Deep Learning Assisted Classification of Wind Exposure Categories

Pallava Kodali
CPP Wind Engineering Consultants
pkodali@cppwind.com

Christian Rohr
CPP Wind Engineering Consultants
crohr@cppwind.com

ABSTRACT

Accurately determining Wind Exposure Categories (AS/NZS 1170.2:2021) is a crucial step in wind engineering analysis. For each wind direction considered, these categories characterize the roughness upwind of a project site, significantly influencing wind loading and environmental conditions on and around structures. Traditionally, this process has been manual, relying on the visual inspection of satellite orthoimagery to assign categories based on ground surface roughness. This method encompasses extensive areas, often extending from 30 to 100 kilometres outward from the site in multiple directions (up to 36) and is applied to patches of land ranging from 6 to 25 square kilometres. However, this approach is subjective allowing inconsistencies. The interpretation of satellite imagery in terms of wind exposure can vary widely among wind engineers leading to potential discrepancies. To address these limitations, this paper introduces a novel, fully automated method that leverages deep learning and convolutional neural networks (CNNs) to classify satellite orthoimagery into wind exposure categories. The methodology harnesses the capabilities of pre-trained CNNs, such as ResNet, which have been exposed to millions of images across thousands of categories. Fine-tuning these models on domain-specific training images has generally been shown to surpass the performance of custom CNNs trained from scratch. Initial validation shows the trained model is capable of correctly detecting 99.33% of Wind Exposure Categories, although development is ongoing.

INTRODUCTION

Drawing inspiration from previous research (Helber et al., 2017) that demonstrated the successful application of a fine-tuned ResNet-50 model on the EuroSAT RGB image dataset (~ 10,000 images) for land use classification, this study adopts a similar approach. However, it employs the more complex ResNet-152 model, chosen for its enhanced ability to discern deeper and subtler features in satellite imagery, despite its higher computational demands. This research utilizes the popular UC Merced Land Use dataset for its smaller, more manageable size and its higher spatial resolution, comprising 2100 images across 21 categories. In the preliminary test, these categories were reclassified into 4 simple levels of roughness (very rough, rough, smooth, and very smooth) to align with most of the wind exposure assessment needs. These upwind exposures are determined by the model's distinction of natural topography, vegetation, and constructed facilities. A patch of land never seen by the trained ML model is randomly selected and recoloured by one of these 4 levels of roughness. By leveraging advanced deep learning techniques and fine-tuning strategies elaborated below, the proof of concept presented here paves a possible path for automating the classification of wind exposure categories, towards a more consistent, efficient, and objective method for assessing wind exposure when studying the roughness conditions around and safety/reliability of structures under varying wind conditions.

DATASET PARAMETERS

Released on the 28th of October 2010, The UC Merced Land Use Dataset has been widely used for research purposes and benchmarking land use classification models. This is a 21-class land use image dataset with 100 images for each of the following classes:

1. Agricultural
2. Airplane
3. Baseball Diamond
4. Beach
5. Buildings
6. Chaparral
7. Dense Residential
8. Forest
9. Freeway
10. Golf Course
11. Harbor
12. Intersection
13. Medium Residential
14. Mobile Home Park
15. Overpass
16. Parking Lot
17. River
18. Runway
19. Sparse Residential
20. Storage Tanks
21. Tennis Court

Each image measures 256x256 pixels and the pixel resolution of this public domain imagery is 1 foot (0.3 m). Yi Yang et al. (2010) manually extracted these images from the USGS National Map Urban Area Imagery collection for various urban areas around the country.

TRAINING STRATEGY INSPIRATION

Advantages of Pre-trained CNNs

Pre-trained CNNs suit a complex image classification task such as wind exposure category assessment for the following benefits they provide:

- As observed by Marmanis et al. (2016), instead of starting from scratch, the weights learned by the pre-trained CNNs after seeing millions of images can serve as a starting point. Later, these weights can be fine-tuned for the specific image classification task. This allows us to leverage the previous knowledge and modelling and avoid pitfalls such as overfitting especially when the available data are limited.
- Having been exposed to a large number of images across diverse categories, pre-trained CNNs are able to resolve both high-level (prominent) and low-level (subtle) features in images. Zhang et al. (2016) have shown that they generalise better and are domain agnostic, which means they can be fine-tuned to image datasets in any field and perform well.

