CRACK DAMAGE DETECTION IN UNMANNED AERIAL VEHICLE IMAGES OF CIVIL INFRASTRUCTURE USING PRE-TRAINED DEEP LEARNING MODEL

Kasthurirangan Gopalakrishnan[1], Hoda Gholami[2], Akash Vidyadharan[3], Alok Choudhary[4], Ankit Agrawal[5]

[1,4,5] *Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois, USA*

[2,3] *InfraDrone LLC, Des Moines, Iowa, USA*

Abstract: The field of computer vision based civil infrastructure defect detection is constantly evolving with steady advances being made in sensing technologies, hardware, and image processing techniques. Although a number of image processing techniques have been proposed over the years with varying degrees of success, real-world situations (e.g., lack of background illumination, shadow changes) pose a significant challenge to the wide adoption of such techniques, especially for routine analysis. With the emerging application of Unmanned Aerial Vehicles (UAVs) or drones for civil infrastructure inspection and condition monitoring, an added level of complexity is introduced in vision-based crack detection as the process of acquiring images using UAVs is not yet standardized resulting in images of widely varying sizes, resolutions, blurriness, etc. In recent years, Deep Learning, a generalized form of Deep Neural Network (DNN) algorithms that can learn very complex mappings between inputs and outputs directly from the data, has achieved huge success in diverse fields such as automatic speech recognition, image recognition, Natural Language Processing (NLP), drug and materials discovery, etc. However, the large number of hidden neurons and layers used in DNNs result in computationally-intensive matrix and vector computations involving millions of parameters, requiring the use of high-performance computing systems. Also, it is practically impossible to get labeled "big data" samples in many domains to be able to train an entire DNN from scratch. In such situations, the use of a pre-trained deep learning model and fine-tuning it to the novel task at hand with smaller datasets, has shown to be successful across domains. In this paper, we propose the use of pre-trained deep learning models with transfer learning for crack damage detection in UAV images of civil infrastructure. The robustness of the proposed approach is tested on a small set of real-world, complex UAV-sourced infrastructure images not used during training and validation. The results show that the proposed method can rapidly and easily achieve up to 90% accuracy in finding cracks in realistic situations without any augmentation and preprocessing.

Keywords: crack detection, unmanned aerial system (UAS), drone, deep learning, convolutional neural networks, random forest.

---

[1] Corresponding author: rangan@northwestern.edu

1. Introduction

The 2017 American Society of Civil Engineers (ASCE) Infrastructure Report Card assigned a 'D' grade for United States infrastructure, which includes bridges, roads, rail, levees, energy, etc. (ASCE, 2017). Efficient condition monitoring strategies can aid engineers in developing appropriate scheduling of infrastructure maintenance and repair activities leading to significant reduction in infrastructure life-cycle maintenance costs. In recent years, Unmanned Aerial Vehicles (UAVs), commonly referred to as drones, are witnessing unprecedented growth and applications ranging from disaster management and construction surveying to infrastructure inspection, etc. (Floreano and Wood, 2015). On August 2, 2016, the White House Office of Science and Technology Policy (OSTP) in the United States hosted the first-ever workshop on Drones and the Future of Aviation to advance and celebrate the potential of UAVs. In this workshop, a series of actions were initiated to promote the safe integration and adoption of UAVs across the US, including $35 million in research funding by the National Science Foundation (NSF) (which includes research on UAVs for monitoring and inspection of critical infrastructure systems). In this paper, consistent with the common usage, the terms UAV, Unmanned Aerial System (UAS), and drone are used interchangeably and they all refer to the same thing.

The UAVs can be used for a variety of inspection and asset management applications ranging from structural health monitoring of bridges, buildings and wind turbines to the damage assessment of infrastructure after disaster events such as Hurricanes Katrina, Wilma and Ike;

Typhoon Morakot; and the 2009 L'Aquila, 2010 Haiti and 2011 Japan Earthquakes. Inspecting civil and transportation infrastructure manually is an expensive and relatively dangerous task. The use of UAV for infrastructure inspection purposes can make this task significantly easier, safer and cost effective (Vidyadharan *et al.,* 2017). Many times a lift or a crane is used to allow manual access to inspect the 'hard to reach' locations where damages may have occurred. UAVs provide the accessibility and flexibility to navigate around complex structure and collect data with equal or higher quality to that collected by manual inspectors. The data collected by UAVs can span across an array of data types such as 3-D rendering using photogrammetry, High Definition (HD) visual imagery, Light Detection & Ranging (LiDAR) scanning, infrared thermography, multispectral modelling and Non Destructive Evaluation (NDE) using Ultrasonic Testing (UT). These methods can be used for monitoring both the microscale and macroscale defects of infrastructures resulting in detailed inventory, survey, and condition assessment of civil and transportation infrastructure systems.

