2023-06-05

Poly-GAN: Regularizing Polygons with Generative Adversarial Networks

Lasith Niroshan
Technological University Dublin, d19126805@mytudublin.ie

James Carswell
Technological University Dublin

Funder: Technological University Dublin College of Arts and Tourism

Poly-GAN: Regularizing Polygons with Generative Adversarial Networks

Lasith Niroshan[1] [0000-0002-9868-8338] and
James D. Carswell[2] [0000-0002-4766-7297]

Technological University Dublin, Dublin, Ireland
[1] d19126805@mytudublin.ie, [2] james.carswell@TUDublin.ie

Abstract. Regularizing polygons involves simplifying irregular and noisy shapes of built environment objects (e.g. buildings) to ensure that they are accurately represented using a minimum number of vertices. It is a vital processing step when creating/transmitting online digital maps so that they occupy minimal storage space and bandwidth. This paper presents a data-driven and Deep Learning (DL) based approach for regularizing OpenStreetMap building polygon edges. The study introduces a building footprint regularization technique (*Poly*-GAN) that utilises a Generative Adversarial Network model trained on irregular building footprints and OSM vector data. The proposed method is particularly relevant for map features predicted by Machine Learning (ML) algorithms in the GIScience domain, where information overload remains a significant problem in many cartographic/LBS applications. It addresses the limitations of traditional cartographic regularization/generalization algorithms, which can struggle with producing both accurate and minimal representations of multisided built environment objects. Furthermore, future work proposes a way to test the method on even more complex object shapes to address this limitation.

Keywords: Geographic Information System · Polygon Regularization · Generative Adversarial Networks · OpenStreetMap

1 Introduction

Polygon *regularization*, also known as simplification/generalization, is a vital image processing step when digitising object shapes predicted from aerial/satellite imagery using Machine Learning (ML)/Deep Learning (DL) change detection algorithms. Overcoming the formation of built environment objects (e.g. buildings) with stair-like edges or otherwise irregular shapes (Fig. 1.a) remains a challenging topic much discussed in the computer vision and GIScience/Digital Earth domain [9].

Many related studies propose statistical or ML/DL based solutions for polygon detection/regularization. For example, in our previous work, an ML-based spatial change detection mechanism (called *OSM-GAN*) was proposed using satellite images and OpenStreetMap (OSM) vector data [12,13]. The output of this approach is an outline of any changed building object(s) in raster format, which then needs to be vectorised prior to uploading into online mapping platforms like OSM [14]. However, the outlined object first needs to be simplified (regularized) to represent the same map feature using a minimum number of vertices before the vectorisation process begins (Fig. 1.b). This new research extends previous work by appending to the DeepMapper [10] automated map update workflow an additional ML-based regularization process (called *Poly*-GAN) designed to produce fully vectorized map features having a minimum number of nodes.

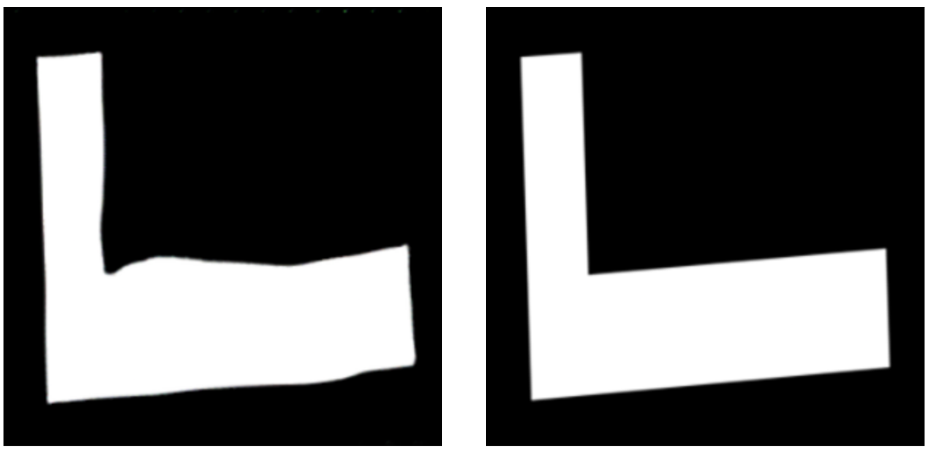

(a) DL predicted object containing irregular edges and numerous nodes (193).

(b) A simplified (regularized) version of the object containing a minimum number of nodes (6).

Fig. 1. Non-regularized (a) and regularized (b) building footprints.

In other words, this investigation mimics the ML-based OSM-GAN change detection mechanism already used by *DeepMapper* to now include *Poly*-GAN, another Generative Adversarial Network (GAN) model – this time specifically trained to the purpose of polygon regularization.

