

# Accurate Detection of Illegal Dumping Sites Using High Resolution Aerial Photography and Deep Learning

Chirag Padubidri  
CYENS Center of Excellence  
Nicosia, Cyprus  
c.padubidri@cyens.org.cy

Andreas Kamilaris  
CYENS Center of Excellence,  
Nicosia, Cyprus  
Pervasive Systems, University of Twente

Savvas Karatsiolis  
CYENS Center of Excellence  
Nicosia, Cyprus

**Abstract**—Urban waste impacts human and environmental health. Waste management has become one of the major challenges faced by local governing authorities. Illegal dumping has become an important problem in many cities around the world. Effective and fast detection of illegal dumping sites could be a useful tool for the local authorities to manage urban waste and keep their administrative zones clean. Remote sensing based on satellite imagery or aerial photography is a key technology for dumping management, aiming at locating illegal waste sites and monitoring the required actions after the detection.

This study focuses on developing a method for detection and reporting illegal dumping sites from high-resolution airborne images based on deep learning (DL). Due to data unavailability for training a DL model, we use synthetic images. The trained model is evaluated based on a real-world dataset containing images from the city of Houston, USA. The results show that the proposed method solves the problem with high precision and constitutes a useful tool as part of a complete solution targeting dumping management by authorities.

**Index Terms**—Aerial photography, Deep learning, Dump Site, Waste Detection

## I. INTRODUCTION

Today, the majority of human population tends to migrate to urban areas. This trend results in the formation of massive cities of increased needs for physical resources (e.g. food, water, energy, materials). As a result, large quantities of waste are being released in the urban environment. According to Eurostat [1], waste produced by municipalities in EU-27 countries increased by 7.5% between 1995 and 2019. By 2050, daily per capita waste generation in high-income countries is predicted to rise by 19%, whereas it is expected to be increased by 40% or more in low- and middle-income countries [2]. It is evident that an increase in the overall waste generation is directly correlated with an increase in illegal dumping [3]. Illegal dumping increase may result in serious health problems. According to a study [4], illegal open waste dumping causes significantly more respiratory and eye illnesses. Moreover, trash disposal sites serve as a breeding ground for insects such as mosquitoes and flies, as well as for disease-carrying animals such as rats, skunks, and opossums which can spread life-threatening illnesses to nearby residents [5], [6].

Waste management, from consumption to disposal, is a significant challenge faced by governmental authorities around the world. Various smart waste management systems, such as community-based ones [7], [8], IoT-based [9], [10] and volunteer-mobile-app-based reporting systems [11] have enabled governmental and local authorities to improve the management and the treatment of the growing volume of waste generated, especially in big cities. However, less attention has been paid on unstructured waste (illegal dumping), which seems to be an open problem around the world to date [3]. There is a need for automatically detecting these sources of waste as fast as possible, for reasons of hygiene and environmental protection.

The current availability of high-resolution imagery from satellites/airborne (i.e. unmanned aerial vehicles - UAV, aeroplanes), together with the advancements in computer vision (CV), especially the successful use of Deep learning (DL) in CV applications, provides new opportunities for better monitoring of illegal dumping. Dumping site detection is generally a challenging problem in CV/DL, because the datasets openly available for training DL models to classify landscapes do not contain sufficient images of dumping sites, being heavily biased towards other non-dumping site classes (e.g. urban infrastructures, buildings and houses, forests, wetlands, pasture land, agricultural parcels, water bodies, etc.). This fact has motivated the authors to employ synthetic data to enrich existing datasets of remote sensing-based imagery. The proposed method applies an iterative process to generate synthetic data. In each iteration, new synthetic images were generated to cover diverse data distribution. The results show that the proposed method is effective and suitable for dealing with problems having highly imbalanced or limited data.

## II. RELATED WORK

Related work spans two different fields: a) detection and mapping of illegal dumping, and b) synthetic data generation.