The main challenge for visual inspection of large civil infrastructures using UAVs is to estimate the appropriate distance to be maintained between the UAV equipped with camera and the target to be able to achieve reasonable detection sensitivity (Ellenberg *et al.,* 2014). However, in recent years, the availability of small-size, lightweight, high-resolution cameras make it possible for UAVs to acquire HD images of even micro-cracks.

The valuable data such as HD images obtained using UAVs can be used to train machine learning and deep learning algorithms to make the process of identifying

defects on the infrastructure from images efficient. Such images collected using UAVs can come from different types of infrastructure such as roads, bridges, silos, wind turbines and concrete walls.

We propose the use of a pre-trained Deep Convolutional Neural Network (DCNN) model with transfer learning for automated crack detection in UAV images. Deep Learning (DL) has recently been used for road crack detection but since it typically requires large amounts of labeled training data, here we propose the use of pre-trained DCNN with transfer learning for automated crack detection, and demonstrate its advantages. The rest of the paper is divided as follows. A review of state-of-the-art vision-based automated crack detection methods is discussed first, with specific focus on UAV visual inspection, under Related Works followed by a brief introduction to the concept of DCNN with transfer learning for cross domain image analysis. The Research Study is then presented which includes description of the datasets used, the proposed methodology, the experiments carried out and the results.

2. Deep Learning in UAV based Visual Inspection: A Brief Review

The rapidly ageing civil infrastructure systems pose a serious challenge to the safe operations, structural and functional health management, especially in the context of ever-increasing budgetary constraints faced by system operators. In recent years, the UAVs have emerged as a time-saving and cost-effective alternative for carrying out vision-based real-time Structural Health Monitoring (SHM) through full field mapping of large civil structures (Sankarasrinivasan *et al.,* 2015). Pereira and

Pereira (2015) proposed embedded image processing system implementations for automatic recognition of cracks in building facades using UAV images. They concluded that a combined aerial-based and ground-based detection approach would prove to be a viable approach: the embedded particle filter on the Raspberry-Pi in UAV can be used for preliminary detection of crack, and, if the probability of crack presence is higher, that image would be stored for further processing at the ground station using a Sobel operation filter (Pereira and Pereira, 2015).

(Ellenberg *et al.,* 2014) proposed the use of unmanned multicopter equipped with a 16MP camera for visual inspection of masonry buildings and preliminary detection of cracks. Some challenges arising from roughness of surface, masonry patterns, edges from doors, windows, etc. were reported with the application of an edge detection algorithm to UAV images of masonry walls. To ensure accurate localization of cracks in the images, (Ellensberg *et al.,* 2014) suggested that the angle at which the UAV and gimbal are oriented and the GPS location should all be recorded accurately. Even then GPS readings can be inaccurate due to interference from the structure being observed and one solution to that could be to stay farther away from the structure.

Although machine learning based techniques have been successfully applied to object detection in UAV images in the context of SHM, the use of Deep Learning (DL) for prognostics and damage assessment of civil infrastructure systems is still an emerging area of research, mainly due to the difficulties in generating good quality labeled training image datasets (Sarkar *et al.,* 2016; Yokoyama and Matsumoto, 2017). Some recent studies, not exclusive to the

SHM context, report on the use of DL for UAV-based applications. For instance, (Lee *et al.,* 2017) proposed a novel object detection method wherein the images captured using the UAV are processed in the off-board computing cloud using the state-of-the-art Faster Regions with CNNs (R-CNNs), while low-level object detection and short-term navigation takes place onboard. They tested their approach with a Parrot AR.Drone 2.0 as a low-cost, lightweight drone in an indoor environment and reported that their hybrid approach is able to approach real-time performance even when detecting hundreds of object categories, despite the unpredictable communication lags resulting from the use of cloud-based computation resources.

More recently, (Carrio *et al.,* 2017) provided a thorough review on the reported uses and applications of DL for UAVs, including the major challenges for the application of DL for UAV-based solutions. Although the UAVs can be equipped with a variety of sensor payloads (such as LiDAR and other non-destructive sensing devices), (Carrio *et al.,* 2017) found that most successful applications of DCNN to UAVs were with respect to visual data (images) owing to the low-cost, lightweight, and low power consumption of image sensors.