Training Strategy

Helber at al. (2017) proposed a step-by-step method to take advantage of pre-trained CNNs:

1. Start with a ResNet-50 CNN model which has been pre-trained on the ImageNet Large Scale Visual Recognition Challenge 2012 (ILSVRC2012) image dataset by Russakovsky et al. (2015) which has 1.28 million training images, 50,000 validation images, and 150,000 test images. These images are categorised into 1,000 categories.
2. Freeze all layers of the ResNet-50 model.
3. Unfreeze the final fully connected layer and allow it to recategorize its learning from 1,000 categories to EuroSAT’s 10 categories across 27,000 images (uneven distribution, some categories have more images than others).
4. In phase 1 of training, set this final fully connected layer to have a high learning rate of 0.01 so that the model pulls information from the EuroSAT dataset to make aggressive adjustments to its learned weights.
5. In phase 2 of training, unfreeze all layers in the network and reduce the global learning rate to be between 0.001 and 0.0001 so that the model pulls information from the EuroSAT dataset to make modest adjustments to its learning. This protects what the model has learned in phase 1 while also making fine adjustments to achieve a more sophisticated understanding of the data.

Results

This two-phase strategy has allowed ResNet-50 to achieve up to 2% higher classification accuracy in an 80/20 split setting over the single-phase strategy where the entire network was unfrozen and trained. This is shown in Table 1 below:

Table 1: Benchmarked classification accuracy (%) of ResNet-50 on different splits

	10/90	20/80	30/70	40/60	50/50	60/40	70/30	80/20	90/10
Single Phase	75.06	88.53	93.75	94.01	94.45	95.26	95.32	96.43	96.37
Two Phase	-	-	-	-	-	-	-	98.57	-

IMPROVED TRAINING STRATEGY

Pros and Cons of ResNet-152 over ResNet-50

ResNet-152 has, as the name suggests, 152 neural layers in it, which is 102 more than that of ResNet-50. This allows ResNet-152 to learn more complex patterns and subtler features that are prevalent in high resolution satellite imagery like UC Merced Land Use Dataset. When fully connected, these 152 layers also allow the model to better understand feature clusters and complex shape hierarchies. While this is all true, ResNet-152 is more computationally intensive and uses more memory. If not careful, ResNet-152 is at a higher risk for overfitting. However, the strategy explained below mitigates this effect.

Training Strategy

A slightly more sophisticated training strategy is proposed:

1. Start with a ResNet-152 CNN model, pre-trained on the ILSVRC2012 image dataset.
2. Freeze all layers of the ResNet-152 model.
3. Unfreeze the final fully connected layer and allow it to recategorize its learning from 1,000 categories to UC Merced's 21 categories across 2,100 images (even distribution, all categories have the same number of images).
4. In phase 1 of training, a third of the total number of training iterations are completed based on the following mechanism while also automatically saving the best model and preventing overwriting at each training iteration.
   a. Specify that the optimizer should update the parameters of the model's final fully connected layer in the CNN.
   b. Set the initial learning rate to be high. The learning rate as mentioned before is a parameter that controls how much to change the model in response to the estimated error each time the UC Merced data is studied when training.
   c. Implement a weight decay that adds a penalty to the loss function for categories that receive too much emphasis to due large, accumulated weights in the model. This can help discourage bias towards any one wind exposure category.
   d. Implement a scheduler that reduces the learning rate when a metric has stopped improving, which is useful for fine-tuning a model in the later stages of training.
   e. Tell the scheduler to reduce the learning rate when the model stops getting better at understanding the difference between wind exposure categories, i.e., when the model stagnates.

f. Implement a factor by which the learning rate will be reduced. E.g., if the learning rate is 0.5 and the factor is 0.5, the new learning rate will be 0.25.

g. Implement some patience into the model to wait before reducing the learning rate if no improvement is seen. This gives the model more nuance in that it allows enough training iterations before significant changes are made.

h. Set the lower bound on the learning rate. No matter how many reductions occur, the learning rate will not go below this value.

5. In phase 2 of training, all layers in the network are unfrozen and the remaining two thirds of the total number of training iterations are completed based on the following mechanism while also automatically saving the best model and preventing overwriting at each training iteration.

   a. Specify that the optimizer should update all parameters across the entire model, not just the final fully connected layer like in phase 1.

   b. Set the initial learning rate much lower than phase 1, which is intentionally quite small because lower learning rate can lead to slower but more stable convergence, reducing the risk of overshooting adjustments made after studying the images in each training iteration.

   c. Implement a very small penalty (much lower than phase 1) term that helps prevent bias towards wind exposure categories with large weights, encouraging more generalised behaviour on unseen data.

   d. As in phase 1.

   e. As in phase 1.

   f. As in phase 1 but much smaller than phase 1 to allow extremely fine adjustments to the learning sensitivity.

   g. As in phase 1 but with twice or thrice the patience before learning rate was reduced.

   h. As in phase 1 but much lower than phase 1. This prevents the learning rate from becoming so small that the model no longer effectively learns.