### A. *Illegal dumping sites' monitoring and mapping*

A few studies have applied automated classification techniques for identifying and mapping illegal waste disposal from

aerial photography. These studies used either the direct spectral signature of waste itself or indirect spatial signatures (stressed vegetation near the waste dumping). Specifically, in [12], the authors proposed waste dumping location detection in landfills by using the multi-temporal Landsat thermal images. They harness thermal remote sensing i.e. ten years Landsat images) to measure the land surface temperature (LST), which aids in outlining the waste dumping regions within a landfill. By combining the multi temporal LST contours, a probability map was created to indicate the possible location of waste dumping within the studied landfill. The results derived during the summer and winter seasons both yielded an overall accuracy of 72%. The proposed solution in [12] is more suitable for mapping landfill, which emits heat during the aerobic and anaerobic phases.

Other studies used photographs to identify and map illegal waste disposal sites [13]–[15]. For example, Dabholkar et al. [15] designed a smart illegal dumping detector, which uses deep learning (DL) to recognize various types of frequently dumped wastes in images taken from streets. The authors trained a DL model with a few images showing frequently dumped waste from hundreds of images denoting illegal dumping provided by the city of San Jose. An edge computing station was installed, which run the DL model for captured images of individual dumping hot spots, sending the images to the server only when an image contained frequently dumped wastes

Finally, in [14], an automatic solution for the detection of clandestine waste dumps was proposed, using unmanned aerial vehicle (UAV) images from the Saint Louis area of Senegal, West Africa. The authors use very high-resolution UAV images (on the order of a few centimetres). The proposed solution used DL object detection for detecting waste dumping. The results showed that the model recognizes well the areas concerned, but presents difficulties on some areas lacking clear ground truths.

### B. Synthetic data generation

An important challenge and major limitation faced by the research community in this application domain, is the shortcoming of aerial imagery datasets containing dumping sites. The authors have also faced this challenge, since most of the publicly available datasets found involving high-resolution satellite imagery were heavily biased towards other classes (Non-dumping sites). Manual inspection of satellite photos and annotation of dumping sites is difficult and costly (i.e. large amounts of mostly non-dumping site data need to be carefully checked since dumping sites are scarce in the datasets).

The lack of data can be dealt with the use of synthetic data and heavy data augmentation, i.e. either producing artificial data from scratch or using advanced data manipulation techniques to produce novel and diverse training examples. In [16], the authors trained a DL model with synthetic images and then tested the model with real-world data from aerial images produced by unmanned aerial vehicles (UAV), focusing on

two different applications: forest fire detection and buildings' counting.

One common approach for generating synthetic data is to harness software development kits for gaming and scene generation. For example, UnrealCV [17] is an open-source plugin for the popular game engine Unreal Engine 4 that provides features that allow to set the camera location and the field of view, to set objects in a scene together with their positions, to set lighting parameters, to modify properties of objects such as material, etc. As another example, Blender<sup>1</sup> is a free open-source 3D creation suite that supports 3D pipeline—modelling, animation, simulation, rendering, compositing and motion tracking, video editing and 2D animation pipeline.

A Blender-based synthetic scene generator for recognizing objects inside a refrigerator was proposed in [18], showing improved results with a fully convolutional version of GoogLeNet [19], adapted for object detection. Similarly, the work in [20], [21] had successfully created realistic datasets using the Grand Theft Auto V video game. The authors claim that by capturing the communication between the game and the graphics hardware, they were able to cut the labelling costs by several orders of magnitude.

The novelty of our work lies in the application of synthetic images in classification problem, which has important differences than the technique discussed in related work [14], [15]. The results show that the proposed method solves the problem with good precision with minimum annotation. To the best of our knowledge, no other work has focused yet on data generation approaches for applications involving aerial photography or satellite imagery and dumping site classification.

## III. METHODOLOGY

An illustration of the proposed methodology for dumping site detection is shown in Figure 1, depicted as a system pipeline. This pipeline consists of a classification DL model which takes geo-referenced cropped images of size  $256 \times 256$  as input. The model predicts the class of the image, i.e. either dumping site or not. For each image used as input where a dumping class has been predicted, the image coordinates are converted into geo-coordinates (latitude and longitude). These geo-coordinates are then transformed to a GIS layer, which points the exact locations of detected dumping in the satellite images. This process is explained in detail in Section III-D. The GIS layer could then be used as a web or mobile application, a visualization data story or a visually-enhanced report, helping the authorities take actions, e.g. to inform the waste collection vehicles to clean the waste dumping and also to monitor the cleaning process.

### A. Dataset description

To demonstrate the proof of concept of the proposed approach, the 2018 IEEE GRSS Data Fusion Challenge - Urban land use and land cover classification dataset was used [22].

