identify heterogeneous facilities across multiple sectors. These nine facility types

were selected to represent a range of CI functions. They represent facilities critical to the energy, water, transportation, healthcare, and chemical sectors. Three of these sectors (energy, water, and transportation) are designated by the U.S. Department of Homeland Security as lifeline sectors. Loss of a lifeline sector facility will have a direct impact on the resilience of the affected facility and any interdependent facilities [19].

Table 1. Selected CI sectors and facilities for identification using proposed machine-learning model.

Sector	Facility
Energy	• Hydroelectric Dams
	• Natural Gas Generation Plants
	• Solar Farms
	• Substations
Water	• Potable Water Tanks
	• Water Treatment Plants
Transportation	• Airports
Healthcare	• Hospitals
Chemical	• Petrol Terminals

As inputs for the model training, facility locations were obtained from Idaho National Laboratory’s All Hazard Analysis (AHA) database. AHA is a methodology and application to collect, store, and model function, commodity, and service flows of interconnected systems to facilitate scalable and repeatable assessments of system behaviors suitable for vulnerability, consequence, and risk analysis [20]. Facility locations were then overlaid with the U.S. Department of Agriculture’s National Agriculture Imagery Program’s (NAIP) most recent data layer. The NAIP data set was selected for four fundamental reasons: (1) a 3-year refresh rate of the data set, (2) a resolution ranging between 2 m and 0.5 m, (3) an average of 10% or less cloud cover in gathered images, and (4) coverage of the contiguous United States [21]. Single facility images were extracted from NAIP using a combination of manual and automated techniques. An example of NAIP imagery data with AHA overlay is shown in Figure 2. The example shows an airport, as identified and indicated by the green dot on the image. Total number of datapoints for each facility ranged from 479 to 2292 unique images as shown in Table 2. Data was then randomly sorted along a 20% testing and 80% training split.

Table 2. Total images for each CI facility type.

Facilities	Total Images
Airports	1000
Potable Water Tanks	514
Hospitals	999
Hydroelectric Dams	637
Natural Gas Generations Plants	1538
Petrol Terminals	2292
Solar Farms	500
Substations	479
Water Treatment Plants	634



Figure 2. NAIP imagery example with AHA overlay indicated with a green dot.

3.2. Model Selection

Numerous deep learning models and approaches are currently available to researchers working in the field of imagery classification. These approaches range from widely available and utilized architectures to advanced multimodal deep learning and cross-modality learning frameworks that allow for complex in-depth analysis of the imagery they classify [22]. This work utilized the former approach, given the encouraging results of the widely available architectures demonstrated throughout the training and testing process and the desire to make the work easily replicable to a wide audience. Additionally, we implemented deep learning rather than shallow methods (traditional machine learning) because of the complexity of the datasets. One challenge of automated CI analysis is that the facilities themselves are diverse as are the background (i.e., surrounding geography). During the exploratory phase of this project, the team implemented a range of network depths and network types and found that acceptable accuracy was not achievable with anything less than a deep and fully connected network.

The initial stage of model selection was exploration-based, where the DenseNet-201, DenseNet-161, ResNeXt-101, and Resnet-152 were implemented and assessed for performance. These assessments focused on accuracy, training loss, validation loss, and a qualitative assessment of LIME explanations. Based on preliminary findings, a DenseNet-161 architecture was implemented for the final model. The DenseNet architecture was developed by Huang et al. [16] and implements a densely connected CNN, where each node is fully connected to every other node in a series. Unlike residual styles, such as ResNet and ResNeXt, dense CNN blocks do not utilize skip connections. Instead, they are designed for efficiency by implementing shallow sub-networks separated by convolution and pooling layers that simplify the data. DenseNet's robust communication between nodes is computationally expensive but facilitates the assessment of complex data, such as the complex imagery data generated via remote sensing as used in this work. Based on the preliminary results and this architectural style that is less prone to overfitting and requires fewer parameters to develop an accurate model than alternative methods (Huang et al., 2017), this model architecture was therefore selected and used here.

4. Results and Analysis

This section describes the results of the model based on the DenseNet-161 architecture. Included are the overall accuracy and training loss results as well as accuracy results for

individual CI facilities by type from cross validation. Next, the explainability activities conducted during the model development process are described, along with implementation and outcomes of both LIME and SHAP explainability frameworks. Resultant dataset analysis based on the explainability outcomes are then discussed.

