OPEN ACCESS Freely available online

Journal of Remote Sensing & GIS

Research Article

A Machine Learning Classification to Identify Houses Suitable for Electric Vehicle Charging from Remotely Sensed Imagery

James Flynn*1, Giannetti C2

1Materials and Manufacturing Academy, Swansea University, Fabian Way, Skewen, Swansea SA1 8EN

2Swansea University, Fabian Way, Skewen, Swansea SA1 8EN

ABSTRACT

Over the past decade, Deep Convolutional Neural Networks (DCNN's) have emerged as a powerful tool for the classification of remotely sensed imagery. In this multi-disciplinary paper, we demonstrate a novel application of machine learning in the field of remote sensing by developing a workflow to survey urban areas for residential properties suitable for electric vehicle charging. A fine-tune transfer learning approach is presented as a new method for analysing remotely sensed image data. A unique dataset comprised of Google Street View images sourced from multiple UK towns and cities is used to train can compare three neural networks and represents the first attempt to classify residen- tial driveways from streetscape imagery using machine learning. When testing the full workflow on two urban areas the full system achieves accuracies of 87.2% and 89.3% respectively. This proof of concept demonstrates a promising new application of deep learning in the field of remote sensing, geospatial analysis, and urban planning.

Keywords: Transfer Learning, Electric Vehicles, Remote Sensing

INTRODUCTION

Machine learning and computer vision algorithms have proven to be powerful tools for classification of remotely sensed imagery. Over the past decade, as distributed smart systems and the Internet-of-Things have become increasing widespread, new technologies have emerged to deal with the vast amount of data produced by these systems. The rapid development of smart devices, artificial neural networks, and computer vision has meant recent approaches to image classification in remote sensing rely increasingly on machine learning techniques Srivastava et al. (2019); Leung and Newsam (2012); Kang et al. (2018); Huang et al. (2018). Computer vision algorithms have proven to be especially powerful tools due to their versatility, scalability and low cost Enr ́ıquez et al. (2017). Despite the large amount of research into Land Use and Land Cover (LULC) classification using remote sensing imagery, the identification and classification of private off-street parking represents a significant gap in the literature.

Following the Paris Climate Agreement in November 2016, the UK government made a series of announcements to tackle greenhouse gas emissions. The production and sales of petrol and diesel engine vehicles are to be banned in the UK by 2040, and Clean Air Zones (CAZ) are to be introduced across major cities and local authorities Department for Environment Food and Rural Affairs (b,a). Battery Electric Vehicles (BEVs) and Plug-in Hybrid Electric Vehicles (PHEVs) are exempt to these government restrictions and, depending on CO2 pipeline emissions, plug-in vehicle grants are available for UK customers Palmer et al. (2018). These factors, combined with their low running costs and increasing affordability, make BEVs and PHEVs serious contenders in the future of UK sustainable personal mobility. The UK government recognize that availability of accessible and affordable home charging options is key to increasing the uptake of plug-in vehicles. At present, the available evidence suggests that most plug-in vehicle owners will carry out the largest proportion of their charging at home for Low Emission Vehicles (2019). For this reason, the updated 'UK Electric Vehicle Homecharge Scheme' was introduced this year offering a grant of up to £500 to support customers towards the purchase and installation of a home Electric Vehicle (EV) charging point for Low Emission Vehicles (2019). In order to

***Corresponding Author:** Flynn J, Materials and Manufacturing Academy, Swansea University, Fabian Way, Skewen, Swansea SA1 8EN; Tel:+ 07443637510, Email: 827380@swansea.ac.uk

Received date: November 18, 2020; **Accepted date:** August 27, 2021**;** **Published date:** September 8, 2021

Citation: Flynn J (2021) A Machine Learning Classification to Identify Houses Suitable for Electric Vehicle Charging from Remotely Sensed Imagery. J Remote Sens GIS.Vol.10 .p147.

qualify for the scheme, customers who are the registered keeper, lessee or have primary access to an electric vehicle must have designated private off-street parking accessible at all times for Low Emission Vehicles (2019).