This paper is structured into five sections, including this Introduction. Section 2 follows by providing more background on the topic of polygon regularization and the challenges it presents for image analysis. Section 3 provides a detailed explanation of the *Poly*-GAN modelling mechanism, highlighting its advantages and innovations. It describes the complete polygon regularization methodology including data preparation, pre- and post-processing, and refining polygon edges with GANs. Section 4 presents some experimental results and a discussion of their implications. Finally, Section 5 concludes the paper and discusses potential future research directions.

2 Background

Traditional cartographic regularization approaches have issues with producing both accurate and minimal representations of complex built environment structures. These algorithms depend on pre-defined threshold values that need to be discovered empirically or consist of deterministic rules (IF-THEN-ELSE) instead of rules derived from the data – as in Machine Learning approaches. Consequently, traditional methods to address these issues can adversely affect polygon regularization accuracy, while ML/DL-based methods can produce more consistently reliable results.

Briefly, traditional regularization algorithms commonly include the *Lang simplification algorithm* [6], the *Ramer-Douglas-Peucker* (RDP) algorithm [16], the *Zhao-Saalfed algorithm* [27], *Opheim simplification* algorithm [15] and the *Reumann-Witkam* algorithm [17]. Figure 2 illustrates the behaviour of each algorithm applied to a sample non-regularized building footprint. In addition to these algorithms, *perpendicular distance simplification*[1] was also implemented and tested in this study.

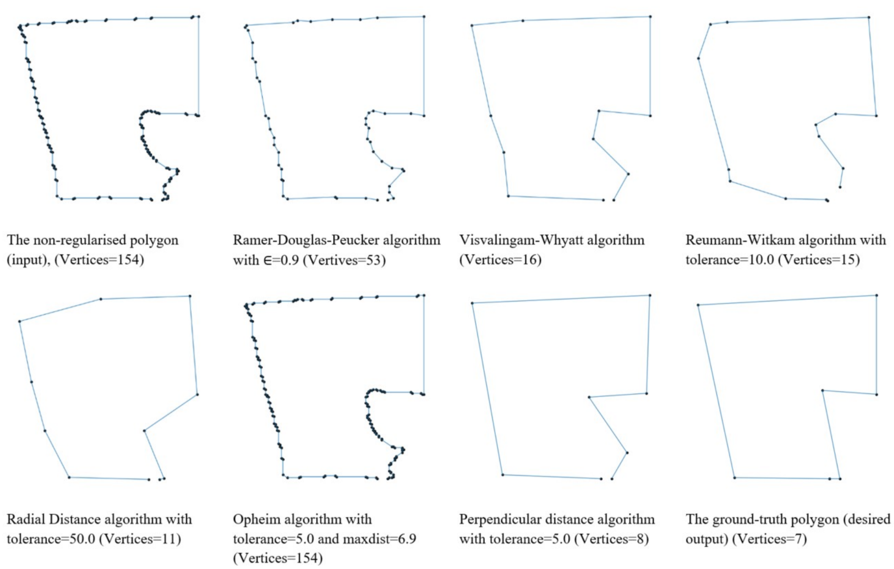

Fig. 2. A comparison of traditional regularization algorithms on a noisy polygon in terms of node reduction, shape simplification, and edge smoothness. The total number of nodes (vertices) produced in each case is also given.

In 2012, Sohn et al. proposed a regularization method for building rooftop models using LiDAR data [20]. The proposed method is based on Minimum

[1] https://psimpl.sourceforge.net/perpendicular-distance.html

Description Length (MDL) theory and comprises two stages: 1) Hypothesis generation and 2) Global Optimisation. Lu et al. (2018) used a richer convolutional feature (RCF) network to create an edge probability map (i.e. a map of building footprints) and an edge refinement process according to morphological analysis of the topographic surface [8]. These approaches primarily tried to improve the edge refinement process by addressing issues such as stair-like noise, isolated points, and outliers using a non-maximal suppression (NMS) algorithm. Their methods outperform the NMS algorithm in terms of accuracy.

A study by Zhao et al. investigated a data-driven approach for boundary regularization along with a machine-learning approach to building extraction from satellite images [26]. This approach includes three steps: 1) Initial modelling, 2) Hypothesis generation, and 3) Minimum Description Length (MDL) optimisation. At the initial modelling, the Ramer-Douglas-Peucker algorithm was applied to simplify the initial polygon. In the next phase, local hypothetical models were generated using a set of temporal points and lines. Finally, MDL optimisation was performed to assert the optimal hypothesis among local hypotheses.