3. Deep Convolutional Neural Networks and Transfer Learning

Deep Neural Networks (DNNs), more commonly referred to as Deep Learning (DL), employ deep NN architectures to automatically learn hierarchy of features from raw input data without the need for feature engineering (Hinton *et al.,* 2006). Loosely inspired by how the mammalian brain uses different areas of the cortex to abstract different levels of features when given an input percept, deep learning methods are characterized by deep architectures with several hidden layers that allow them to learn many levels of abstraction, as opposed to shallow architectures with 1 or 2 hidden layers (Xie *et al.,* 2017).

(Hinton *et al.,* 2006) proposed Deep Belief Networks (DBNs) with a new unsupervised training method called the layer-wise-greedy-training, which gave rise to the popularity of deep learning methods. Going beyond traditional machine learning and artificial intelligence approaches that depend on human-crafted features, DL technologies, with their ability for automatic high-level feature abstraction combined with exponential growth of data, have already produced breakthrough results in Natural Language Processing (NLP), speech recognition, computer vision, and materials informatics (Agrawal and Choudhary, 2016). There are four main DL architectures, namely the Restricted Boltzman Machines (RBMs), DBNs, Autoencoder (AE), and Deep Convolutional Neural Networks (DCNNs or Deep ConvNets) (Liu *et al.,* 2017).

DCCNs have shown to be highly effective in processing visual data, such as images and videos. DCNNs take raw input data at the lowest level and transforms them by processing them through a sequence of basic computational units to obtain representations that have intrinsic values for classification in the higher layers (Bai, 2017; Goh *et al.,* 2014). A DCNN typically consists of three layer types (Fig. 1): convolution layers, subsampling layers, and fully connected layers. A convolutional layer is parameterized by the number of channels, kernel size, stride factor, border mode, and the connection table. The convolution layer

takes the input image and applies convolution filter on it to produce the output image or the filter response. Multiple convolutional layers are used to take into consideration the spatial dependencies among image pixels. The subsampling layer is used to make the neural network more invariant and robust. The max-pooling method is commonly used for the subsampling layer as it has shown to lead to faster convergence and better generalization. It is common to use multiple fully-connected layers after several rounds of convolution, and the resulting structure of the last convolutional layer is flattened before connecting to the following fully-connected layer.
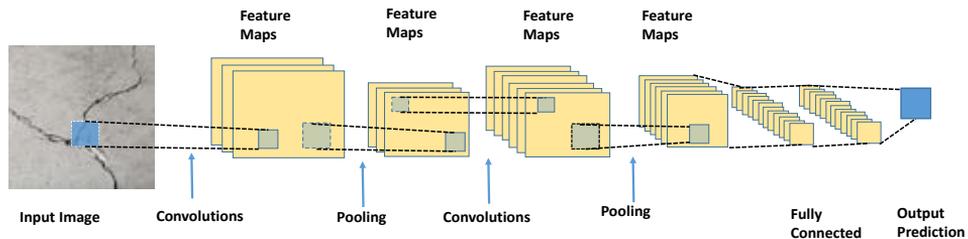
Fig. 1.

A Schematic of Convolutional Neural Network (CNN) Architecture

DCNNs typically require large annotated image datasets to achieve high predictive accuracy. However, in many domains, acquisition of such data is difficult and labeling them is costly. In light of these challenges, the use of 'off-the-shelf' DCNN features of well-established DCNNs such as VGG-16, AlexNet, and GoogLeNet pre-trained on large-scale annotated natural image datasets (such as ImageNet) have been shown to be very useful for solving cross domain image classification problems through the concept of transfer learning and fine-tuning (Shin *et al.,* 2016). In DCNN, representations learned at different layers of the network correspond to different levels of abstractions present in the input images. The initial layers extract edges and color information while the latter layer filters encode shape and texture. The idea behind transfer learning is that it is cheaper and efficient to use deep learning models trained on "big data" image datasets (like ImageNet) and "transfer" their learning ability to new classification scenario rather than train a DCNN classifier from scratch (Bar *et al.,* 2015). With adequate fine-tuning, pre-trained DCNN has been shown to outperform even DCNN trained from scratch for some medical imaging applications (Shin *et al.,* 2016; Tajbakhsh *et al.,* 2016).