Even the most extensive mapping agencies in the UK such as Ordnance Survey (OS) and Open Street Maps (OSM) lack accurate data describing detailed features of individual residential buildings. While both datasets do include some building attribute data such as building type, active frontage, square footage, etc., some of this information can be difficult to obtain for various reasons. Firstly, obtaining data relating to street frontage using traditional survey methods is often both costly and time consuming, therefore this information is difficult to keep up to date Law et al. (2017). Secondly, in the late 1960's following advances in aerial surveying, drawing and printing techniques, as well as the introduction of digital technologies, OS began to manually digitalize their historical paper maps into what is now the OS Master Map. These techniques, particularly the use of aerial remote sensing, make it difficult to classify buildings or identify their frontal features without ancillary data Hussain et al. (2007). This is especially true for identifying areas suitable for EV charging points as it is difficult to determine whether a property has road access from an aerial view, and integral/semi-detached garages may be concealed. Thirdly, mapping agencies such as OSM that predominately rely on Volunteered Geographical Information (VGI) find that the accuracy and attention to detail of various measurements differs between areas depending on the skills and patience of the contributor Haklay (2010). Furthermore, the quality of VGI coverage drops considerably in suburban areas where there are fewer contributors Haklay (2010). These challenges make it difficult for companies to obtain data on residential areas suitable for EV charging points.

In this paper we overcome these challenges by combining data from OSM and Google Street View (GSV) to develop an end-to-end workflow to survey small to medium urban areas for residential properties suitable for EV charging. This paper also presents the first case of using deep CNN's to classify residential off-street parking from remotely sourced imagery. The paper is structured as follows. In section 2 we review the current literature of machine learning in remote sensing. Following the literature review, section 3 outlines the methodology to identify residential properties for EV charging. This methodology is presented as a 5 step, end-to-end workflow with each step described in detail within the section. The methods used to select, train, test and optimise the two CNN's that are utilised within this workflow are also discussed in the methodology. Section 4 includes the results of the CNN optimisation as well as a demonstration of the full workflow on two urban areas. Finally, sections 5 and 6 include the discussion and conclusion of our findings.

Related Work

Deep Learning in Remote Sensing

Over the past decade, Deep Convolutional Neural Networks (CNN's) have emerged as a powerful tool for remote sensing applications due to their superior performance in visual pattern recognition, object recognition image classi- fication and image segmentation Kampffmeyer et al. (2016); Zhang et al. (2018); Schmidhuber (2015). Based loosely on the perception mechanisms of the visual cortex, a CNN architecture consists of multiple layers of artificial neu- rons in a stacked arrangement to perform three main operations: convolution, non-linearity and pooling/subsampling Kampffmeyer et al. (2016); Schmidhuber (2015). For image classification, a series of multispectral images are fed into the first layer in the form of 2-D arrays. The following sequential layers are then represented as input and an output feature maps calculated by alternatively stacking convolutional and pooling layers. The final layer is a fully connected layer in which classification is performed. Alexnet, GoogLeNet, VGG, and ResNet are the most common pre-trained network architectures in the reviewed literature.

Alexnet is the most commonly used network for remote sensing and LULC tasks due to its high efficiency and performance Krizhevky et al. (2007) . Alexnet is an 8 layer CNN developed by Krizhevsky et. al in 2007 as part of the ImageNet LSVRC-2010 contest. Alexnet was the first network to incorporate the Rectified Linear Activation Function (ReLU) which marked a major algorithmic change that greatly improved the performance of feed-forward networks and permitted the development of very deep neural networks Goodfellow et al. (2016). The network features five convolutional layers and three max pooling layers that result in approximately 60 million parameters, making it the least complex network reviewed in this section. Due to this relatively low complexity, modified versions of AlexNet have been developed to run on embedded devices in real time Amato et al. (2017). In their 2017 paper, Amato et al. (2017) develop a CNN for parking lot occupancy detection based on AlexNet's architecture which was able to run in real time on a Rasbperry Pi 2 model B.