Zorzi and Fraundorfer (2019) proposed an approach inspired by style transfer techniques that utilise adversarial losses to generate accurate building boundaries [28]. The proposed regularization model is trained using satellite images of Jacksonville, Florida and OpenStreetMap building footprints. Three types of loss functions were used to achieve regularized and visually pleasing building footprints. However, some inconsistent occurrences presented in the visual outputs, such as rotation, skew, and false detections, which were inherited from the initial feature mask predictions.

In 2020, Wei et al. proposed a polygon simplification algorithm for complex conditions such as different building shapes, image resolutions, and low-quality image segmentation since image segmentation quality significantly impacts regularization results [23]. The suggested algorithm consists of two components. First, coarse adjustment, contains empirical rules to remove any segmentation errors. Second, fine adjustment, intends to refine the polygon using Awrangjeb's approach [1] for extracting building polygons from point clouds. A relationship between polygon regularization and morphological filtering was demonstrated in a study conducted by Xie et al. (2020) [24]. The proposed methodology explored the possibility of using morphological building features to generate more realistic boundaries. The concept of using morphological features in this task is also explored in our proposed *Poly*-GAN method.

Each of the related approaches discussed above have their own advantages and disadvantages. When working with geospatial data, it has been shown that the method used must be specifically adapted to the characteristics of the area being mapped. In other words, classification performance depends on the morphological features and other local spatial data attributes. Such uncertainty encouraged this study to utilise state-of-the-art Deep Learning techniques and crowdsourced geo-data to propose a more reliable regularization mechanism. As such, this paper introduces *Poly*-GAN, a novel solution inspired by image-to-image translation for addressing the polygon regularization problem that uses large volumes of crowdsourced geo-data to train data-hungry real-world ML models.

3 Methodology

The following Sections describe in more detail the proposed GAN-based polygon regularization algorithm (*Poly*-GAN), including pre- and post-processing procedures.

3.1 Data Preparation

Data preparation is a necessary part of ML-based research, particularly for Deep Learning models [7]. To develop an effective DL model for polygon simplification, both noisy and corresponding ground truth data are essential for training. This study uses OSM-GAN predicted building polygons for noisy data, while OSM building footprints were used as the ground truth component for training polygon simplification models.

After the OSM-GAN model generates the feature-map for a given satellite image, the noisy predicted building polygons are separated using an *instance segmentation* process. Following the creation of these predicted footprint segments, a lookup process is launched to retrieve the corresponding OSM building footprint in a local vector map database. A filtering mechanism removes anomalous data (i.e. incorrect OSM-GAN polygons) from the dataset when both components are present. Finally, the filtered data is combined into a 600x300 pixel image for input into the DL model training process. The building footprints were not zoomed to the size of the input images to maintain consistency of prediction results. Fig. 3 (left half) illustrates some of the different building objects predicted by the OSM-GAN model together with their “ground truth" counterparts.

3.2 *Poly*-GAN Modelling

The Generative Adversarial Network (GAN) modelling architecture consists of two main components (each a Neural Network): the *Generator* and the *Discriminator*. The Generator component of the GAN is responsible for producing new (but fake), regularized polygons by learning from a dataset of previously regularized polygons (ground truth data). The Discriminator component then attempts to determine whether the generated polygons are real or fake to improve the Generator’s ability to produce realistic polygons. The two components work together in a competition, with the Generator trying to produce more realistic polygons and the Discriminator trying to become better at identifying fake polygons.

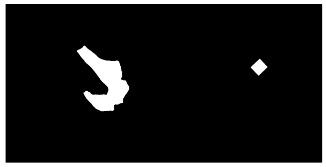

(a) A wrong prediction by OSM-GAN.

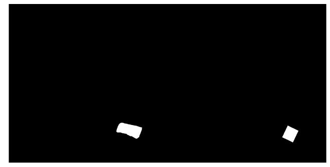

(b) OSM-GAN predicts a footprint in the wrong orientation.

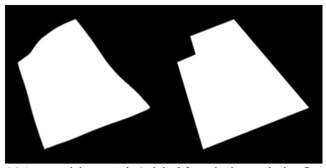

(c) A zoomed data sample (original footprint is stretched to fit into the image size).

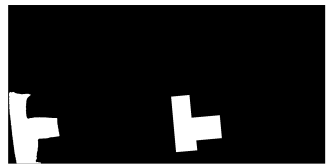

(d) An accepted data sample according to the filtering rules.

Fig. 3. Different forms of data samples used in this study. All left side objects were predicted by OSM-GAN model; right side objects are "ground truth" building footprints taken from OSM. First two samples (a and b) were not accepted and filtered out of the *Poly*-GAN training dataset.