Results

This improved two-phase fine-tuning strategy has allowed ResNet-152 to achieve an average of 99.33% classification accuracy across UC Merced Land Use Dataset's 21 categories in an 80/20 split setting. The percentages below show how often the model has identified images in each category correctly.

Agricultural: 100%	Forest: 100%	Overpass: 100%
Airplane: 100%	Freeway: 98%	Parking Lot: 100%
Baseball Diamond: 100%	Golf Course: 100%	River: 100%
Beach: 100%	Harbor: 100%	Runway: 100%
Buildings: 100%	Intersection: 99%	Sparse Residential: 98%
Chaparral: 100%	Medium Residential: 95%	Storage Tanks: 99%
Dense Residential: 98%	Mobile Home Park: 100%	Tennis Court: 99%

TEST & DISCUSSION

The model trained using the improved strategy can distinguish between the above 21 categories 99.33% of the time. However, from a wind exposure perspective not all land use categories above warrant their own wind exposure category. Therefore, the above 21 classes are reclassed into the following 4 roughness classes as shown in table 2 below.

Table 2: 21 Land use categories reclassed into 4 Levels of wind exposure roughness

Very Smooth	Smooth	Rough	Very Rough
Tennis Court Golf Course Baseball Diamond Chaparral Harbour Beach River	Runway Freeway Intersection Parking lot	Sparse Residential Storage Tanks Airplane Overpass Agricultural	Buildings Dense Residential Medium Residential Mobile Home Park Forest

These categories were reclassified into 4 levels of roughness to align with most of the wind exposure assessment needs. Figure 1 below illustrates how every 6000 square meter patch of land in a randomly selected and untrained 100 square kilometre area in Austin, TX, USA, is recoloured by the ML model by one of 4 levels of roughness.

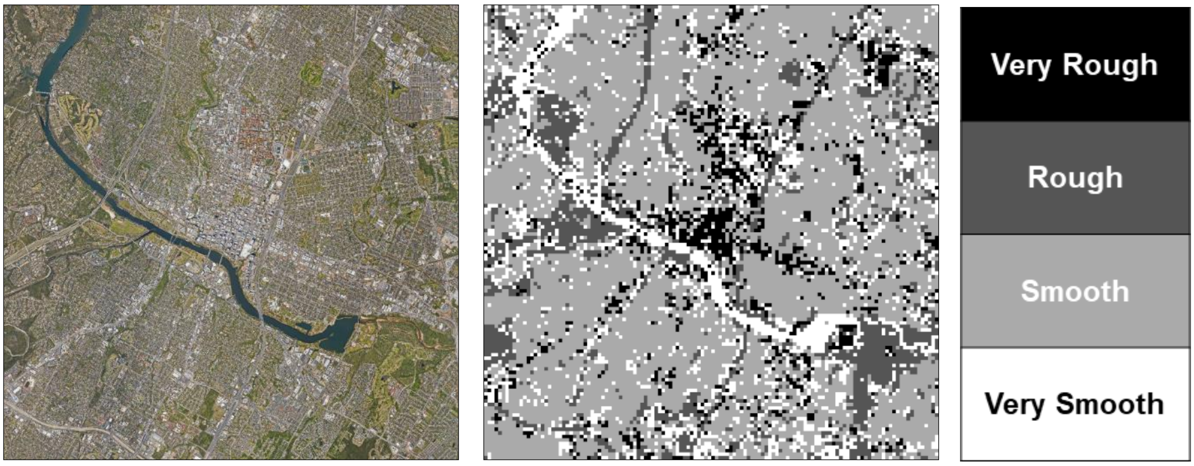

Figure 1: ML model recolours previously unseen area (left) by 4 levels of roughness (right)

As seen in the image above, the river is appropriately given a ‘Very Smooth’. Open areas and green belts such as parks, parking lots, and roads were given a ‘Smooth’ rating. Forested areas, sparse to medium residential suburbs, high overpasses, and crop areas were given a ‘Rough’ rating. And finally, dense suburbs and commercial precincts were given a ‘Very Rough’ rating. While not perfect, a model that was trained on just 2,100 images understands the distinctions between 21 different types of land use imagery, which performed well at distinguishing 4 broad roughness categories. With synthetic inputs and custom labelling to create more robust datasets in the future, the model can eventually learn to assess 8 exposure categories in line with wind engineering studies as shown in table 3 below.