<sup>1</sup><https://www.blender.org/>

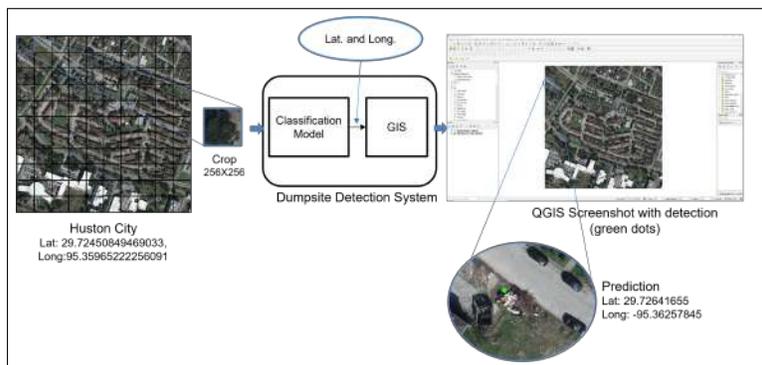


Fig. 1. Block Diagram of the proposed methodology for dumping site detection.

The multi-resolution and multi-modal optical remote-sensing dataset was released by the Hyperspectral Image Analysis Laboratory and the National Center for Airborne Laser Mapping (NCALM) of the University of Houston. In this paper, we use the high resolution RGB of the dataset. The dataset consists of 14  $11920 \times 12020$  resolution images which span between 29.7271 N to 29.7225 N and -95.3627 W to -95.3195 W (WGS-84) covering urban and semi urban region of the city of Houston.

### B. Data pre-processing

During data preparation stage, we extracted image patches for training and testing using sliding window-based cropping with zero stride, resulting in 29,000 image patches. The image patches were visually examined and classified into two classes, dumping and background class. Using this method, we obtained 176 image patches with garbage dumping, which serve as the ground truth information during the evaluation process (see Section IV). Out of a remaining 28,824 ( $29000 - 176$ ) image patches, 18,000 were randomly picked from background class for the experiment. 2000 image patch was used for synthetic data generation and 16,000 as Non-dump class image

1) *Synthetic data generation*:: Blender, an open-source software for creating 3D environments, was used to generate synthetic data. Blender allows the user to write scripts and add-ons in Python to automate the data generation process. The freely available dumping 3D objects [23], [24], provided by Blender, were used in the proposed work. A 2000 images were selected from background class and manually annotated with dots where possible dumping pile may occur. Using the Blender's Python modules, the dumping 3D objects were placed on the images based on the location determined in the previous step, generating incidents of *dumping* classes to properly train the DL classification models involved (Section III-C). The background class images were only used as background for generating synthetic images, they were not used in the training or the validation set. We note that the synthetic dumping sites contained a variety of 3D objects placed in the piles randomly.

Data generation was based on a positive feedback loop, i.e. an iterative process was followed where a baseline model

(Section III-C1) was trained at each iteration and then the results were examined. In case some dumping sites were not detected by the DL baseline model, due to different materials disposed which were not covered adequately by the Blender's 3D objects, then more synthetic data was generated in order to better approximate the data distribution of dumping sites, piles and materials. The data generation iteration was stopped based on the baseline model performance (here, the iterative process was followed approximately 5 times) and the final result was the generation of 2,000 synthetic dumping images. The Figure 2 shows the comparison between model performance in terms of the materials depicted in the synthetic data. Based on these results, we decided to use the following material distribution in the synthetic data. (a) Plastic (Number of items in each generated image varied between, 0-30%), (b) Wood (Number of items, 0-35%), (c) Packages and Paper waste (Number of items, 0-15%), (d) Bottles and Cans (Number of items, 0-20%).

### C. Dumping classification

Two different classification models were developed, a) A basic convolutional neural network (CNN) classification model, b) A deeper model with residual blocks. These two models are described below.

1) *Basic CNN classification model*:: The basic CNN model is a baseline model consisting of three hidden blocks (Figure 3.a). Each block consists of a convolutional layer, followed by pooling and batch normalization. After each block, the input size was halved by setting the stride equal to 2 and the filter size was doubled. At the final block, the output was flattened before passing through the final dense layers. The output of the dense layers applies a sigmoid activation. This model was used to create a baseline and provide feedback for the generation of the synthetic data.

