SPATIOTEMPORAL ROCKFALL DETECTION USING POINT-BASED NEURAL NETWORKS

Thanasis Zoumpekas, Anna Puig, Maria Salamó
Department of Mathematics and Computer Science
University of Barcelona, Spain
{thanasis.zoumpekas,annapuig,maria.salamo}@ub.edu

ABSTRACT

Rockfall poses a significant threat to human life and infrastructure in mountainous regions, necessitating effective detection and mitigation strategies. Sensor technologies, such as Terrestrial Laser Scanners, are being widely utilized to periodically scan mountain terrains and cliffs acquiring 3D point clouds. Current research on intelligent rockfall detection utilizes pre-computed features and machine-learning models, clearly lacking hidden geometric properties inherent in 3D point clouds. Also, recent point-based deep learning approaches that focus on geometric feature extraction and end-to-end learning, mainly study datasets with balanced labeled observations, not addressing the rockfall class imbalance in real-world cases. Our approach builds upon advancements in point-based neural networks and integrates spatiotemporal information to enhance performance in detecting rockfall candidate areas in cases where rockfall observations are limited. Addressing the class imbalance issue inherent in rockfall detection, we present results in real-world 3D scans from a cliff in Spain showcasing the effectiveness of our method in identifying rockfall candidate areas.

1 INTRODUCTION

Rockfall is a common natural hazard that poses a significant threat to human life and infrastructure in mountainous regions. Effective rockfall detection and mitigation strategies are essential for risk reduction and prevention of potential disasters (Bourrier et al., 2009; Agliardi et al., 2009). Advances in sensor technology and machine learning have shown promise for improving rockfall detection accuracy and efficiency (Abellán et al., 2014; Jaboyedoff et al., 2012; Zoumpekas et al., 2021; Blanco et al., 2022). Currently, LiDAR scanners and photogrammetry tools are highly utilized to capture high-resolution 3D point clouds digitizing the geometry of rocky slopes and mountainous areas (Jaboyedoff et al., 2012; Westoby et al., 2012).

In rockfall detection tasks, point clouds are captured periodically to acquire information about changes in the cliff or other rock masses. Analyzing these point clouds through multi-temporal comparison and inspection is time-consuming and prone to errors based on factors like human skill and sensor sensitivity. While machine learning has enhanced automation in rockfall detection, current approaches overlook hidden geometric properties in 3D scans and rely on pre-computed features (Lague et al., 2013; Zoumpekas et al., 2021).

Point-based neural networks, such as PointNet (Qi et al., 2017a), and PointNet++ (Qi et al., 2017b), have shown potential in 3D object detection, segmentation, and recognition tasks (Guo et al., 2021). Their ability to directly process 3D point clouds and learn global and local geometric features enhances accuracy and efficiency in geometric deep learning tasks (Qi et al., 2017a;b). Given their success in learning from intricate geometric structures, point-based neural networks are particularly well-suited for the analysis of complex and irregular rock masses, making them a valuable tool in tasks where a 3D spatial understanding is important.

Despite several studies addressing the detection of rockfall using point-based neural networks, such as Farmakis et al. (2022; 2023), the class imbalance nature of the problem has not been fully investigated, as rockfall candidates are typically limited compared to other classes such as vegetation in certain mountainous areas and geological contexts. Also, although Zoumpekas et al. (2021) and

Blanco et al. (2022) have attempted to address this issue, they have relied on machine learning and resampling strategies on pre-computed features rather than point-based neural networks.

In this paper, we address the issue of class imbalance and we propose **(i)** a process to augment limited labeled data samples integrating spatiotemporal neighborhoods of points and **(ii)** a rockfall detection framework utilizing well-known point-based neural networks. Our proposal shows that integrating spatiotemporal information leads to performance improvements in rockfall detection using point-based deep learning models in cases where rockfall candidates are limited. Our results on 3D point clouds captured by a Terrestrial Laser Scanner of a cliff in Spain demonstrate the effectiveness of the proposed approach in detecting rockfall candidates.

2 Related Work

Traditional methods for rockfall detection and prediction rely on continuous monitoring of the terrain using various sensors and physical modeling (Gallo et al., 2021; Alvioli et al., 2021; Samodra et al., 2016). However, these methods require significant data acquisition, processing resources, and collection of geotechnical parameters to generate realistic models. Moreover, they are often limited by sensor placement and coverage.