4. Research Study

The goal of this study was to develop a simplified, crack detection model from UAV images of civil infrastructure through the application of deep transfer learning approach. Accordingly, the experimental study involved acquisition of UAV visual data (images containing cracks and no cracks) and the development of automated, UAV-vision-based crack detection system through the application of ImageNet pre-trained VGG-16 DCNN to the UAV images.

4.1. UAV Images Dataset

In this study, we used a Hexacopter UAV (Fig. 2) to collect close-up images of few common civil infrastructure systems (like storage silos, local roadways, etc.) due to attributes such as high stability and precise control. This hexacopter was equipped with state-of-the-art guidance and control systems to validate precision controls and enable autonomous flight capabilities so as to minimize the role of a human operator in the loop. It is capable of carrying a payload of 15 lbs with an endurance of 30 minutes. The HD visual inspection is carried out using the Hexacopter-I UAV with a 30 MP high definition Canon EOS 5D Mk IV DSLR camera mounted on a 3-axis rotatable gimbal with live video transmission, which allows the inspector to change the direction of the camera and focus to get better pictures of the defects on the structure. Examples of HD images of cracks and corrosion on different structures, as well as those with no cracks, captured by this camera are shown in Fig. 3 and Fig. 4.

Fig. 2.
Hexacopter UAV Equipped with High Definition Camera Used for Visual Inspection in this Research Study
Source: authors

The primary classes of image categorization would be the cracked subset (80 images) and uncracked subset (50 images) of images. The images obtained from UAV can vary based on proximity to the surface, angle and can be collected more often, thereby allowing for the flexibility of multiple iterations of the same defect present on the infrastructure.

Fig. 3.
Examples of UAV Images of Infrastructure with Cracks
Source: authors

Fig. 4.
Examples of UAV Images of Infrastructure with No Cracks
Source: authors

4.2. Crack Detection Approach Using Pre-trained Deep Learning Model

We used the Keras deep learning framework (Chollet, 2015) that includes pre-trained deep learning models made available alongside weights within Keras Applications. Specifically, we used the Keras implementation of the VGG-16 model (16-layer DCNN developed by the Visual Geometry Group (VGG) at the University of Oxford), with weights pre-trained on ImageNet database in this study (Pal, 2016; Simonyan and Zisserman, 2014). A schematic illustration of the VGG-16 architecture is shown in Fig. 5. The ImageNet database, built upon the hierarchical structure of WordNet, contains more than 3.2 million cleanly annotated images distributed over 5,247 categories (Deng *et al.*, 2009). Since the pre-trained VGG-16 model has learned to extract features from images that can distinguish one image class from another, they have shown to achieve excellent performance even when applied to image recognition and classification datasets in

other domains (Pal, 2016; Simonyan and Zisserman, 2014). When the fully-connected classifier (or the bottleneck layer) is removed from the pre-trained VGG-16 network, it can be used as a deep feature generator for producing semantic image vectors for our UAV images. These semantic image vectors can then be trained and tested using another classifier (like Neural Networks [NN], Support Vector Machine [SVM], Random Forest [RF], etc.) for predicting the class labels ('crack' or 'no_crack').
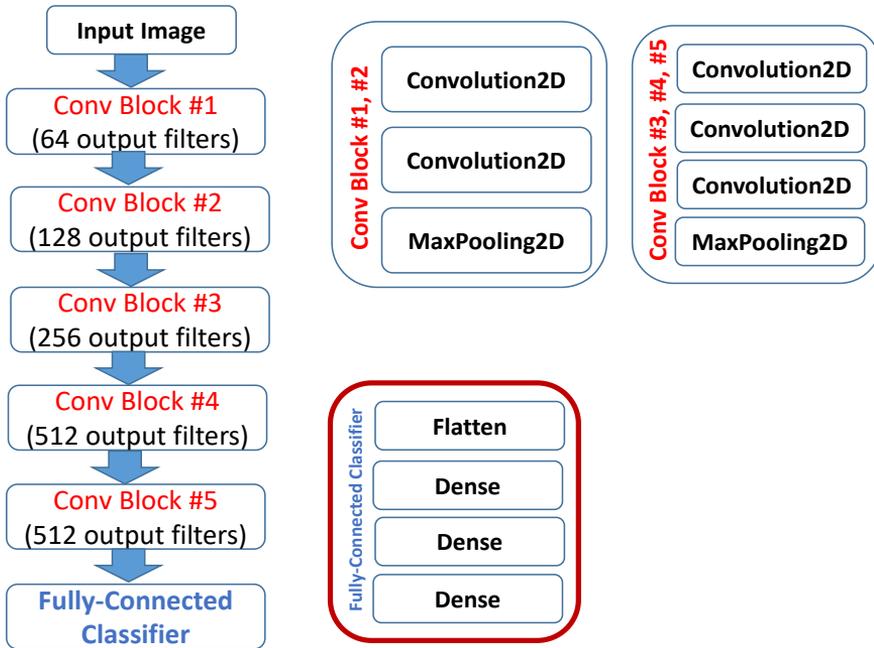
Fig. 5.
A Schematic Illustration of VGG-16 Architecture