2) *Residual block classification model*:: The residual block classification model consists of two types of residual blocks, as shown in Figure 4: (a) A typical residual block (RBLK) (Figure 4.a), (b) A down-sampling residual block (DRBLK) (Figure 4.b).

A typical residual block (RBLK) contains two convolutional layers at the data path and a convolutional layer with a

TABLE I  
TEST RESULTS OF THE DL MODELS ON REAL-WORLD DATASETS.

Model	Natural Image	
	Precision	Recall
<i>Basic CNN model</i>	0.98	0.90
<i>Residual model</i>	0.97	0.92

kernel size of 1 at the residual connection path. The down-sampling residual block differs in the stride used for the first convolutional layer and the skip connection. Using a higher stride in these convolutions, the previous feature maps are down-scaled by a factor of 2 in the first convolutional layer of the block and the skip connection, which reduce the input feature map to half the input size, resulting in a smaller feature map.

The residual block classification model consists of these two pairs of residual blocks. Each pair comprises a RBLK followed by a DRBLK, as shown in Figure 3.b. The output of the final residual block is flattened and fed to the dense layers. The final output layer consists of a single neuron with a Sigmoid activation. The residual architecture was selected for its good performance on classification tasks [25].

3) *DL model training*:: The dataset used for training the DL models consisted of 16,000 non-dumping images and 2,000 synthetic dumping images (see Section III-B1). An additional 20% (35 images) of real-world dumping site images were added to the training dataset. Simple augmentations were applied on the patches during training: rotations of 90, 180 and 270 degrees, colour channel small-value shifting, blurring, brightness adjustment and vertical/horizontal shifting. All parameters were initialized with the He normal technique [26]. The dataset was randomly split into two sets: a training set (80%) and a validation set (20%). The validation set was used for hyperparameter fitting. A separate real-world dataset was then used as the testing set, to assess the model’s performance (see Section III-A) on real images. Both DL models were trained using the Adam optimizer and the binary cross-entropy loss function, using an Nvidia RTX2060 GPU.

#### D. Geographic information system data output

A dumping site detection pipeline was created using Python (Figure 1), as mentioned before. The inputs to the pipeline was the real-world dataset used for testing (2018 IEEE GRSS Data Fusion Challenge) and the output was a GIS vector file with the exact locations of the dumping sites detected. Since the model operates on patches of size  $256 \times 256$ , the original RGB image of size  $11920 \times 12020$  was divided into smaller patches (as described in III-B).

The image patches were given as input to the DL classification models and a class for each patch was predicted (i.e. 1 for *Dump* class and 0 for *non-dump* class). The sliding window patch coordinates (x, y) were recorded only for images classified as dumping sites. These coordinates were used to determine the pixel coordinates [X, Y] of the image patch with respect to the HR\_IMG.

$$[X, Y] = \left[ w \times \left( \frac{2x + 1}{2} \right), h \times \left( \frac{2y + 1}{2} \right) \right] \quad (1)$$

The patch coordinates [X, Y] were converted into geo-coordinates  $[X_{geo}, Y_{geo}]$  using the Python GDAL library [27]. The image patch ID, class and geo-coordinates were plugged to a vector layer (CSV). The vector file was uploaded as a layer file to some GIS software (QGIS [28]). The QGIS pinpoints the locations of dumping sites on an actual geographical map, which could then be used to inform the authorities to take appropriate measures. As future work, a web application will be developed for authorities to visualize and monitor the dumping as well as actions taken remotely.

## IV. RESULTS

Both models were trained with synthetic data and small percentage of real-world data but were tested only on real-world data, i.e. aerial images. The results from the test are shown in Table I, considering the metrics of precision and recall. The precision measures the percentage of images that were correctly classified as dumping sites, and the recall measures the percentage of actual dumping site images that were correctly classified as dumping sites.

For the basic CNN model, precision and recall scores of 0.98 and 0.90 respectively were recorded for the real-world images. Similarly, for the residual model, precision and recall scores of 0.97 and 0.92 were recorded for the real-world images respectively. Both models showed a similar performance, but the residual model had lower number of false positives, i.e. it had less number of non-dumping images classified as dumping sites. On the other hand, the basic CNN model had a lower number of false negatives, i.e. misclassification of dumping images as non-dumping images.