The main idea behind image-to-image translation in a ML context is that a given input image (e.g. sketch/outline of an object) translates or transforms into another higher-level representation (e.g. photo-realistic image) of the set of input information. Isola et al. [4] presented several generalized uses of Conditional GAN based image-to-image translation, such as labels-to-street scenes, black&white images-to-colour images, sketches-to-photos, and especially aerial images-to-maps, which is important for this study. Pix2Pix is their implementation of image-to-image translation, which is freely available for use on GitHub[2].

This study uses an updated version of Pix2Pix in its polygon regularization modelling experiments. Once data samples were prepared and split into *train* and validation categories, they were then uploaded to Kay[3], Ireland's national supercomputer, to train the proposed *Poly*-GAN regularization model. Kay cuts model training times from a few days on a typical "gamer-spec" laptop to a few hours – allowing to train and test multiple models relatively quickly with different hyperparameters and datasets [11]. The following hyperparameters (Table 1) were applied to train the *Poly*-GAN model.

[2] `https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix`

[3] `https://www.ichec.ie/about/infrastructure/kay`

Table 1. Values for key hyperparameters used in the *Poly*-GAN model training process.

Name	Value	Description
batch_size	1	Number of images in a batch.
gan_mode	vanilla	The type of GAN objective. (i.e. vanilla, lsgan, wgangp)
init_gain	0.02	Scaling factor for the network.
init_type	normal	Network initialisation method.
input_nc	1	Number of input image channels.
output_nc	1	Number of output image channels.
lr	0.0002	Initial learning rate for adam optimisation.
lr_policy	linear	Learning rate policy.
n_epochs	100	Number of epochs with the initial learning rate.
ndf	64	Number of discriminative filters in the first conv layer.
ngf	64	Number of generative filters in the first conv layer.
netD	basic	The type of the discriminator architecture.
netG	unet_256	The type of the generator architecture.
pre-process	resize_and_crop	Scaling and cropping of images at load time

3.3 A Combined Regularization Method

Poly-GAN aims to improve the accuracy and efficiency of polygon regularization by combining data-driven and DL-based regularization methods into a single solution. Using a Generative Adversarial Network in this approach allows for a supervised learning technique, where the network can learn from a dataset of previously regularized polygons. Once an irregular building footprint is obtained (predicted) from the OSM-GAN change detection process, a pre-processing step stores its geo-referenced coordinates. Simultaneously, the Ramer-Douglas-Peucker (RDP) algorithm processes the building's shape to reduce the number of redundant polygon nodes and simplify it for input to the GAN-based regularization modelling procedure.

The refined/simplified polygon is then further processed to generate an input data sample for training the *Poly*-GAN model. An overview of the process is shown in Figure 4 and consists of two phases: the training phase and the prediction phase. The training phase involves the use of a dataset of ground-truth building footprints to train the *Poly*-GAN model. The prediction phase involves the use of the now trained *Poly*-GAN model to regularize the predicted building footprints obtained from the OSM-GAN change detection process.

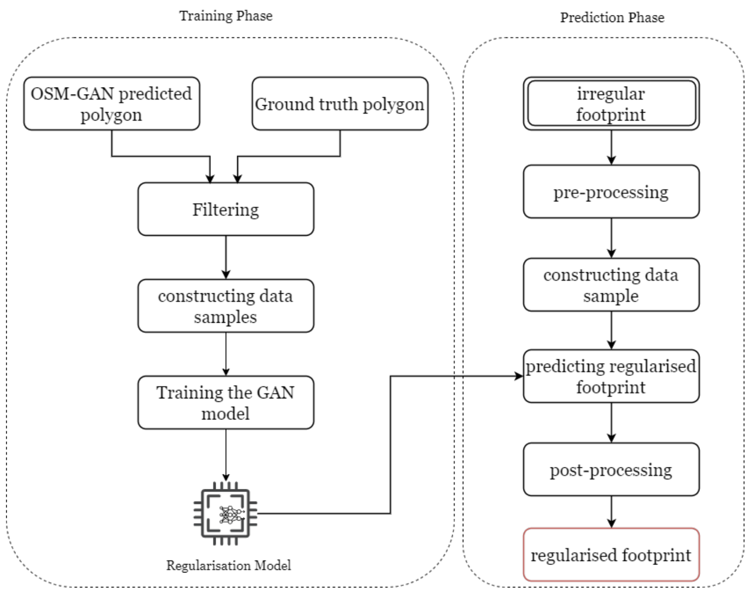

Fig. 4. Schematic diagram of the polygon regularization process linking the *Poly*-GAN model training phase to the (predicted) building regularization phase.