4.3. Results and Discussion

In this study, the truncated VGG-16 DCNN is used as deep feature generator for our UAV images of civil infrastructure. We train only the final classifier layer using ImageNet pre-trained VGG-16 DCNN features for the extracted semantic image vectors. The overall framework of the proposed methodology is presented in **Error! Reference source not found.**. Various experiments were conducted to assess the pavement crack detection performance of the pre-trained VGG-16 DCNN network with transfer learning. For the final classifier layer, we employ and compare specific machine learning classifiers that have shown promising results in previous vision-based crack detection studies, such as Neural Networks (NN), Support Vector Machine (SVM), and Random Forest (RF). Apart from these, we also included Logistic

Regression (LR) and Extremely Randomized Tree (ERT) (Geurts *et al.,* 2006) in our investigations.

The following designations are used for the various deep transfer learning model configurations evaluated in this study for automated crack detection from UAV images of civil infrastructure (specific parametric values for each of the machine learning classifiers were chosen based on the successful results obtained from a previous study by (Gopalakrishnan *et al.,* 2017)):

- TL + NN (AM): Single NN layer classifier trained on ImageNet pre-trained VGG-16 DCNN features for UAV images (optimizer: 'adam' (AM) as implemented in Keras);

- TL + RF (300): Random Forest classifier trained on ImageNet pre-trained VGG-16 DCNN features for UAV images (number of trees: 300);
- TL + ERT (300): Extremely Random Trees classifier trained on ImageNet pre-trained VGG-16 DCNN features for UAV images (number of trees: 300);
- TL + SVM (L,5): Support Vector Machines classifier trained on ImageNet pre-trained VGG-16 DCNN features for UAV images (kernel type = 'linear' [L]; C = 5.0);
- TL + LR (L1): Logistic Regression classifier trained on ImageNet pre-trained VGG-16 DCNN features for UAV images (penalty = 'L1').
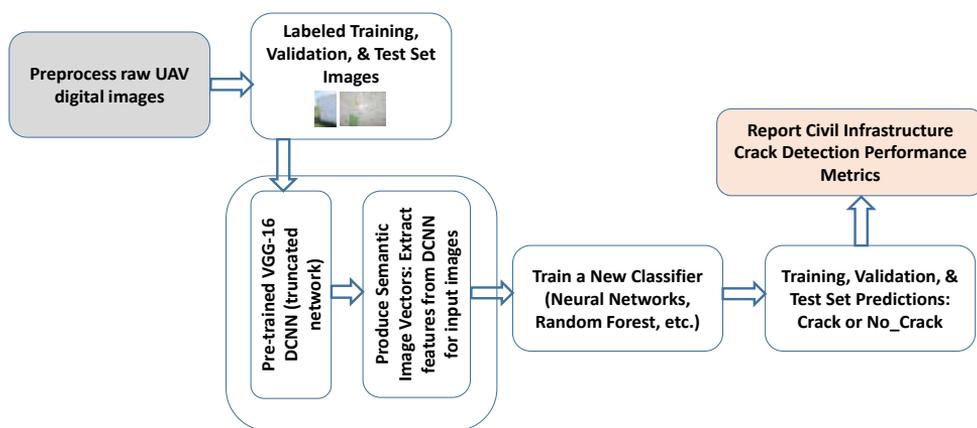
Fig. 6.
Overall Framework of the Proposed Deep Transfer Learning Approach applied to UAV Images

Deep transfer learning was implemented with the Keras (Chollet, 2015) deep learning framework, using a Intel® Core™ i7-5600U CPU on 64-bit Windows® 10 OS. The ImageNet pre-trained VGG-16 DCNN implemented within Keras Applications takes in a default image input size of 224 x 224. Other image dimensions not smaller than 48 x 48 are valid as well. A 224 x 224 image input dimension (which was used in this study as well) results in 25,088 deep transfer learning features that serve as inputs to the final classifier.