Christian et al. (2015) introduced the architectural concept of "Inception Modules" which was used to develop their GoogLeNet network which won the ILSVRC14 competition. By enabling multiple feature extractors to branch

from a single layer, the inception modules allow for filters of different resolutions to operate in parallel. With different sized filters at each layer, more accurate spatial information is retained while the number of parameters is significantly reduced when compared to similarly performing non-inception architectures. This means that despite its 22 layers, GoogleNet has only 4 million parameters, making it less-sensitive to over-fitting. Castelluccio et al. (2015) use GoogleNet and CaffeNet to perform LULC classification on aerial images. In this paper the focus is on how the method of training effects the final accuracy of the network. GoogleNet outperformed CaffeNet at all tasks and it was found that fine-tune transfer learning gave the best results compared to training the network from scratch or using feature vector.

The concept of a Residual Network (ResNet) was introduced in 2016 by He et al. (2016) as part of an attempt to overcome degradation when building very deep networks. Degradation refers to the tenancy for training error to suddenly increase as a networks depth is increased beyond the point where accuracy saturates. This issue is overcome by introducing feed-forward shortcut connections around multiple stacked layers within the network architecture. This simplifies the process of mapping the identity function onto the output of the skipped layers, meaning additional layers can be added without increasing the number of parameters or computational complexity. In Cai et al. (2018), a modified ResNet50 network was demonstrated to outperform existing unsupervised methods at measuring urban tree cover from GSV images.

In Kang et al. (2018), CNN's are compared in their ability to classify buildings from GSV and aerial images at several cities in North America. The paper looks at four CNN's: VGG16, AlexNet, ResNet18 and ResNet34. All 4 networks showed high levels of error when trying to distinguish between houses and garages. The paper identifies the source of this error being due to the high number of images of residential houses with an integral garage. VGG16 performed the best at the classification task followed by the ResNet variants and AlexNet. In this paper we do not consider VGG networks as they are too computationally expensive.

The pre-trained networks discussed in this section are all trained on ImageNet, a huge dataset containing millions of images and 1000 categories Li et al. (2009). When re-purposing these pre-trained networks for remote-sensing applications, a common challenge is the lack of publicly available training data Castelluccio et al. (2015). This issue is common across many different fields and solutions have been widely explored in the literature Yosinski et al. (2014); Donahue et al. (2014); Zhu et al. (2011). The most common method of adapting pre-trained networks is fine- tune transfer learning whereby the user only adapts a small number of high-level layers to the new task. This means that the network retains knowledge of lower level features that are widely transferable across tasks. Castelluccio et al. (2015) demonstrates how a networks trained using the fine tune approach outperforms the same network trained from scratch on the same dataset in both final accuracy as well as computation time.

Deep learning challenges and limitations

As yet, no attempts have been made to use CNN's to identify residential properties suitable for EV charging from streetscape imagery. However, the challenges of such a task are widely discussed in the literature. One major problem is occlusion. Residential off-street parking facilities are often relatively small features when compared to car parks, residential buildings and industrial buildings that most classifiers deal with. Such features can easily be obscured by structures, trees or vehicles in the foreground. This issue can be resolved by using multiple GSV images from different angles and positions although this may not always be possible Enr ́ıquez et al. (2017).

The second major problem involves the diversity of the class. Some structures, such as religious buildings, are easier to classify as they often contain unique, distinctive features that not associated with other classes e.g. spire or a dome. Private driveways vary greatly in size, shape, texture and their location relative to the associated property, making them difficult to distinguish from other features such as a front garden or pathway. Furthermore, while one would assume that garages could be easily distinguished by a large, metallic, square door, previous attempts to classify garages have proven difficult and are often misclassified as houses as the two often appear in the same image. This issue is highlighted in Kang et al. (2018) where four different CNN's demonstrate poor accuracy when attempting to identify garages with a significant portion misclassified as houses.