Pre-Processing As mentioned, the pre-processing step for the proposed building regularization procedure begins by storing the geo-referenced coordinates of the predicted building footprint generated by the OSM-GAN change detection algorithm. However, analysis of the OSM-GAN results show that the generated footprints can contain a considerable number of redundant nodes, which can affect the accuracy of polygon regularization. In this case, the RDP algorithm is applied to perform node reduction while preserving the overall shape of the predicted building footprint.

The RDP algorithm is a popular algorithm for polygon simplification that aims to reduce the number of vertices in a polygon while preserving the shape of the polygon as much as possible. The simplification is performed by identifying and removing redundant vertices that are not essential to the overall shape of the polygon. The algorithm works by iteratively removing vertices that are within a specified distance tolerance, called *epsilon*, from a straight line connecting the start and end vertices of the polygon. The RDP algorithm is simple, fast, and efficient, making it a popular choice for polygon regularization in GIScience, cartography, and computer vision domains [21]. Additionally, this algorithm is easy to implement and does not require complex parameter tuning, making it useful for various dataset types. In Figure 5, the footprints generated by the OSM-GAN change detection algorithm are compared with the simplified footprints produced by RDP.

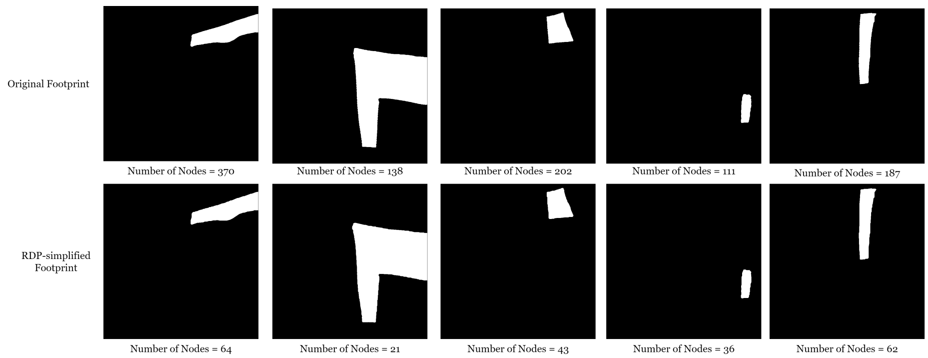

Fig. 5. Comparison of original input footprints returned from OSM-GAN algorithm (top row) and RDP-simplified footprints (bottom row). The captions show the number of nodes in each.

After applying the RDP algorithm, the pre-processed footprint is saved into a 300x300 pixel black&white image; this is used as the inference data sample in the next step. This pre-processing procedure ensures that the input data to the GAN-based regularization algorithm is a cleaned, pre-processed version of the initially predicted (noisy) footprint. This helps to produce a more accurate and reliable regularization of the building's footprint in the next step.

Refining Polygon Edges with Generative Adversarial Networks The above pre-processed building footprint is then combined with an empty mask, resulting in an image that is 300x600 pixels in size. The empty mask, which is 300x300 pixels, serves as a placeholder for the generator to fill in with the regularized polygon. By providing the generator with a clear distinction between the building footprint and the area to be regularized, the GAN can focus on making changes to the specific area of the image that needs to be regularized. This process continues until the Generator produces polygons (i.e. buildings) that are indistinguishable from real ones. The final *Poly*-GAN regularization model developed in this work also reshapes irregular building footprints in order to make them more "regular". The result of this combined approach is dependent on the quality of the dataset used for training the GAN, as well as the specific architecture and hyperparameters chosen for the GAN. Finally, the predicted footprint is passed through a post-processing procedure to assess its quality and create an OSM-acceptable changeset. Figure 6 illustrates the prediction results from the GAN-based regularization model.

Post-Processing Once *Poly*-GAN regularization simplifies the building footprint, the result is passed through a post-processing procedure to refine the footprint further. Overall, the post-processing steps aim to ensure that the building footprint is accurate, minimal, reliable, and meets the mapping

Fig. 6. Regularized building footprints after applying the GAN-based regularization method to map objects shown in bottom row of Fig.5. The *Poly*-GAN architecture learns to reshape the irregular building footprints while preserving the overall shape of the building and other important features such as corners.

conventions of OSM. The output of this step is a high-quality regularized building footprint that can be uploaded to the online OSM dataset.

The first step in the post-processing phase is to extract the GAN-regularized footprint from the 300x600 pixel image and re-apply the coordinates that were recorded in the pre-processing phase. Then a *perpendicular distance algorithm* (PD) is applied to the extracted footprint since experiments show this to be a reliable and accurate method for producing building footprints with minimal nodes while preserving the overall shape of the building (See Figure 2).