4.1. Model Accuracy

With the developed model, we evaluated the results by both accuracy and training loss as shown in Figure 3. Accuracy is measured by the proportion of correctly identified facilities when testing data is used in conjunction with a trained model; training loss is the summation of incorrect predictions that occurred during a training epoch. The lower the training loss, the more accurate the model should be. Results for both the training data (80% training set) and validation data (20% testing set) are shown. From Figure 3, the results show that after an initial training phase, the highest overall validation accuracy of 82% was achieved at Epoch 33. During model training, the model accuracy and training loss improved significantly between Epochs 1–8 and only improved slightly with additional training. The values converge near their best performance around Epoch 15, suggesting that the model is slightly underfit and additional improvements in performance likely will not be achieved by processing the data over additional epochs. Accuracy (Figure 3a) and loss (Figure 3b) patterns for Epochs 9–50 include a substantial amount of noise which is likely caused by the inherent complexity of remotely sensed imagery data [23]. This is caused by the diversity of the CI facilities themselves as well as the diverse background pixels relating to climate and landscape diversity throughout the United States.




Figure 3. Model accuracy by epoch (a) and validation loss by epoch (b).

The best accuracies for each class (i.e., each CI facility type) are ranked as follows: airports (97%), hydroelectric facilities (96%), solar farms (94%), hospitals (93%), potable water tanks (93%), substations (91%), petroleum terminals (86%), natural gas generation plants (78%), and water treatment plants (78%). These results are presented as a confusion matrix in Figure 4. The confusion matrix presents the predicted (horizontal axis) compared to true (vertical axis) data labels. The frequency of intersections between predicted values and true values are represented by a color bar, where the most frequent intersections (most accurate) are indicated by the darker blue and the least frequent (least accurate) are white. The results indicate airports, solar farms, hydroelectric facilities, substations, finished water tanks, and hospitals as achieving greater than 90% accuracy (i.e., where more than 90% of the true and predicted labels for that class are the same).