Table 3: Land use & terrain type with respect to roughness length (CPP)

Description	Roughness Length (m)
Ocean | River | Lake	0.002
Runway | Open Country	0.02
Open fields with some obstructions | Scattered buildings with open land	0.15
Suburbia	0.2
Deciduous trees | Residential	0.35
Evergreen Trees | High Density Residential (e.g. LA) | Multi-story Commercial Areas	0.5
Campuses (4-5+ stories)	0.75
Centres of Cities	0.75 - 1.5
Downtown Large City (NY)	2.0

The category assessment model described above is one part of the larger whole. When it becomes capable of discerning 8 exposure categories as shown in table 3 above, the system will be further developed to trace terrain changes to determine fetch lengths and compute overall roughness in each fetch based on all terrain categories and their distributions observed in that fetch. Post-processing of the categorized imagery will act as an important intermediary, with terrain change spatial resolution to decrease with distance from the project location.

An example of the manual process is shown below in Figure 2. This can be both subjective and difficult to check. Automation of the initial category assignments will enable engineering effort to be spent on interpretation of the computed fetches and checking by minimising the time-consuming manual categorization. The goal is to provide better inputs to the ESDU 01008 Computer program for wind speeds and turbulence properties: flat or hilly sites in terrain with roughness changes (ESDU, 2010).

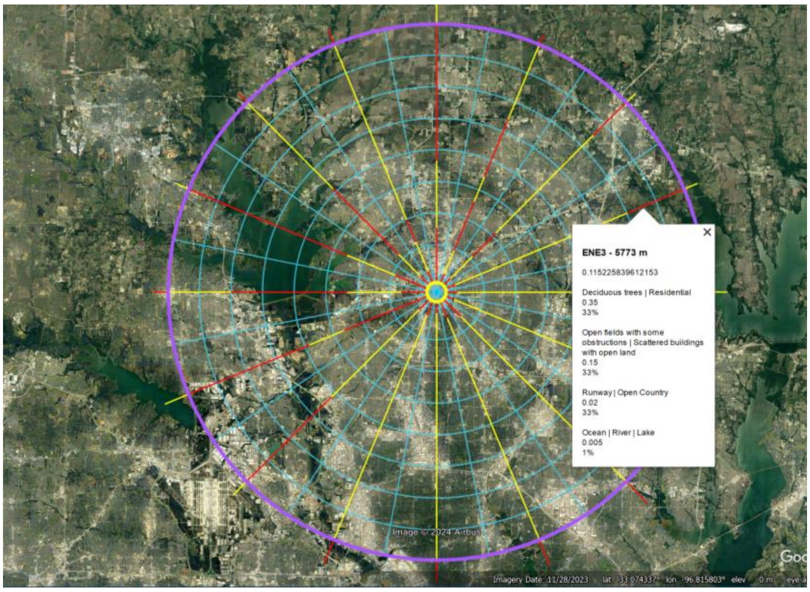

Figure 2: Manually traced major terrain change fetches, fetch lengths, and average exposures

SUMMARY

A new machine learning training strategy was developed to pave a potential path to automate wind exposure category identification just using satellite imagery. UC Merced Land Use Dataset was used to train a model and a test was performed where a randomly selected 100 km^2 land previously unseen by the model was recoloured based on 4 levels of roughness (very smooth, smooth, rough, and very rough). The model is demonstrated to correctly classify 99.33% of the terrain.

REFERENCES

[1] Yi Yang and Shawn Newsam, "Bag-Of-Visual-Words and Spatial Extensions for Land-Use Classification," ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM GIS), 2010.

[2] D. Marmanis, M. Datcu, T. Esch, and U. Stilla. Deep learning earth observation classification using imagenet pretrained networks. IEEE Geoscience and Remote Sensing Letters, 13(1):105–109, 2016.

[3] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016.

[4] Helber, Patrick & Bischke, Benjamin & Dengel, Andreas & Borth, Damian. (2017). EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. 10.1109/JSTARS.2019.2918242.

[5] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV), 115(3):211–252, 2015.

[6] ESDU 01008 (2010) "Computer program for wind speeds and turbulence properties: flat or hilly sites in terrain with roughness changes," Engineering Sciences Data Unit, London, UK

ACKNOWLEDGEMENTS

The UC Merced Land Use Dataset is based upon work supported by the National Science Foundation under Grant No. 0917069.