For the single-layer NN classifier implemented in Keras, 256 neurons were used in the hidden layer and a dropout value of 0.5 was used. The 'relu' activation was used in the hidden layer and 'softmax' activation in the output layer. The image batch size was set to 32 and all models were trained for up to 50 epochs. Among the total number of UAV images (80 with 'crack' label and 50 with 'no_crack' label), 80% was used for training and validation and the rest 20% were used for testing. Among the 80% of the images dataset used for training and validation, 10% was used for validation alone. The early stopping criteria was used with the final model being the one with low validation loss. Traces of training and validation accuracy and loss (average categorical cross-entropy error) during the training of single-layer NN classifier (optimizer: 'adam') with ImageNet pre-trained VGG-16 DCNN deep image features are shown in Fig. 7.
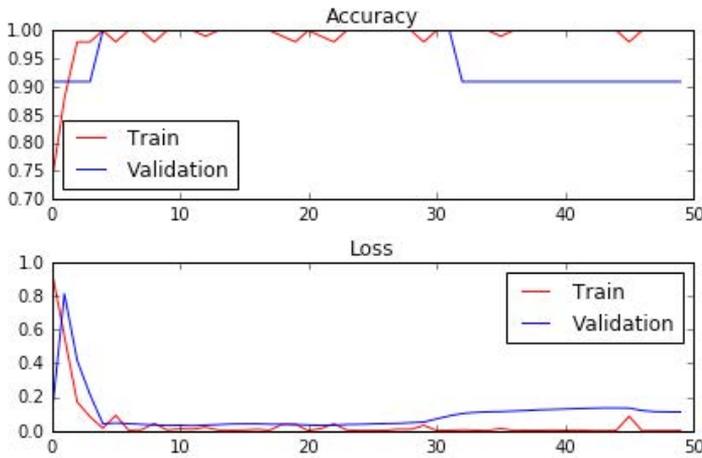
Fig. 7.

Training and Validation Accuracy (top) and Loss (bottom) during Training Single-layer Neural Network (NN) with Pre-trained Deep Learning Model Semantic vectors (x-axis represents Epochs and y-axis represents Accuracy in the Top Figure and Loss in the Bottom Figure, respectively)

Standard implementations in 'scikit-learn' machine learning library in Python were used for all other classifiers. A comparison of pavement crack detection results using various deep transfer learning models is presented in Table 1. Classification performance metrics include the accuracy, precision, recall, F1-score and Cohen's Kappa score (Blackman and Koval, 2000).

Table 1

Comparison of Crack Detection Results on UAV Test Set Images of Civil Infrastructure using Pre-trained Deep Learning Model with Various Final Classifiers

Final Classifier	Model	Accuracy	Precision	Recall	F1-score	Cohen's Kappa Score
Single-layer Neural Network (NN)	TL + NN (AM)	0.89	0.91	0.89	0.89	0.788
Random Forest (RF)	TL + RF (300)	0.86	0.86	0.86	0.85	0.622
Extremely Randomized Trees (ERT)	TL + ERT (300)	0.82	0.87	0.82	0.82	0.650
Support Vector Machine (SVM)	TL + SVM (L,5)	0.75	0.84	0.75	0.74	0.515
Logistic Regression (LR)	TL + LR (L1)	0.89	0.91	0.89	0.89	0.788