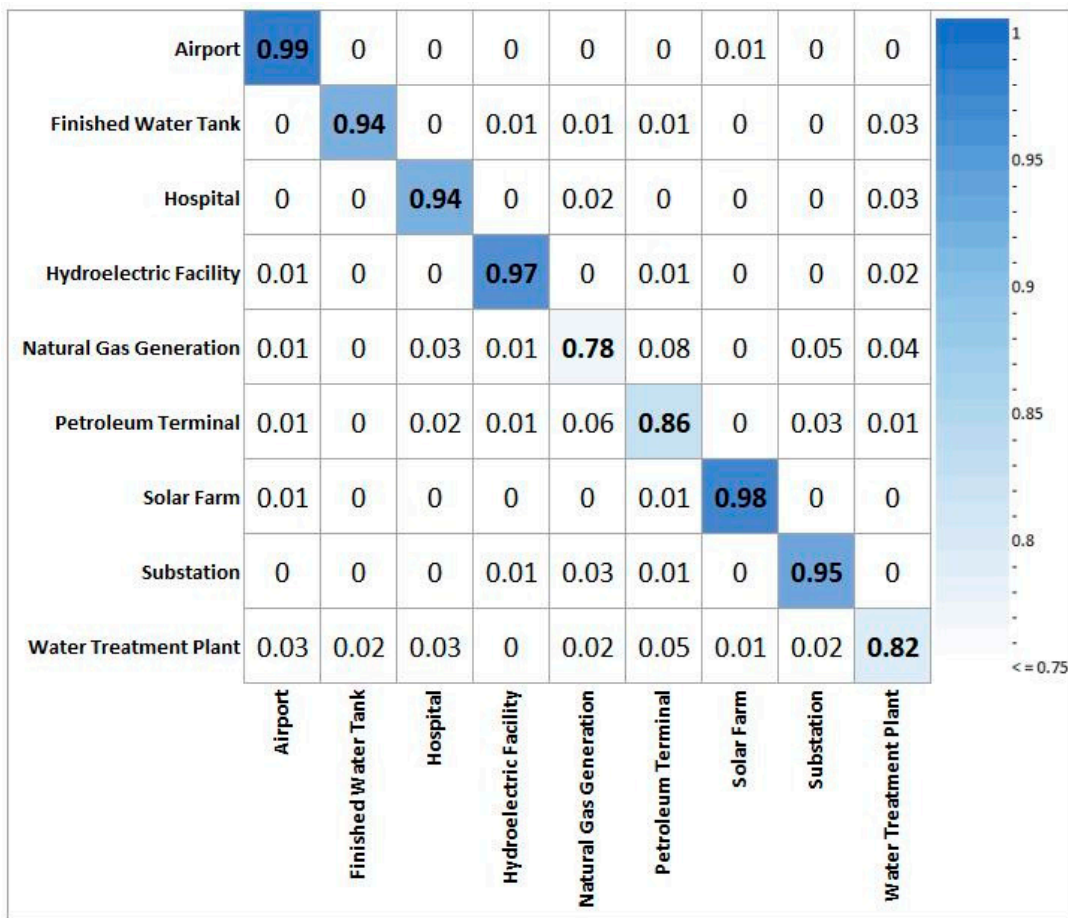


Figure 4. Confusion matrix for the nine CI facility types.

4.2. Cross-Validation

To further evaluate model accuracy, we conducted a k-fold cross-validation analysis [24]. Based on a series of test runs with k values ranging from 5–50, we established that of k of 10 folds was appropriate. During each iteration of k, data for each facility type was randomized k times and then split into training and testing data sets along an 80/20 split, respectively. The CNN model was then run k times. The overall cross-validation process is shown in Figure 5.

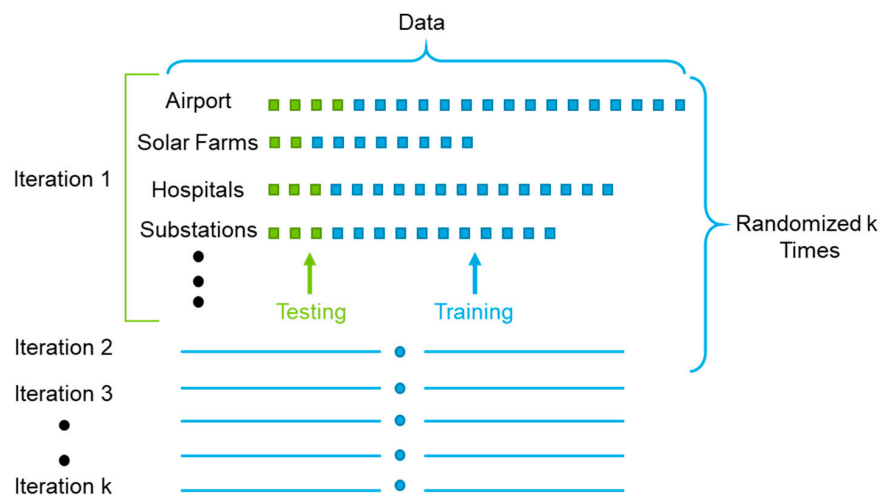


Figure 5. Visual depiction of the cross-validation process for CI dataset.

Table 3 presents the accuracy results averaged over all runs [25]. The similarity between the initial analysis results and the results from cross-validation indicate that the accuracy results are consistent, and the model results are unbiased relative to the data distribution. In addition, the high accuracy in cross-validation indicates the generalizability of the approach to new datasets, particularly for the identification of airports, hydroelectric dams, solar farms, potable water tanks, and substations facilities.

Table 3. Cross-validation average accuracy results.

Facilities	Average Accuracy (k = 10)
Airports	97
Hydroelectric Dams	96
Solar Farms	94
Hospitals	93
Potable Water Tanks	93
Substations	91
Petrol Terminals	86
Natural Gas Generation Plants	78
Water Treatment Plants	78
Overall Model Average	90

4.3. Explainability

While many machine-learning models operate as black boxes, in addition to the accuracy results presented in the previous sections, key to this work is our explainability analysis of the model outcomes. Explainability is a process that can assist in determining why a machine-learning model produced a certain output given a unique input, “explaining” how a trained model came to its conclusions. This provides a window into an otherwise black box process. Different machine-learning models and approaches utilize different implementations of explainability [26]. For our purposes, we utilize explainability to ensure the trained model is detecting the correct CI facilities for each class, guard against any unknown bias present in the training data set and provide a level of certainty in the model’s conclusions. Additionally, we integrate our explainability approaches into the overall modeling pipeline to establish a basis of trust for potential non-expert users to view and understand model classifications. This trust in the model’s conclusions is particularly important for CI applications, where asset and facility identifications have lifeline-critical implications, and where information is to be used by CI owners, operators, and emergency response personnel. The following two sections describe the analysis outcomes from implementing the LIME and SHAP explainability frameworks.