Recent advances in remote sensing have led to the development of enhanced approaches for rockfall detection and prediction, leveraging point cloud data. Research efforts to identify changes in the cliff's surface and rockfall candidates typically use the *m3c2* method (Winiwarter et al., 2021) for change detection, followed by clustering-based feature extraction and statistical analysis (Blanco et al., 2022). More recent approaches to address rockfall detection have shifted to point cloud classification and segmentation (Farmakis et al., 2022; 2023; Weidner et al., 2019; Zoumpekas et al., 2021; Blanco et al., 2022; Battulwar et al., 2020).

Machine learning-based approaches such as Weidner et al. (2019), utilize classification models such as random forests. Other studies such as Zoumpekas et al. (2021) utilize multiple classifier models, resampling strategies, and feature selection processes to tackle class imbalance in rockfall candidates. However, due to the lack of further geometric understanding of the machine learning models, research has been focused on point-based deep learning approaches. Such efforts offer several advantages, such as higher accuracy and robustness, and local geometry understanding emphasizing their ability to detect rockfall candidates that may not be detected by traditional or machine learning methods (Farmakis et al., 2023).

In a closely related study to ours, Farmakis et al. (2022) proposed a method for rockfall detection using PointNet (Qi et al., 2017a) and PointNet++ (Qi et al., 2017b), achieving high accuracy on two different datasets, the CN Rail Ashcroft Mile 109.4 (Mile 109) and White Canyon (WCW). In addition, they extend their work using graph-based and convolutional-based models in similar set-ups showing promising results in detecting rockfalls (Farmakis et al., 2023). However, they rely on datasets containing nearly balanced instances of rockfall candidates in comparison to vegetation or other classes. Thus, despite the recent progress in point-based approaches for rockfall detection and susceptibility assessment, there are still several challenges that need to be addressed, such as robustness in detection accuracy with limited rockfall data samples. Also, it is worth noticing the significance of learning with class imbalance in earth vision tasks (Bai et al., 2023).

3 Our Approach

We formulate the problem as a 3D scene segmentation task and we detail our approach in two main phases. In the first phase, we create a densely labeled 3D scene representing the spatiotemporal distribution of the labeled points (see Sec. 3.1), and in the second one, we learn point-based models to segment the rockfall candidate areas (see Sec. 3.2).

3.1 Creation of a Densely Labeled 3D Scene

Starting from a small dataset of labeled data samples, i.e. sparsely labeled 3D point cloud, and multiple temporal 3D scans (3D scans acquired periodically), our goal is to create a densely labeled large-scale 3D scene by augmenting the labeled points of the sparse point cloud to form a larger set of points based on per-point spatiotemporal neighborhoods.

We do this by selecting and labeling points in the temporal point clouds with the label of their nearest point from the small set of labeled samples. For convention, we refer to the labeled points in the small set as *centroids*. Formally, we begin with a set of N labeled centroids, denoted by $\{\mathbf{c}_i \mid 1 \leq i \leq N\} = \{\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_N\}$. We denote the sequence of T temporal 3D point clouds as $\{\mathbf{P}^{(t)} \mid 1 \leq t \leq T\} = \{\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \ldots, \mathbf{P}^{(T)}\}$. Specifically, for each $\mathbf{c}_i$ with label l_i, we find the k closest points in each $\mathbf{P}^{(t)}$ and assign them the label l_i.

To efficiently find the k closest points to each centroid c_i, we represent each temporal point cloud $\mathbf{P}^{(t)}$ as a *KDTree*, a common way to represent point clouds (Zeng & Gevers, 2019), and we iteratively query the centroids. We use the Euclidean distance as our distance function d between the points and utilize a fixed radius sphere r to select the points around each centroid from the temporal point clouds. Precisely, given a radius r, we denote the set of points in $\mathbf{P}^{(t)}$ within distance r from $\mathbf{c}_i$ is as $\mathcal{Z}^{(t)}$, i.e. the points that satisfy the $d(\mathbf{c}_i, \mathbf{x}_j^{(t)}) < r$, where $\mathbf{x}_j^{(t)}$ is the j-th 3D point of $\mathbf{P}^{(t)}$. Then, we select k points from $\mathcal{Z}^{(t)}$ that are closest to $\mathbf{c}_i$, creating the following set of densely labeled points:

$$\mathcal{L}^{(t)} = \mathbf{x}_{j_1}^{(t)}, \mathbf{x}_{j_2}^{(t)}, \ldots, \mathbf{x}_{j_k}^{(t)}, \tag{1}$$

where $\mathbf{x}_{j_1}^{(t)}, \mathbf{x}_{j_2}^{(t)}, \ldots, \mathbf{x}_{j_k}^{(t)}$ are the k points in $\mathcal{Z}^{(t)}$ that minimize $d(\mathbf{c}_i, \mathbf{x}_j^{(t)})$ subject to the condition that $d(\mathbf{c}_i, \mathbf{x}_j^{(t)}) < r$. In simpler terms, we select k points from the temporal point clouds that are closest to each centroid $\mathbf{c}_i$, within a distance r from $\mathbf{c}_i$, and assign them the label of $\mathbf{c}_i$.