MATERIALS AND METHODS

This section introduces our method for remotely surveying urban areas to identify residential properties suitable for EV charging. The method, illustrated by the workflow diagram in Figure 1, is split up into five main steps. In steps 1 and 2, a MATLAB code is used to request the geographical and image data for the area to be surveyed. This image data is then filtered and sorted by a CNN in step 3 before being passed onto a second CNN in step 4 which identifies all images suitable for EV charging. Finally, in step 5 a mailing list of all locations identified as suitable for EV charging is presented as well as a map of their locations. These steps are described in further detail in this section, as well as the methods used to select, train, test and optimise the two CNN's that are utilised within this workflow.

Figure 1: Diagram showing the total workflow of the proposed method.

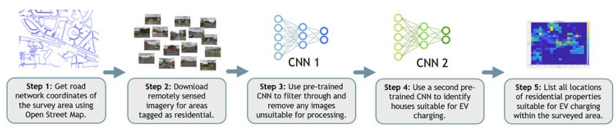

Training and Testing Data Acquisition

In order to train a CNN to recognize properties that are suitable for EV charging point installation, a bespoke dataset was created for both training and testing using GSV data collected from eight different UK towns and cities, as shown in Figure 2. These towns and cities are selected due to their varying lo- cal histories, population densities, and geographical locations, which in turn affects the age, density, and style of housing. This allows us to build a diverse dataset, despite its small size, thus maximising the accuracy of the overall system when testing in mul- tiple areas. Open Street Maps (OSM) provides open source geographical data and was used to acquire the geographic information for these urban areas Con- tributors (2019). At each location, described by rect- angular boundaries of longitude and latitude, data is exported in XML format before using the method de- veloped by Filippidis (2013) to extract the road net- work coordinates . For every coordinate on the road network four separate images were downloaded us- ing the Google Street View API at headings of 0, 90, 180 and 270 degrees, each with an image size of 600× 400 pixels. The pitch was kept constant at 0 de-grees. This same method is used to acquire both the training and testing data, although for the testing data the road network coordinates were interpolated to re- duce the distance between nodes and ensure full coverage of the road network.

Testing data was gathered from three locations: Oswestry, Petersfield and Birmingham

Figure 2: A map showing the locations of the towns and cities from which the training and testing data is sourced.

The Oswestry data is used for hyper-parameter optimisation and network selection, while the Petersfield and Birm- ingham datasets are used as case studies to demonstrate the application of the workflow to survey an urban area for homes suitable for EV charging.

Table 1: A breakdown of the image data used to retrain CNN 1.

Category	Blackpool	Peterborough	Swansea	Cwmbran	Colchester	Exmouth	Total
Car Parks	349	424	308	405	90	117	1,693
Trees and Foliage	318	1,658	573	311	348	402	3,610
Road Views	354	388	434	334	337	375	2,222
Residential Front View	1,251	937	947	1,019	972	1,033	6,159
Totals	2,272	3,407	2,262	2,069	1,747	1,927	13,68

Table 2: A breakdown of the image data used to retrain CNN 2.

Category	Blackpool	Peterborough	Swansea	Cwmbran	Colchester	Exmouth	Totals
EV Suitable	627	743	295	451	898	759	3,773
EV Unsuitable	542	379	411	94	322	415	2,732
Totals	1,169	1,122	706	545	1,220	1,174	6,505

Challenges of the data

Using OSM and GSV to source our data has several advantages, most importantly the fact that they are open source. However, there are certain characteristic of the data that present challenges for detailed surveying. OSM relies on crowd sourced data to develop and update maps. Despite quality assurance tools being in place, mistakes are not uncommon and information such as road names are often missing. OSM also lacks detailed data on individual properties such as plot boundary polygons, house numbers and names, as well as the presence of any external features such as garages and driveways. This, combined with the blurring of house numbers in GSV data, makes it difficult to determine the address for any specific property suitable for home EV charging without using auxiliary datasets.