4.4. LIME Implementations

LIME is a model-agnostic approach utilized in the explanation of machine-learning classification models [27]. When applied to an image classification model, LIME begins its analysis by dividing an image into superpixels or defined regions within the given image. A linear regression model is then trained based on the probabilities of correct classifications produced by turning off and on various superpixels. The results of the linear regression model are then used to apply positive or negative weights to each superpixel region. These weights correlate with how important a region is in the classification of an image. Figure 6 shows an example of LIME’s weighted superpixels feature applied to a substation. In the rightmost part of the figure, the darkest colors indicate areas of higher correlation, indicating the region was weighted heavily in the model’s classification process.

LIME was utilized in our process to validate model classifications by running 100 random samples from each class through LIME to confirm that the classification model was correctly classifying images based on the CI present in the sample image. Figure 7 shows example LIME results for a solar farm, water treatment plant, substation, petrol terminal, airport, and hydroelectric dam. For each pair of images for each facility type, the lefthand images show

the highlighted superpixels defined by LIME. In the righthand images, red indicates areas of negative correlation, and blue indicates areas of positive correlation between superpixel region and probability of a correctly classified image, with values ranging between -1 and 1 .




Figure 6. Example of LIME's weighted superpixels feature applied to a substation.

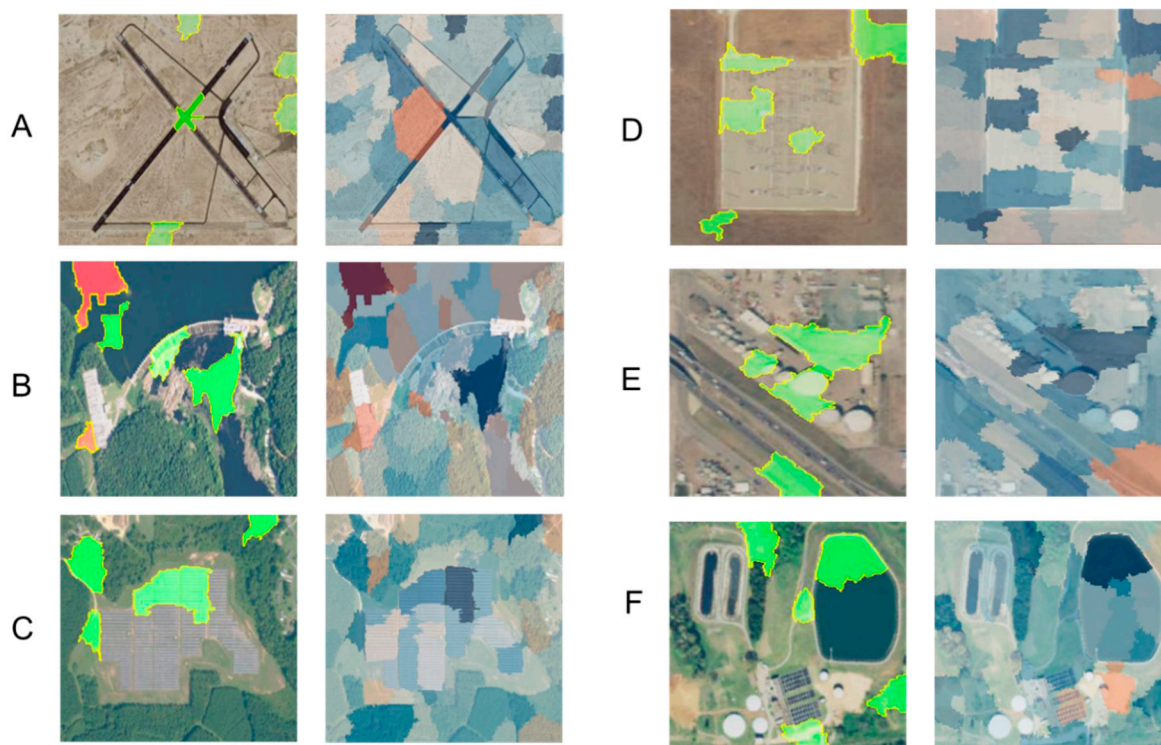


Figure 7. Sample LIME results, with red indicating areas of negative correlation and blue indicating areas of positive correlation for (A) airports, (B) hydroelectric dams, (C) solar panels, (D) substations, (E) petrol terminals, and (F) water treatment plants.