Among these 5 classifiers, the single-layer NN classifier (with Adam optimizer) and logistic regression classifiers (penalty = 'L1') trained on ImageNet pre-trained VGG-16 DCNN features yielded the best performance, followed by RF and ERT classifiers. This result is further confirmed by the Area-Under-the-ROC-Curve (AUC) values for the different classifiers in Fig. 8.
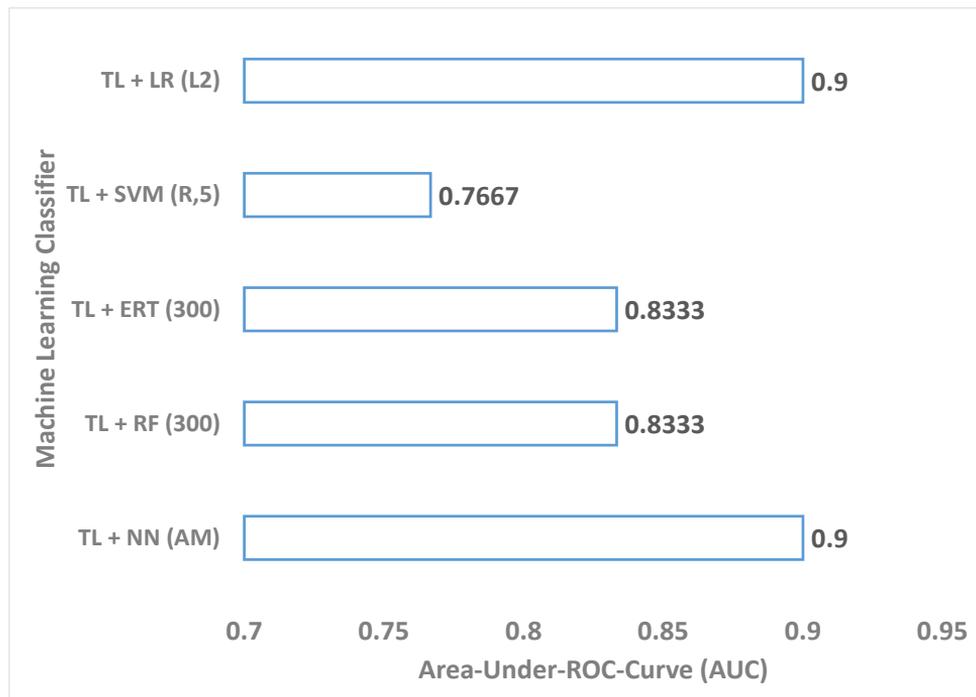
Fig. 8.
Comparison of Area-Under-the-ROC-Curve Values

4. Conclusions

In this paper, we explored and evaluated the deep transfer learning approach, *viz.*, use deep learning models trained on "big data" image datasets (ImageNet) and "transfer" their learning ability to automatic crack detection from complex UAV images of civil infrastructure systems. A single-layer neural network classifier and logistic regression classifier trained on ImageNet pre-trained VGG-16 DCNN features yielded the best performance. The results show that the proposed method can rapidly and easily achieve up to 90% accuracy in finding cracks in realistic situations without any augmentation and preprocessing.

Our future research efforts will focus on addressing the following challenges:

- Since there is no publicly available big UAV images dataset of civil infrastructure systems, or the available big image datasets (such as ImageNet) have low contextual similarity with our images of interest, we will try to develop a module for collecting and preprocessing UAV images in order to create a well-organized, clean and big dataset of images for predictive analysis;
- We'll extend our classification algorithm to include object detection which makes it possible to detect cracks and crop the crack region and send it for further image processing instead of having the whole image in the memory;
- Explore other classification methods with or without the use of pre-trained deep learning models to continuously improve the algorithms;
- Extend our algorithms to detect and classify other types of defects such as corrosion, dents, etc.;

- Combining our UAV based deep learning algorithms with big data techniques such as Hadoop, MapReduce to handle complex computing;
- Use of these deep learning techniques on data obtained not just from a HD camera, but also from infrared thermal cameras and 3D rendered models obtained from LiDAR scanners and photogrammetry would be a major breakthrough in identifying defects of various types and dimensions from UAV images of civil infrastructure systems.

Acknowledgements

References

Agrawal, A.; Choudhary, A. 2016. Perspective: Materials informatics and big data: Realization of the "fourth paradigm" of science in materials science, *APL Materials* 4(5): 1-9.

ASCE. 2017. American Society of Civil Engineers (ASCE) 2017 Infrastructure Report Card: Roads. Available from internet: < https://www.infrastructurereportcard.org/>.

Bai, S. 2017. Growing random forest on deep convolutional neural networks for scene categorization, *Expert Systems with Applications* 71: 279–287.

Bar, Y., Diamant, I., Wolf, L., Lieberman, S., Konen, E., Greenspan, H., 2015. Chest pathology detection using deep learning with non-medical training. *In Proceedings of the IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, 294–297.

Carrio, A.; Sampedro, C.; Rodriguez-Ramos, A.; Campoy, P. 2017. A Review of Deep Learning Methods and Applications for Unmanned Aerial Vehicles, *Journal of Sensors* 2017(3296874): 1-13.