Another issue with this method of data acquisition lies in the Google Street View API. At every point in the road network, 2 images need to be downloaded: an image directly facing the properties on either side of the road at 90 degrees and 270 degrees relative to the vehicles direction of travel. However, in section 3.1 it was stated that four images were downloaded at each coordinate with headings of 0, 90, 180, and 270. The reason for this is that the GSV headings are unreliable, and because panoramic images do not fit the input size of the most commonly used CNN's we must download all four angles to ensure a suitable image of the adjacent property is retrieved. The result is that roughly half of the image downloads contain angles unsuitable for processing, with views that look directly down the road making it difficult and often impossible for the network to identify the features of houses that identify them as suitable for EV charging.

Data Pre-Processing

A large proportion of images downloaded from any area are likely to contain images that are unsuitable for pro- cessing, particularly in areas with lower housing density. To overcome these challenges and remove any unsuitable images, a pre-processing step is included in the workflow to remove images of car parks, trees and foliage, unsuitable headings, as well as any duplicate images. With these im- ages removed this should only leave images of buildings taken at suitable headings that can be passed on for pro- cessing. This is achieved using a CNN, know henceforth as CNN 1, trained using the dataset shown in Table 1. CNN 1 was tested and optimised using a sample of the Oswestry dataset containing 200 randomly selected images in each of the 4 categories.

Figure 3: Results of the grid search to optimise the hyper-parameters for the three networks.

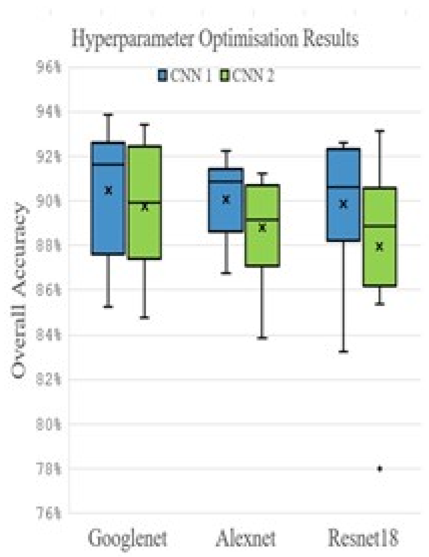

Network Selection

Two CNN's are used in the workflow. CNN 1 is used to filter through all image downloads from the target area and pass on any images of buildings taken at suitable headings

on to CNN 2 which is then used to identify all properties for EV charging. A breakdown of the image data used to train CNN 2 is presented in Table 2. For both networks, since the training datasets are not large enough to train a network from scratch, a fine-tune transfer learning approach is adopted. For image classification this approach makes use of existing networks that have already been pre-trained on large image datasets and therefore have already leant certain low-level features such as object edges, shapes, corners and intensity.

Based on related work identified in the literature, it was decided that for both CNN 1 and CNN 2 the performance of three different pre-trained deep CNN's is to be compared: Resnet-18, Googlenet, Alexnet. When comparing and optimising network performance, the same two hyper-parameters are optimised using a grid search approach by varying the mini batch size at 32, 64 and 128 while the initial learn rate is varied at 1e-3, 1e-4 and 1e-5. For all runs the learn rate drop period is set at 15 epochs with a drop factor of 0.1 and the L2 Regularization is set at 1e-5 to reduce over-fitting. All other hyper-parameters are set at the MATLAB defaults.

Data Post-Processing

For testing of the Petersifeld and Birmingham datasets, we modify the 'openstreetmap' interface code from Filip- pidis (2013) to request image data for every 20m of the road network. This ensures that we download all GSV image of the entire road network. When each image is downloaded from Open Street Maps additional properties associated with the requested coordinate are also downloaded and tagged to each image, such as road name and other address data if available. Once an image is identified by the network as suitable for EV charging these features are evaluated and returned to the user highlighting all suitable locations within the area as well as a mailing list of all roads ranked by the frequency of EV suitable properties.