Table 4 gives the LIME results for the nine CI facility classes. LIME provides the top three predictions for a given image. In Table 4, "First Guess" gives the accuracy percentage of LIME's first guess out of the test set. "Overall" provides the correct estimations percentage across the top three predictions. Comparing the results shown in Table 4 with those in Table 3, LIME's performance was similar to the overall model accuracy. For the LIME analysis, noting that a sample size of 100 is a smaller testing data set than was used for

cross validation, this could account for variability in the predicted classes. For the potable water tanks and hospitals classes, LIME accuracy was lower than DenseNet-161 model accuracies. This was attributed to the general ambiguity of the features within both classes. Hospitals appear as generic buildings and potable water tanks appear as circles. As LIME is designed to distinguish unique features in an image, this results in suspected inaccuracies when LIME establishes superpixels for classification for these classes.

Table 4. LIME results for the nine CI facility classes.

	LIME	
	First Guess	Overall
Airports	99	100
Hydroelectric Dams	93	99
Solar Farms	95	99
Hospitals	68	80
Potable Water Tanks	11	41
Substation	88	98
Petrol Terminals	80	100
Natural Gas Generation Plants	40	90
Water Treatment Plants	60	96

4.5. SHAP Implementations

SHAP is another model-agnostic approach for explainability that utilizes cooperative game theory to determine which features of an image are crucial in the classification process. When using images, the pixels can be grouped into regions, distributing the predictions in the regions. For our purposes, we utilized SHAP with DeepExplainer, which is considered an enhanced version of the DeepLIFT algorithm. DeepExplainer approximates the SHAP values when going over several background samples by summing the difference between the expected model output based on the passed background samples and the current model output.

Similar to the analysis conducted with LIME, 100 images were randomly selected and analyzed with SHAP, returning the top three classification predictions for each image. SHAP denotes the correlation between a pixel and the model's weighting of the pixel when classifying the image with pink highlighting when positive and blue highlighting when negative as shown in Figure 8. The accuracy results from SHAP are shown in Table 5. SHAP provides the top three predictions for a given image. In Table 5, both first guess and overall accuracy values are shown. "First Guess" gives the accuracy percentage of SHAP's first guess out of the randomly selected 100 images; "Overall" gives the percentage of correct estimations across the top three predictions. A notable difference in the SHAP results from LIME was SHAP's poor performance across all but two classes (airports and hydroelectric dams) for first guess accuracy. When the top three classifications are considered, SHAP's results improve but still underperforms compared to both model and LIME accuracy. Locating the cause of the SHAP's accuracy discrepancies would require further investigation. However, given the performance of LIME, the LIME-explainability framework is better suited for analysis of CI imagery data and is recommended for use in the overall modeling and analysis pipeline, as detailed in Figure 1.

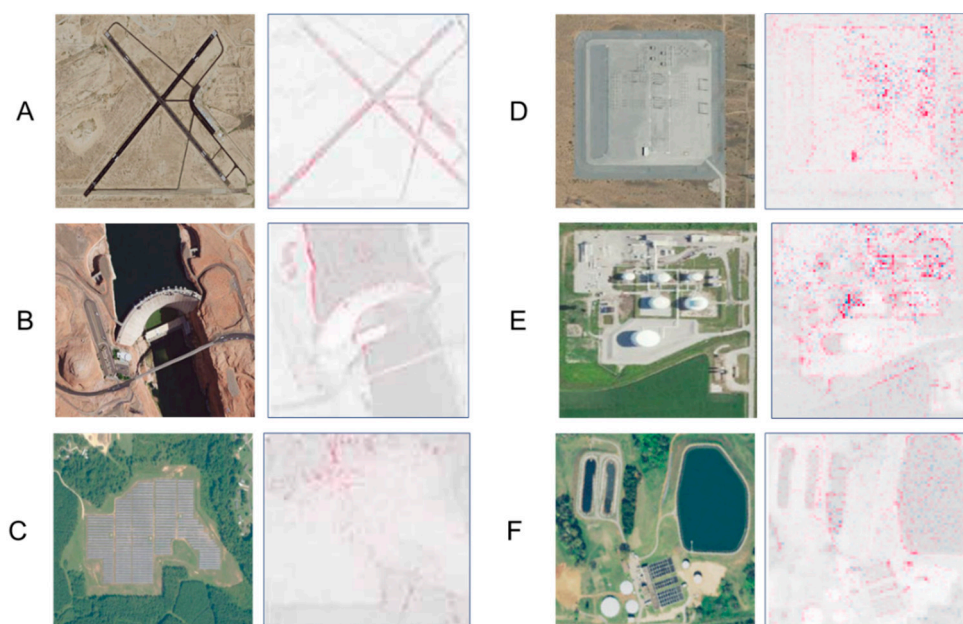


Figure 8. Sample SHAP results, with red indicating areas of positive correlation and blue indicating areas of negative correlation for (A) airports, (B) hydroelectric dams, (C) solar panels, (D) substations, (E) petrol terminals, and (F) water treatment plants.