This algorithm works by measuring the perpendicular distance between a point and a line and removing the point if it falls within a specified threshold distance from the line (Figure 7). The algorithm does not require complex parameter tuning, which makes it useful for various dataset types, and it is relatively fast and efficient, which makes it suitable for use in real-time applications.

After the PD algorithm is applied, the *Poly*-GAN generated footprints are re-georeferenced using the geo-data stored during the pre-processing step. Re-georeferencing the footprints using the pre-processing stored geo-data assumes that the original geo-data is accurate and that the GAN-generated footprints are aligned with the original raster image.

Finally, an OSM-acceptable changeset is built, containing any modifications made to the *Poly*-GAN simplified footprint that passed the post-processing procedure. The post-processing procedure aims to ensure that the resulting footprint is of high quality and meets the requirements of OSM mapping conventions. This may include editing the positions of building vertices, merging or splitting polygons, and adding/removing attributes for polygons. The output from this step is a reliable, accurate, and regularized building footprint which can be uploaded to the online OSM dataset.

4 Results and Discussion

In the ML domain of generative models, such as GANs, it is suggested that qualitative analysis is often more informative and effective than quantitative analysis, as it provides a more comprehensive understanding of the generated samples [5,22]. Quantitative analysis metrics, for instance, *Fréchet Inception*

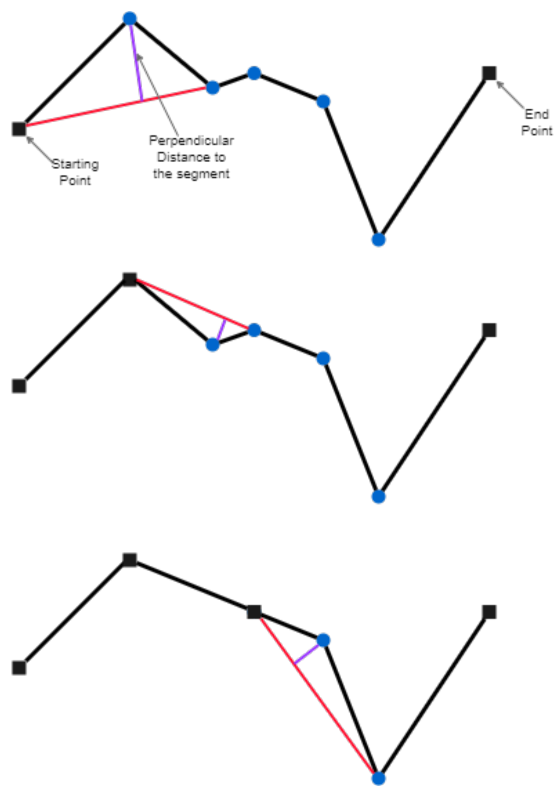

Fig. 7. The perpendicular distance algorithm for polygon simplification. Initially, the first and third nodes are used to define a line segment, and the perpendicular distance to the second node is calculated and compared against a given threshold. After that, the algorithm is moved to the next node pairs which are the second and fourth. If the calculated distance is lower than the threshold the node gets removed from the polygon, and so on.

Distance [3] or *Inception Score* [18], can provide a good indication of the quality of the generated data samples, but these metrics only provide a limited view of generated samples.

When working with GANs, qualitative analysis can help to understand the biases that the model has learned from the training data, which can be important for avoiding unwanted consequences. In relation to *Poly*-GAN, qualitative analysis was used in the experiments, which has several advantages. For example, a human expert could check at a glance if a building footprint is complete, covers the correct area, has the correct shape, or adheres to OSM mapping conventions. Figure 8 depicts both the original footprint and the *Poly*-GAN regularized result. For possible use in future research and applications, both non-regularized (containing all OSM-GAN predicted vertices) and regularized (containing only *Poly*-GAN regularized vertices) building footprints are saved in *GeoJson* format. This format ensures that the footprints can be easily accessed and integrated into a wide range of GIS software and LBS applications.