## V. DISCUSSION

Results have shown that the proposed DL models address the problem of detecting dumping sites quite well, showing high accuracy and very good performance. The results show that the proposed solution has high classification accuracy with a very low number of misclassified samples. Due to unavailability of any relevant work in synthetic data in dumping site detection, we couldn’t compare our results with any state-of-art technique. But the method used in this paper constitutes a novel approach for detection of dumping site from high-resolution aerial imagery. The proposed solution tackles a challenging problem in computer vision in generally, i.e. the problem of having to deal with highly unbalanced datasets and extremely under-represented classes of interest. The approach used for generating additional data for training the DL models works quite well, allowing to create diverse training samples that capture the whole spectrum of the data distribution of the real-world dataset used, for the city of Houston, USA.

We strongly believe that by exploiting the proposed method, potential users and stakeholders such as governmental agencies, municipalities and local authorities have the following options and opportunities:



Fig. 2. Graph showing model performance vs synthetic data generation. During each trail (from left to right) new material usually found in waste were added which leads to better model performance

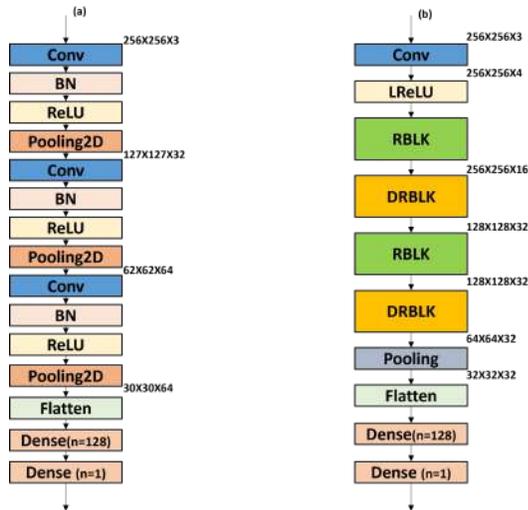


Fig. 3. The architecture of the DL classification models: a) Basic CNN classification model, b) Residual block classification model. *BN* stands for batch normalization, *n* is the number of neuron. The basic CNN uses "valid" padding for all convolutions, while the residual network uses "same" padding.

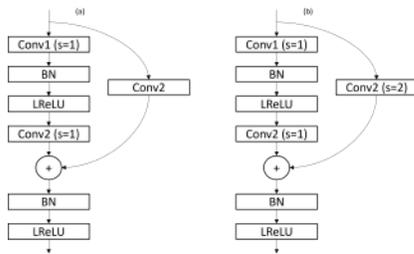


Fig. 4. The architecture of the two types of residual blocks used in the proposed models: (a) the typical residual block (RBLK). (b) The down-sampling residual block (DRBLK) using a stride of 2 at both the first convolutional layer and the skip connection. *LReLU* stands for parametric ReLU and *s* is the stride of the convolutional layer.

- Detect illegal dumping accurately and as fast as possible, within few days from the occurrence of the activity (i.e. time frequency depends on the next flight of a UAV/aeroplane or the passing by a satellite).
- Alleviate the environmental impact of illegal dumping, by understanding where illegal dumping occurs and when, taking cleaning measures and monitoring the measures taken.
- Fighting dumping effectively, before it becomes a potential source of environmental pollution and contamination.

- By recording the dumping locations, visualizing them in space and time, allows detecting hot-spots of frequent illegal dumping activity. This could facilitate more targeted campaigns for raising pro-environmental and responsible behaviour of the local citizens.
- Encourage, promote and empower waste reuse and recycling, building successful and enlightening stories of circular economy. Waste materials such organic waste, construction work, plastic, wood, electrical appliances and metals could be reused in a range of industries, such as agriculture, construction, manufacturing, etc.

As a more complete solution suitable for commercial use, the outputs of the proposed method could be fed to a GIS application, which would take the outputs from the DL model and place the coordinates where (and when) dumping has happened on a map of different cities, regions or countries, which would then be divided in different administrative zones (i.e. regions, municipalities, communities and/or villages). Each agency, organization or body responsible for some zone would have access only to dumping activity inside that zone, being able to perform the necessary actions.