Table 5. SHAP results for the nine CI facility classes.

	SHAP	
	First Guess	Overall
Airports	95	100
Hydroelectric Dams	98	100
Solar Farms	1	54
Hospitals	64	85
Potable Water Tanks	14	68
Substation	0	2
Petrol Terminals	75	98
Natural Gas Generation Plants	54	92
Water Treatment Plants	1	4

4.6. Dataset Analysis

Considering the range of accuracy results across the CI facility types from Table 3 combined with the outcomes from the explainability analysis, we examined more closely the results for those classes with less than 90% accuracy. Of the original nine classes, the training data for potable water tanks, natural gas generation plants, petroleum terminals, substations, and water treatment plants exhibited less than 90% accuracy. Class outputs were examined using LIME and SHAP to determine which features were being misidentified. Using this approach, we determined that a large source of class confusion was originating from poor training data image quality. Once removed from the data set, two of five classes' (potable water tanks and substations) accuracy levels were increased to a 90% or greater accuracy when tested with cross-validation (Table 3). Data removal was based on two metrics: clarity of an image and the amount of noise or competing non-related class features in an image. While this was performed manually, it is only a one-time effort and does not need to be repeated for use of the data analysis pipeline. Low levels of accuracy in the

remaining three classes were attributed to commonality in imagery data between classes. When tested in isolation from other like classes, class accuracy improved to above the 90% accuracy threshold.

Beyond the aforementioned data quality assurance, data set size is limited by the number of locations where critical infrastructure exists. For example, there are substantially fewer airports than there are substations and so you're left with the option to either have substantially different dataset sizes or to limit the number of locations included from the larger class. An additional option is to introduce synthetically generated data, but as this is a benchmarking study, that is beyond the scope of the work presented here. The final challenge that should be addressed is the geographic diversity across the United States, resulting in a wide range of landscape types in the imagery surrounding infrastructure. This effectively adds noise to the data because the model has no way of knowing what pixels are part of the background versus those that are representative of components of interest as it develops the models. Unfortunately, current methods rely on rectangular areas of interest which makes it difficult to develop a model if the shape of the target of interest is non-rectangular or is positioned at an inconvenient orientation. A potential solution to this in future work is to implement feature masking or even manual reorientation of features to reduce the number of background images included in the training data.

5. Conclusions

This work provides a foundational understanding of how effective deep learning is for CI analysis. Presently, CI analysis is a labor-intensive activity that depends on consistent manual assessments by subject matter experts. This is problematic during crisis conditions when efficiency is key to effective response, such as when a natural disaster occurs. This paper does not solve those problems but provides a baseline understanding of the effectiveness of convolutional neural networks for CI applications. This work benefits from the All Hazards (AHA) database, which includes the most extensive geospatial and dependency-focused data source for critical infrastructure within the United States. Even with AHA, there are still several challenges remaining in this domain, especially relating to the number of data points available and the impacts of geographic diversity.

The method detailed in this work produced a model trained to recognize the nine classes of interest from open-source satellite imagery. It achieved a high degree of accuracy from open-source imagery data. The integration of a trust mechanism using LIME and SHAP provides potential users with a high degree of confidence, particularly with LIME, when assessing model classifications. The work presented here is the first instance of using explainable CNNs in the identification of specific CI facilities in diverse CI sectors. Both the trained model and explainability approaches provide a repeatable and reliable method for identifying the nine classes of CI for which the model was trained. In practice, the method could be utilized in additional CI research and analysis to identify previously unknown CI facilities. The method is transferable for use with new data as updated aerial imagery becomes available. The model is easily rerun with new data to provide timely updated information of the construction or decommissioning of different CI facilities. Given an updated imagery set for classification and proper CI baseline for an area, the model could be utilized for increasing situational awareness of CI assets for disaster preparedness and response.