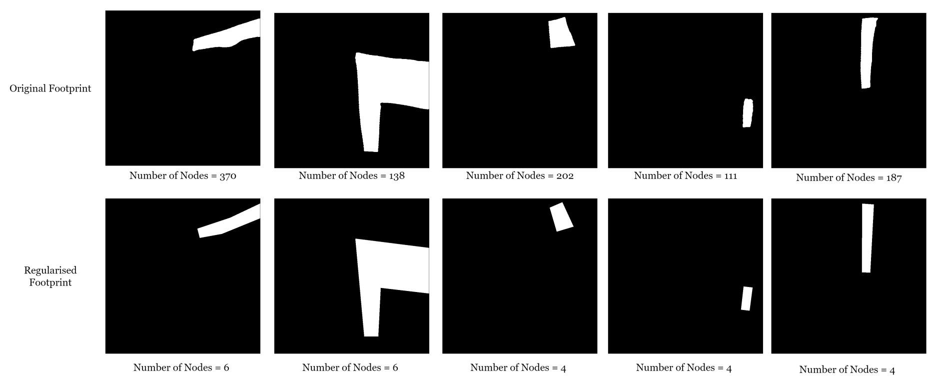

Fig. 8. Comparison of original (OSM-GAN predicted) irregular building footprints (top row) with their corresponding *Poly*-GAN regularized and simplified polygons (bottom row) obtained using a combined approach of data-driven and GAN-based regularization methods. The number of nodes in the final polygons is also shown, indicating object complexity reduction while preserving the building's overall shape.

Polygon regularization is important in the context of digital maps, such as OpenStreetMap, because it helps to improve the accuracy and uniformity/consistency of online map data [2,19]. The regularization process can correct errors and inconsistencies in the map data and make it more consistent with the real world. In the case of OSM, map data is often contributed by many volunteers, which can lead to inconsistencies in the data. For example, building footprints can be irregular and inconsistent in shape, size, and orientation. Regularization can correct these inconsistencies and make the building footprints more uniform with surrounding real-world buildings. Polygon regularization can also improve the quality of map data used for further analysis in other mapping applications - like 3D modelling, solar exposure/energy consumption analysis, and emergency planning. Additionally, regularization can reduce the overall complexity of map data, which can make it easier to work with and improve the performance of LBS applications by helping to reduce information overload [25].

It is important to note that the quality of the *reshaping* produced by *Poly*-GAN is dependent on several factors. The dataset used for training the GAN should be of high quality, containing a diverse set of building footprints that are representative of the real-world buildings that the GAN will be used to regularize. Additionally, the architecture and hyperparameters of the GAN should be carefully chosen to ensure that it can effectively reshape the footprints while preserving the overall shape of the building and important building features such as corners.

One limitation of the method presented is that it has not yet been tested on circular objects or other complex building shapes, such as buildings with empty areas (holes) inside. These types of shapes can present unique challenges for regularization and simplification, as they may require different threshold values or algorithms compared to more uniform, orthogonal building shapes. This limitation highlights the need for further research and experimentation to develop methods to effectively handle these types of shapes. In this regard, the effectiveness of the approach should be trained and tested on larger geographic datasets to optimise the architecture and hyperparameters of the GAN model and further improve the reshaping performance to suit the built environment where it is used.

5 Conclusions

The experiments carried out in this study show promising results in terms of extracting the key building vertices needed to preserve the overall shape of a building, and the final regularized polygon being acceptable for updating online mapping platforms such as OSM.

This combined polygon regularization approach integrates data-driven and Deep Learning-based methods. It was applied to irregular building footprints obtained from OSM-GAN change detection, where a pre-processing step involved storing geo-referenced coordinate data and applying the RDP algorithm (*epsilon* = 0.9) to reduce redundant polygon nodes. The reduced polygon was then used as input for the *Poly*-GAN algorithm, which aimed to regularize building footprints while preserving their overall shape. As a post-processing step, the regularized polygons were further refined using the Perpendicular Distance algorithm to extract key building vertices (e.g. corners) while continuing to maintain the overall shape of the building.

The work presented in this paper is a part of a comprehensive online map updating solution (called *DeepMapper*) that endeavours to provide an end-to-end automated workflow for populating OSM [10]. The complete solution is in final stages of development and includes several important components: geo-data (raster/vector) crawling and indexing, GAN-based change detection, GAN-based regularization, quality analysis, and OSM changeset creation and map updating. This fully integrated prototype aims to improve the accuracy, consistency, and efficiency of the online map-updating process by automating many of the manual tasks that VGI mappers typically carry out. This research is a step forward in improving the consistency and quality of crowdsourced maps and is expected to lead to further developments (e.g. regularizing complex objects and other map feature types) in the future.

Acknowledgements. The authors wish to thank all VGI contributors involved with the OpenStreetMap project. This research is funded by Technological University Dublin College of Arts and Tourism, SEED FUNDING INITIATIVE 2019–2020. The authors wish to acknowledge the Irish Centre for High-End Computing (ICHEC) for the provision of supercomputing facilities. We also gratefully acknowledge Ordinance Survey Ireland (OSi) for providing both raster and vector ground truth data used to verify accuracy experiments.