Figure 4: Confusion matrices from the best performing networks following the grid search optimisation for CNN 1 (left) and CNN 2 (Right).

Experimental Results

Network Selection

In order to select suitable networks and hyper-parameters for both CNN 1 and CNN 2 a grid search approach was adopted. All runs were performed on a single Nvidia V100 PCI 16Gb GPU. The graph in Figure 3 shows how the three networks compared in their performance following the grid search experiment. For CNN 1, when tested on a sample of the Oswestry dataset containing 200 images per category, Googlenet achieved the highest overall accuracy of 93.9% when using a mini batch size of 32 and an initial learn rate of 0.001. More importantly, as shown by the confusion matrix in Figure 4 this network was 100% accurate at identifying images of Residential Front Views suitable for processing. For CNN 2, when tested on a sample of the Oswestry dataset containing 500 images per category, Googlenet achieved the highest overall accuracy of 93.7% when using the same mini batch size and initial learn rate of 32 and 0.001 respectively. The confusion matrix for the optimised CNN 2 network is shown in Figure 4. Given the results of this experiment, the Googlenet architecture was selected with a mini batch size of 32 and an intial learn rate of 0.001 to be used for both CNN 1 and CNN 2.

Testing the full Workflow

In order to demonstrate the entire workflow on a larger scale we are required to manually label every image in each survey area order to evaluate the performance of the system. To make this feasible all of the images within the test areas are split into 3 categories: 'EV Suitable', 'Unsuitable (Residential Front View)', and 'Unsuitable (Other)'. These three categories allow us to test the performance of both networks. CNN 1 should filter out all images in the 'Unsuitable (Other)' category which includes any images of car parks, trees and foliage and unsuitable headings. As before, images of buildings are then passed onto CNN 2 which analyses these remaining images to identify those suitable for EV charging point installation before the results are then plotted on a the map.

Figure 5: The resultant confusion matrices for CNN 1 and CNN 2 on the Petersfield testing data.

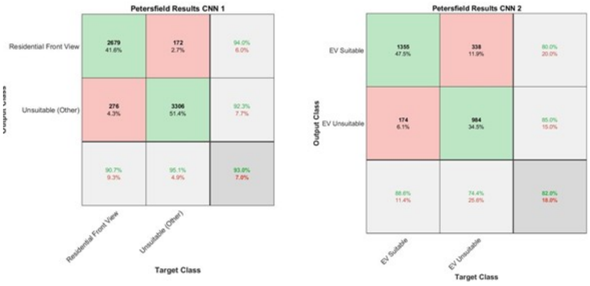

Test Area 1 - Petersfield

Data Acquisition

Petersfield is a rural town in Hampshire, South England with a population of approximately 15,000 people accord- ing to the most recent census data ONS (2011). Image data for the Petersfield test area was requested at 30m intervals along all roads within the boundary that were tagged as 'Residential', resulting in a total of 2,092 locations along the road network shown on Figure 6. After downloading images at each of the four headings and removing any duplicate images this gives a total of 6,433 images in the Petersfield testing dataset. These images are then manually labelled into the 3 test categories also shown in Figure 6.

For the Petersfield test area CNN 1 is 90.7% accurate at identifying images of 'Residential Front Views' suitable for processing, as shown in Figure 5. Following the CNN 1 step, 2,851 images are passed onto CNN 2 for processing. CNN 2 then identifies 1,693 images as suitable for EV charging, achieving an 88.6% accuracy, as shown in Figure

Overall the system recognises 1,355 of the total 1,651 images suitable for EV charging. Within the Petersfield test area 132 streets are identified with properties suitable for EV charging. Figure 7 shows all suitable properties plotted on a map of the road network. Table 7 shows the roads returned by the system as having the highest number of residential properties suitable for EV charging point installation.