#### A. Limitations

Although the proposed solution was able to correctly predict the class for most of the images as dumping or non-dumping, there were few image patches which were wrongly classified. Figure 5 (Row-1) depicts a few examples which were misclassified as dumping. The confusion occurs mainly due to the characteristics of those images, which contain visual features that are very similar to waste dumping. Similarly, Figure 5 (Row-2) shows few examples which have been misclassified as non-dumping while they constitute actual dumping sites. Such issues of eliminating false negatives/positives will be solved by adding more examples to the training, better capturing the whole spectrum of the data distribution of waste, garbage and dumping. Moreover, another limitation is the fact that the proposed approach requires high-resolution imagery to work properly. It is questionable whether the DL classification models would work well with lower-resolution aerial photography or satellite imagery.



Fig. 5. Examples of false predictions of the DL models. Row-1 shows examples of the non-dumping class misclassified as dumping site. Row-2 shows examples of the dumping class misclassified as a non-dumping site.

## B. Future work

Future efforts will focus on improving the classification accuracy of the proposed DL models, as described in Section V-A. Various optimization techniques will be considered, while the generation of more complex synthetic data will be performed, including low-represented incidents of dumping, especially scenarios where the DL models could not correctly detect dumping. More data will be used for training the model to advance its performance and achieve generalization on dumping sites located in diverse regions of the world.

Moreover, following the discussion of Section V in respect to the possibilities and features of a complete solution based on the proposed method for local authorities and municipalities, we plan to develop web and mobile applications, which would then allow clients to visualize and monitor illegal dumping occurrences, taking appropriate measures and actions.

## VI. CONCLUSIONS

This paper presented a method for the accurate detection of illegal dumping sites from high-resolution airborne images based on deep learning. The performance of the proposed method was assessed based on aerial photography of the city of Houston, USA. Due to the unavailability of labelled data required to train the DL models used, synthetically generated images were used, allowing the DL models to successfully learn the problem of detecting waste and dumping occurrence. The results showed that the proposed method solves the problem with high precision, and it constitutes a promising technique for problems dealing with highly imbalanced or very limited data.

## ACKNOWLEDGMENTS

Chirag Padubidri, Andreas Kamilaris and Savvas Karatsiolis have received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 739578 complemented by the Government of the Republic of Cyprus through the Directorate General for European Programmes, Coordination and Development.