References

1. Awrangjeb, M.: Using point cloud data to identify, trace, and regularize the outlines of buildings. International Journal of Remote Sensing **37**(3), 551–579 (2016)
2. Grinberger, A.Y., Minghini, M., Juhász, L., Yeboah, G., Mooney, P.: Osm science—the academic study of the openstreetmap project, data, contributors, community, and applications (2022)
3. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
4. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
5. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
6. Lang, T.: Rules for the robot draughtsmen. The Geographical Magazine **42**(1), 50–51 (1969)
7. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. nature **521**(7553), 436–444 (2015)
8. Lu, T., Ming, D., Lin, X., Hong, Z., Bai, X., Fang, J.: Detecting building edges from high spatial resolution remote sensing imagery using richer convolution features network. Remote Sensing **10**(9), 1496 (2018)
9. Mooney, P., Corcoran, P.: Has openstreetmap a role in digital earth applications? International Journal of Digital Earth **7**(7), 534–553 (2014)
10. Niroshan, L., Carswell, J.: Deepmapper: Automatic updating crowdsourced maps. Poster, Technological University Dublin, College of Arts and Tourism (2020), https://arrow.tudublin.ie/gradcamoth/3
11. Niroshan, L., Carswell, J.D.: Machine learning with kay. AGILE: GIScience Series **3**, 11 (2022)

12. Niroshan, L., Carswell, J.D.: Osm-gan: using generative adversarial networks for detecting change in high-resolution spatial images. In: Geoinformatics and Data Analysis: Selected Proceedings of ICGDA 2022, pp. 95–105. Springer (2022)
13. Niroshan, L., Carswell, J.D.: Post-analysis of osm-gan spatial change detection. In: Web and Wireless Geographical Information Systems: 19th International Symposium, W2GIS 2022, Constance, Germany, April 28–29, 2022, Proceedings. pp. 28–42. Springer (2022)
14. OpenStreetMap: Openstreetmap (2004), https://www.openstreetmap.org, accessed on: 2023-03-04
15. Opheim, H.: Fast data reduction of a digitized curve (1982)
16. Ramer, U.: An iterative procedure for the polygonal approximation of plane curves. Computer Graphics and Image Processing **1**(3), 244–256 (1972). https://doi.org/https://doi.org/10.1016/S0146-664X(72)80017-0, https://www.sciencedirect.com/science/article/pii/S0146664X72800170
17. Reumann, K.: Optimizing curve segmantation in computer graphics. International Comp. Sympo. 1973. Amsterdam pp. 467–472 (1974)
18. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. Advances in neural information processing systems **29** (2016)
19. See, L., Mooney, P., Foody, G., Bastin, L., Comber, A., Estima, J., Fritz, S., Kerle, N., Jiang, B., Laakso, M., et al.: Crowdsourcing, citizen science or volunteered geographic information? the current state of crowdsourced geographic information. ISPRS International Journal of Geo-Information **5**(5), 55 (2016)
20. Sohn, G., Jwa, Y., Jung, J., Kim, H.: An implicit regularization for 3d building rooftop modeling using airborne lidar data. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences **1**(3), 305–310 (2012)
21. Szeliski, R.: Computer vision: algorithms and applications. Springer Nature (2022)
22. Theis, L., Oord, A.v.d., Bethge, M.: A note on the evaluation of generative models. arXiv preprint arXiv:1511.01844 (2015)
23. Wei, S., Ji, S., Lu, M.: Toward automatic building footprint delineation from aerial images using cnn and regularization. IEEE Transactions on Geoscience and Remote Sensing **58**(3), 2178–2189 (2019)
24. Xie, Y., Zhu, J., Cao, Y., Feng, D., Hu, M., Li, W., Zhang, Y., Fu, L.: Refined extraction of building outlines from high-resolution remote sensing imagery based on a multifeature convolutional neural network and morphological filtering. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing **13**, 1842–1855 (2020)
25. Yin, J., Carswell, J.D.: Spatial search techniques for mobile 3d queries in sensor web environments. ISPRS International Journal of Geo-Information **2**(1), 135–154 (2013)
26. Zhao, K., Kang, J., Jung, J., Sohn, G.: Building extraction from satellite images using mask r-cnn with building boundary regularization. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 247–251 (2018)
27. Zhao, Z., Saalfeld, A.: Linear-time sleeve-fitting polyline simplification algorithms. In: Proceedings of AutoCarto. vol. 13, pp. 214–223 (1997)
28. Zorzi, S., Fraundorfer, F.: Regularization of building boundaries in satellite images using adversarial and regularized losses. In: IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium. pp. 5140–5143. IEEE (2019)