The nine classes studied in this work represent a significant advancement on prior work. In future work, the number of facilities investigated can be expanded to include the full range of CI facilities that exist. The current nine classes were chosen to demonstrate method applicability across multiple CI sectors with a focus on lifeline sectors. CI sectors are composed of numerous individual facilities. Additional CI facilities fall outside of the scope of identification by traditional satellite imagery data (e.g., buried pipelines or non-descript buildings). Of the facilities that do fall within the scope of traditional satellite imagery, data availability was a determining factor in facility type selection. If there was not enough location data present for a given CI facility type, it was not included in this

work. If additional data is available, the proposed approach can be utilized to identify those facilities.

Expansion of this work would include incorporating semantic segmentation to allow for finer grain analysis of individual components of identified CI facilities. Successful component identification could lead to the successful identification of dependencies, such as estimations of required treatment chemicals at a water treatment plant or the feasible generation capacity of a power plant. Incorporating semantic segmentation would require an expanded higher resolution data set and expanded classification ability.

Author Contributions: Conceptualization, S.N.E., A.J.B.S., and I.T.; methodology, S.N.E., A.J.B.S., and I.T.; software, S.N.E. and A.J.B.S.; validation, S.N.E., A.J.B.S.; formal analysis, I.T.; investigation, S.N.E., A.J.B.S., and E.M.K.; resources, S.N.E., A.J.B.S., and E.M.K.; data curation, S.N.E., A.J.B.S., and E.M.K.; writing—original draft preparation, S.N.E., A.J.B.S., and E.M.K.; writing—review and editing, S.N.E., A.J.B.S., E.M.K., and I.T.; All authors have read and agreed to the published version of the manuscript.

Funding: This work of authorship was prepared as an account of work sponsored by Idaho National Laboratory (under Contract DE-AC07-05ID14517), an agency of the U.S. Government. Neither the U.S. Government, nor any agency thereof, nor any of their employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Data Availability Statement: Data was derived from United States Department of Agriculture's National Agriculture Imagery Program. Specific testing and training data sets can be obtained by contacting the authors.

Acknowledgments: The authors would like to thank the Department of Energy and Idaho National Laboratory for their support of this work and David Friedman for his invaluable technical support.

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

References

1. Ali, M.U.; Sultani, W.; Ali, M. Destruction from sky: Weakly supervised approach for destruction detection in satellite imagery. *ISPRS J. Photogramm. Remote Sens.* **2020**, *162*, 115–124. [CrossRef]
2. Arrieta, A.B.; Diaz-Rodriguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; Garcia, S.; Gil-Lopez, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2019**, *58*, 82–115. [CrossRef]
3. Chen, F.; Ren, R.; Van de Voorde, T.; Xu, W.; Zhou, G.; Zhou, Y. Fast Automatic Airport Detection in Remote Sensing Images Using Convolutional Neural Networks. *Remote Sens.* **2018**, *10*, 443. [CrossRef]
4. Cutter, S.L. Compound, Cascading, or Complex Disasters: What's in a Name? *Environ. Sci. Policy Sustain. Dev.* **2018**, *60*, 16–25. [CrossRef]
5. Datta, U. Infrastructure Change Monitoring Using Multitemporal Multispectral Satellite Images. *Int. J. Civ. Archit. Eng.* **2020**, *14*, 155–160.
6. Davis, D. National Agriculture Imagery Program Information Sheet. Available online: https://www.fsa.usda.gov/Internet/FSA_File/naip_info_sheet_2015.pdf (accessed on 1 June 2022).
7. Elliott, S.N.; Shields, A.J.; Klaehn, E.M.; USDOE Office of Environment, Health, Safety and Security. *Scramble*; Computer Software; OSTI.GOV: Oak Ridge, TN, USA, 2022. Available online: <https://www.osti.gov/servlets/purl/1861032> (accessed on 1 June 2022).
8. Guo, H.; Yang, X.; Wang, N.; Song, B.; Gao, X. A Rotational Libra R-CNN Method for Ship Detection. *IEEE Trans. Geosci. Remote Sens.* **2020**, *58*, 5772–5781. [CrossRef]
9. Hong, D.; Gao, L.; Yokoya, N.; Yao, J.; Chanussot, J.; Du, Q.; Zhang, B. More Diverse Means Better: Multimodal Deep Learning Meets Remote-Sensing Imagery Classification. *IEEE Trans. Geosci. Remote Sens.* **2020**, *59*, 4340–4354. [CrossRef]
10. Hruska, R.; Klett, M. Knowledge Framework for Critical Infrastructure Analysis. In *IEEE Resilience Week*; IEEE: Piscataway, NJ, USA, 2014.
11. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K.Q. Densely connected convolutional networks. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 2261–2269. [CrossRef]
12. Johansen, C.; Tien, I. Probabilistic multi-scale modeling of interdependencies between critical infrastructure systems for resilience. *Sustain. Resilient Infrastruct.* **2017**, *3*, 1–15. [CrossRef]
13. Li, S.; Xu, Y.; Zhu, M.; Ma, S.; Tang, H. Remote Sensing Airport Detection Based on End-to-End Deep Transferable Convolutional Neural Networks. *IEEE Geosci. Remote Sens. Lett.* **2019**, *16*, 1640–1644. [CrossRef]