Figure 6: The locations of all downloaded images from the Petersfield test area (Left) and a breakdown of the image categories (Right).

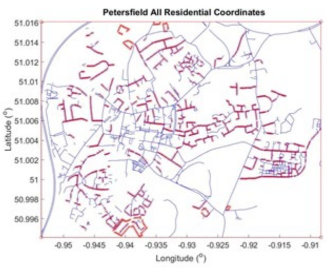

Category	Petersfiel d
EV Suitable	1,651
Unsuitable (Residential Front View) Unsuitable (Other)	3,478 1,304

Figure 7: A plot of all coordinates within the Petersfield test area identified as suitable for EV charging (Left). A list of the top 10 roads most suitable roads for EV charging (Right).

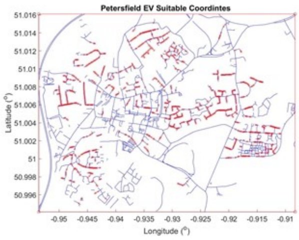

Road Address	Frequency
Woodbury Avenue, Petersfield, GU32	70
Stafford Road, Petersfield, GU32	58
Monks Orchard, Petersfield, GU32	56
Hanger Way, Petersfield, GU31	48
Tilmore Gardens, Petersfield, GU32	45
Marden Way, Petersfield, GU31	37
Rival Moor Road, Petersfield, GU31	36
Durford Road, Petersfield, GU31	36
Cranford Road, Petersfield, GU32	33
Moggs Mead, Petersfield, GU31	33

Figure 8: The resultant confusion matrices for CNN 1 and CNN 2 on the Birmingham testing data.

Test Area 2

Data Acquisition

For the second test area we collected GSV data from a residential area roughly 4 km east of Birmingham city centre, the road network for which is shown in Figure 9. A total of 28,728 coordinates along the road network were downloaded, which resulted in 14,766 in this dataset after removing duplicate images. These images are then manually labelled into the same 3 test categories as before, as shown in Figure 9.

Figure 9: The locations of all downloaded images from the Birmingham test area (Left) and a breakdown of the image categories (Right).

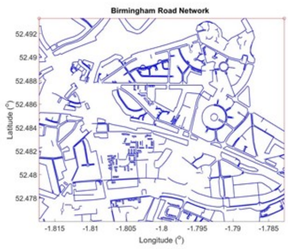

Category	Birmingha m
EV Suitable	5,306
Unsuitable (Residential Front	7,538

Figure 10: A heat map showing the areas with the highest density of EV suitable properties in the Birmingham test area as well as the individual locations of EV suitable properties plotted in red (Left). A list of the top 10 roads most suitable roads for EV charging (Right).

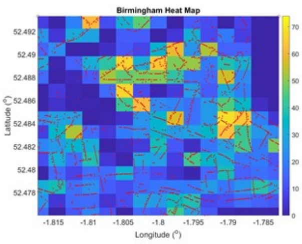

Road Address	Frequency
Stud Lane, Stechford, B33	44
Peplow Road, Stechford, B33	39
Manor Gardens, Yardley, B33	23
Kirton Grove, Stechford, B33	14
Tanfield Road, Stechford, B33	14
Shelly Croft, Stechford, B33	14
Croxton Grove, Stechford, B33	10
Kington Way, Yardley, B33	9
Beswick Grove, Stechford, B33	9
Alcombe Grove, Yardley, B33	8

RESULTS

Once the images are downloaded they are passed through CNN 1 for pre-processing. For the Birmingham test area CNN 1 is 95.1% accurate at identifying images of Residential Front Views suitable for processing, as shown in Figure 8. Following the CNN 1 step, 7,408 images are passed onto CNN 2 for processing which then identifies 5,376 images as suitable for EV charging, achieving an 89.3% accuracy at identifying properties suitable for EV charging, as shown in Figure 8. Overall the system recognises 4,617 of the total 5,306 images suitable for EV charging. Figure 10 shows all EV suitable properties plotted on a map of the road network as well as the roads identified by the system as having the highest number of residential properties suitable for EV charging point installation.