Chollet, F. 2015. keras. GitHub. Available from internet: < https://keras.io/>.

Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; Fei-fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 248-255.

Ellenberg, A.; Kontsos, A.; Bartoli, I.; Pradhan, A. 2014. Masonry Crack Detection Application of an Unmanned Aerial Vehicle. In *Proceedings of the 2014 International Conference on Computing in Civil and Building Engineering*, 1788–1795.

Floreano, D.; Wood, R.J. 2015. Science, technology and the future of small autonomous drones, *Nature* 521: 460–466.

Geurts, P.; Ernst, D.; Wehenkel, L. 2006. Extremely randomized trees, *Machine Learning* 63: 3–42.

Goh, H.; Thome, N.; Cord, M.; Lim, J.H. 2014. Learning Deep Hierarchical Visual Feature Coding, *IEEE Transactions on Neural Networks and Learning Systems* 25: 2212–2225.

Gopalakrishnan, K.; Khaitan, S.K.; Choudhary, A.; Agrawal, A. 2017. Deep Convolutional Neural Networks with transfer learning for computer vision-based data-driven pavement distress detection, *Construction and Building Materials* 157: 322–330.

Hinton, G.E.; Osindero, S.; Teh, Y.-W. 2006. A Fast Learning Algorithm for Deep Belief Nets, *Neural Computing* 18: 1527–1554.

Lee, J.; Wang, J.; Crandall, D.; Šabanović, S.; Fox, G. 2017. Real-Time, Cloud-Based Object Detection for Unmanned Aerial Vehicles. In *Proceedings of the 2017 First IEEE International Conference on Robotic Computing (IRC)*, 36–43.

Liu, W.; Wang, Z.; Liu, X.; Zeng, N.; Liu, Y.; Alsaadi, F.E. 2017. A survey of deep neural network architectures and their applications, *Neurocomputing* 234: 11–26.

Pal, S.; 2016. Transfer Learning and Fine Tuning for Cross Domain Image Classification with Keras. GitHub. Available from internet: < https://github.com/sujitpal/fttl-with-keras>.

Pereira, F.C.; Pereira, C.E. 2015. Embedded Image Processing Systems for Automatic Recognition of Cracks using UAVs. In *Proceedings of the 2nd IFAC Conference on Embedded Systems, Computer Intelligence and Telematics (CESCIT 2015)*, 16–21.

Sankarasrinivasan, S.; Balasubramanian, E.; Karthik, K.; Chandrasekar, U.; Gupta, R. 2015. Health Monitoring of Civil Structures with Integrated UAV and Image Processing System. In *Proceedings of the Eleventh International Conference on Image and Signal Processing (ICISP 2015)*, 508–515.

Sarkar, S.; Reddy, K.; Giering, M.; Gurvich, M. 2016. Deep Learning for Structural Health Monitoring: A Damage Characterization Application. In *Proceedings of the Annual Conference of the Prognostics and Health Management Society*, 1-7.

Shin, H.-C.; Roth, H.R.; Gao, M.; Lu, L.; Xu, Z.; Nogues, I.; Yao, J.; Mollura, D.; Summers, R.M. 2016. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning, *IEEE transactions on medical imaging* 35: 1285–1298.

Simonyan, K.; Zisserman, A. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition, In *Proceedings of the 2015 International Conference on Learning Representation (ICLR 2015)*, 1-14.

Tajbakhsh, N.; Shin, J.Y.; Gurudu, S.R.; Hurst, R.T.; Kendall, C.B.; Gotway, M.B.; Liang, J. 2016. Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning? In *Proceedings of the IEEE Transactions on Medical Imaging*, 35: 1299–1312.

Vidyadharan, A.; Carter, T.; Ceylan, H.; Bloebaum, C.; Gopalakrishnan, K.; Kim, S. 2017. Civil Infrastructure Health Monitoring and Management Using Unmanned Aerial Systems. In *Proceedings of the 2017 International Conference on Highway Pavements and Airfield Technology*, 1-8.

Xie, D.; Zhang, L.; Bai, L. 2017. Deep Learning in Visual Computing and Signal Processing, *Applied Computational Intelligence and Soft Computing* 2017(e1320780): 1-13.

Yokoyama, S.; Matsumoto, T. 2017. Development of an Automatic Detector of Cracks in Concrete Using Machine Learning. In *Proceedings of The 3rd International Conference on Sustainable Civil Engineering Structures and Construction Materials*, 1250–1255.