14. Mokhtarzade, M.; Zoej, M.V. Road detection from high-resolution satellite images using artificial neural networks. *Int. J. Appl. Earth Obs. Geoinf.* **2007**, *9*, 32–40. [[CrossRef](#)]
15. Nie, S.; Jiang, Z.; Zhang, H.; Cai, B.; Yao, Y. Inshore Ship Detection Based on Mask R-CNN. In Proceedings of the IGARSS—2018 IEEE International Geoscience and Remote Sensing Symposium, Valencia, Spain, 22–27 July 2018; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2018; pp. 693–696.
16. Oshri, B.; Hu, A.; Adelson, P.; Chen, X.; Dupas, P.; Weinstein, J.; Burke, M.; Lobell, D.; Ermon, S. Infrastructure Quality Assessment in Africa using Satellite Imagery and Deep Learning. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, London, UK, 19–23 August 2018; Association for Computing Machinery: New York, NY, USA, 2018; pp. 616–625. [[CrossRef](#)]
17. Rinaldi, S.; Peerenboom, J.; Kelly, T. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Syst.* **2001**, *21*, 11–25. [[CrossRef](#)]
18. Security Agency, Infrastructure. *National Critical Functions 2021 Status Update to the Critical Infrastructure Community*; Security Agency: Washington, DC, USA, 2021.
19. Tabian, I.; Fu, H.; Khodaei, Z.S. A Convolutional Neural Network for Impact Detection and Characterization of Complex Composite Structures. *Sensors* **2019**, *19*, 4933. [[CrossRef](#)] [[PubMed](#)]
20. Ribeiro, M.T.; Singh, S.; Guestrin, C. ‘Why Should I Trust You?’ Explaining the Predictions of Any Classifier Marco. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016.
21. U.S. Department of Homeland Security. *NIPP 2013: Partnering for Critical Infrastructure Security and Resilience*; Security Agency: Washington, DC, USA, 2013.
22. U.S. Department of Homeland Security. *A Guide to Critical Infrastructure Security and Resilience*; Security Agency: Washington, DC, USA, 2019.
23. Wu, X.; Hong, D.; Chanussot, J. Convolutional Neural Networks for Multimodal Remote Sensing Data Classification. *IEEE Trans. Geosci. Remote Sens.* **2021**, *60*, 1–10. [[CrossRef](#)]
24. Yang, X.; Sun, H.; Fu, K.; Yang, J.; Sun, X.; Yan, M.; Guo, Z. Automatic Ship Detection in Remote Sensing Images from Google Earth of Complex Scenes Based on Multiscale Rotation Dense Feature Pyramid Networks. *Remote Sens.* **2018**, *10*, 132. [[CrossRef](#)]
25. Yin, S.; Li, H.; Teng, L. Airport Detection Based on Improved Faster RCNN in Large Scale Remote Sensing Images. *Sens. Imaging* **2020**, *21*, 49. [[CrossRef](#)]
26. Zhang, P.; Niu, X.; Dou, Y.; Xia, F. Airport Detection on Optical Satellite Images Using Deep Convolutional Neural Networks. *IEEE Geosci. Remote Sens. Lett.* **2017**, *14*, 1183–1187. [[CrossRef](#)]
27. Zhang, S.; Wu, R.; Xu, K.; Wang, J.; Sun, W. R-CNN-Based Ship Detection from High Resolution Remote Sensing Imagery. *Remote Sens.* **2019**, *11*, 631. [[CrossRef](#)]