DISCUSSION

The system is shown to be highly accurate at identifying EV suitable properties in both the Petersfield and the Birmingham test areas achieving 87.2% and 89.3% respectively. One of the issues with this current method is that a significant proportion of the original downloaded images are wasted. In the Petersfield example, despite only downloading images from areas tagged as residential, only 26% of the images contained properties suitable for EV

charging. We are currently not charged for the small number of images we are downloading (<24,000 per month),

however, when this method is scaled up to be used to survey local authority areas, Google API applies a charge of approximately £0.0007 per image. This challenge will have to be addressed in future research by either using more precise Ordnance Survey data in place of the current OSM data, or by using aerial imagery as an auxiliary dataset. Another challenge of this method is that using MATLAB to process the data makes it difficult to retrieve certain features at some locations. Approximately, 94% of the road names are unable to be retrieved using our current method. This explains the low frequency of suitable houses identified for each road as shown in Figures 7 and 10. This challenge could be addressed by further modifying the Filippidis (2013) code that was used to import the .osm files to improve feature extraction, or Ordnance Survey data could be used in conjunction with GIS software. Using Ordnance Survey data would allow the user to get addresses at an individual property level.

CONCLUSION

In this paper we demonstrate that surveying external building characteristics of small to medium urban areas is possible using freely available open source data. A system is developed capable of identifying residential properties suitable for EV charging. To achieve this we introduce a novel end-to-end workflow that utilises Open Street Mapand Google Street View data which is processed by two separate CNN's. To select an appropriate architecture for these tasks the performance of 3 pre-trained networks is compared by means of a grid search experiment. Googlenet outperformed AlexNet and ResNet18 at both image classification tasks achiving accuracies of over 90%. To validate the performance of the entire workflow we survey the entire road network of two contrasting urban areas for properties suitable for EV charging, returning the addresses of the most suitable roads. In both cases the both CNN's in the workflow achieve accuracies of around 90%. This proof of concept demonstrates a promising new application of deep learning techniques in the field of remote sensing and urban planning.

REFERENCES

1. Amato G, Carrara F, Falchi F, Gennaro C, Meghini C, Vairo C, et al. Deep learning for decentralized parking lot oc- cupancy detection. Expert Systems with Applications. (2017);72:327-334.
2. Castelluccio M, Poggi, G, Sansone C, Verdoliva L. Land Use Classification in Remote Sensing Images by Convolutional Neural Networks (2011);1-11.
3. Christian S, Wei L, Yangqing J. Going Deeper with Convolutions. Contributors, O.2019.
4. Enr´ıquez F, Soria LM, Antonio´ Alvarez-Garc´ıa J, Velasco F, De ´niz, O. Existing Approaches to Smart Parking: An Overview, in: Smart-CT 2017: Smart Cities. 2017;63-74.
5. Filippidis I. Openstreetmap: Interface to OpenStreetMap. 2013.
6. Huang B, Zhao B, Song Y. Urban land-use mapping using a deep convolutional neural network with high spatial resolution multi-spectral remote sensing imagery. Remote Sensing of Environment. 2018;214:73-86.
7. Hussain M, Davies C, Barr R. Classifying Buildings Automatically : A Methodology. GISRUK 2007 Geographical Information Science Research Conference. 343–347.
8. Donahue J, Jia Y, Vinyals O, Hoffman J, Zhang N, Tzeng E, Darrell T, et al. DeCAF: A deep convolutional activation feature for generic visual recognition. 31st International Conference on Machine Learning, ICML. 2014;2:988-996.
9. Goodfellow I, Bengio Y, Courville A. Deep learning. The MIT Press, Cambridge, Massachusetts.
10. Haklay M. How good is volunteered geographical information? A comparative study of OpenStreetMap and ordnance survey datasets. Environment and Planning B: Planning and Design. 2017;37:682-703.
11. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 2016;770-778.