## REFERENCES

- [1] Municipal waste statistics - european statistical office (2021). URL [https://ec.europa.eu/eurostat/statistics-explained/index.php/Municipal\\_waste\\_statistics#Municipal\\_waste\\_generation](https://ec.europa.eu/eurostat/statistics-explained/index.php/Municipal_waste_statistics#Municipal_waste_generation)
- [2] S. Kaza, L. C. Yao, P. Bhada-Tata, F. V. Woerden, , Washington, DC: World Bank, 2018. doi:10.1596/978-1-4648-1329-0. URL <https://doi.org/10.1596%2F978-1-4648-1329-0>
- [3] D. Ichinose, M. Yamamoto, On the relationship between the provision of waste management service and illegal dumping, *Resource and energy economics* 33 (1) (2011) 79–93.
- [4] S. K. Singh, P. Chokhandre, P. S. Salve, R. Rajak, Open dumping site and health risks to proximate communities in mumbai, india: A cross-sectional case-comparison study, *Clinical Epidemiology and Global Health* 9 (2021) 34–40. doi:<https://doi.org/10.1016/j.cegh.2020.06.008>.
- [5] B. G. Meerburg, G. R. Singleton, A. Kijlstra, Rodent-borne diseases and their risks for public health, *Critical Reviews in Microbiology* 35 (3) (2009) 221–270, pMID: 19548807. arXiv:<https://doi.org/10.1080/10408410902989837>, doi:10.1080/10408410902989837.
- [6] *Journal of Architecture, Planning and Construction Management* 9 (1) (Oct. 2019). [link]. URL <https://journals.iium.edu.my/kaed/index.php/japcm/article/view/335>
- [7] M. Colon, B. Fawcett, Community-based household waste management: Lessons learnt from exnora's 'zero waste management' scheme in two south indian cities, *Habitat International* 30 (4) (2006) 916–931, solid Waste Management as if People Matter. doi:<https://doi.org/10.1016/j.habitatint.2005.04.006>.
- [8] H. Pasang, G. A. Moore, G. Sitorus, Neighbourhood-based waste management: A solution for solid waste problems in jakarta, indonesia, *Waste Management* 27 (12) (2007) 1924–1938. doi:<https://doi.org/10.1016/j.wasman.2006.09.010>.
- [9] Q. Zhang, H. Li, X. Wan, M. Skitmore, H. Sun, An intelligent waste removal system for smarter communities, *Sustainability* 12 (17) (2020) 6829. doi:10.3390/su12176829. URL <https://doi.org/10.3390>
- [10] K. Pardini, J. J. Rodrigues, O. Diallo, A. K. Das, V. H. C. de Albuquerque, S. A. Kozlov, A smart waste management solution geared towards citizens, *Sensors* 20 (8) (2020) 2380. doi:10.3390/s20082380.
- [11] Trashout, Trashout - world cleanup day partner. URL <https://www.trashout.ngo/>
- [12] J. Gill, K. Faisal, A. Shaker, W. Y. Yan, Detection of waste dumping locations in landfill using multi-temporal landsat thermal images, *Waste Management & Research* 37 (4) (2019) 386–393, pMID: 30632930. arXiv:<https://doi.org/10.1177/0734242X18821808>, doi:10.1177/0734242X18821808.
- [13] K. Glanville, H.-C. Chang, Remote sensing analysis techniques and sensor requirements to support the mapping of illegal domestic waste disposal sites in queensland, australia, *Remote Sensing* 7 (2015) 13053–13069. doi:10.3390/rs71013053.
- [14] O. Youme, T. Bayet, J. M. Dembele, C. Cambier, Deep learning and remote sensing: Detection of dumping waste using uav, *Procedia Computer Science* 185 (2021) 361–369, big Data, IoT, and AI for a Smarter Future. doi:<https://doi.org/10.1016/j.procs.2021.05.037>.
- [15] A. Dabholkar, B. Muthiyar, S. Srinivasan, S. Ravi, H. Jeon, J. Gao, Smart illegal dumping detection, in: 2017 IEEE Third International Conference on Big Data Computing Service and Applications (Big-DataService), 2017, pp. 255–260. doi:10.1109/BigDataService.2017.51.
- [16] A. Kamilaris, C. van den Brink, S. Karatsiolis, Training deep learning models via synthetic data: Application in unmanned aerial vehicles, *CoRR abs/1908.06472* (2019). arXiv:1908.06472.
- [17] W. Qiu, A. L. Yuille, Unrealcv: Connecting computer vision to unreal engine, *CoRR abs/1609.01326* (2016). arXiv:1609.01326.
- [18] P. S. Rajpura, R. S. Hegde, H. Bojinov, Object detection using deep cnns trained on synthetic images, *CoRR abs/1706.06782* (2017). arXiv:1706.06782.
- [19] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, *CoRR abs/1409.4842* (2014). arXiv:1409.4842.
- [20] S. R. Richter, V. Vineet, S. Roth, V. Koltun, Playing for data: Ground truth from computer games, *CoRR abs/1608.02192* (2016). arXiv:1608.02192.
- [21] S. R. Richter, Z. Hayder, V. Koltun, Playing for benchmarks, *CoRR abs/1709.07322* (2017). arXiv:1709.07322.
- [22] S. P. B. L. S. N. Y. R. Hansch, 2018 ieee grss data fusion challenge – fusion of multispectral lidar and hyperspectral data (2020). doi:10.21227/jnh9-nz89.
- [23] Free3d, <https://free3d.com/3d-models>, accessed on: 2021-03-25.
- [24] Cgtrader, <https://www.cgtrader.com/free-3d-models/blender>, accessed on: 2021-03-25.
- [25] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778. doi:10.1109/CVPR.2016.90.
- [26] K. He, X. Zhang, S. Ren, J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in: 2015 IEEE International Conference on Computer Vision (ICCV), 2015, pp. 1026–1034. doi:10.1109/ICCV.2015.123.
- [27] GDAL/OGR contributors, GDAL/OGR Geospatial Data Abstraction software Library, Open Source Geospatial Foundation (2021). URL <https://gdal.org>
- [28] QGIS Development Team, QGIS Geographic Information System, QGIS Association (2021). URL <https://www.qgis.org>