The labeled points from all the temporal point clouds are then combined to form a complete set of labeled points, denoted by $\mathcal{L} = \mathcal{L}^{(1)} \cup \mathcal{L}^{(2)} \cup \cdots \cup \mathcal{L}^{(T)}$. Please note that all sets in $\mathcal{L}$ contain unique points, i.e. there are no duplicate points or repetitions among them. Finally, we combine these labeled points to generate a densely labeled 3D scene. Then, for training purposes, we randomly sample with a fixed number of points the densely labeled scene to create multiple labeled instances.

3.2 3D Scene Segmentation

We handle the problem of rockfall detection as a 3D scene segmentation problem using per-point supervision. To facilitate the task of rockfall detection, we binarize the labels by encoding them as rockfall candidates (l_i: **1**) and non-candidate (l_i: **0**). As the utilized data has a significant class imbalance, i.e. the rockfall candidates are approximately 0.5% of the total labels, we use a weighted binary cross entropy loss function, where the weight assigned to each class is proportional to its inverse frequency. Specifically, the assigned weight to the rockfall class is set to 0.995, and the weight to the non-rockfall class is set to 0.005. This approach allows us to give more importance to the rare class, which leads to improved performance in similar tasks of class imbalance.

4 Implementation

This section describes the implementation of our approach for point-based rockfall detection.

4.1 Data Description

We utilize the "Degotalls E Section South" dataset from the from the study of Blanco et al. (2022). The dataset consists of 3D point clouds captured by a Terrestrial Laser Scanner, once a year over a period of nine years, from 2007 to 2015. The point clouds have an average number of points equal to 2.5 Million and a minimum distance between two points of 0.01 meters. In addition, we have a sparse set of 5970 labeled points (centroids) as *vegetation*, *rockfall candidate*, and *limit effect*. The set of labeled points refers to 2007-2009 and was labeled through a change detection and statistical analysis approach as explained in Blanco et al. (2022). The sparsely labeled point set has a minimum pairwise distance between points equal to 0.18 meters and a minimum distance between a rockfall and non-rockfall candidate equal to 0.43.

Figure 1: Densely labeled 3D scene and example instances. Following our notation, starting from left to right, we show the generated densely labeled scene $\mathcal{L}$, and three randomly sampled instances. Blue and red colors correspond to non-rockfall and rockfall candidates respectively.

To create a densely labeled large-scale 3D scene ($\mathcal{L}$) from the captured temporal 3D point clouds, as explained in Sec. 3.1, we arbitrarily select $k = 30$ closest points to each centroid and a fixed radius sphere with $r = 0.3$, i.e. smaller than the minimum distance between a rockfall and non-rockfall candidate in the initial labeled point set.

4.2 Training & Evaluation

We use PointNet (Qi et al., 2017a) and two variants of PointNet++ (Qi et al., 2017b), *Single Scale Grouping* and *Multi-Scale Grouping*, to learn global and local geometric point-wise features. For training, we create a dataset of 904 unique instances with 1024 points each by iteratively randomly sampling the generated densely labeled scene ($\mathcal{L}$). An example is shown in Fig. 1. Then, we randomly split these labeled instances into 80% training and 20% testing sets. We train the models for 100 epochs using Adam optimizer with a learning rate of 0.001 and a batch size of 8.

In imbalanced classification, traditional accuracy is not an appropriate metric as it does not account for class distribution, i.e. models can achieve high accuracy by favoring the majority class. For this, we evaluate our approach using Precision, Recall, F1-score, and Balanced Accuracy ($\mathbf{Acc_b}$) on the per-point labels. Please note that while we studied the same mountainous region in Spain as Zoumpekas et al. (2021) and Blanco et al. (2022), we cannot compare our method directly with theirs, because the acquired publicly available dataset from the provided links in Blanco et al. (2022), is a refined dataset and slightly different to the full dataset that is utilized in their studies. However, it is worth noticing that our findings align closely with theirs.

5 Results

In this section, we show our results and discuss the insights. We conduct the following ablation study. The initial sparsely labeled point set refers to cluster centroids labeled after a change detection and statistical analysis approach on snapshots captured in years 2007 and 2009 as explained in Blanco et al. (2022). The limited number of labeled samples in the initial sparsely labeled point set, i.e. 5970, makes it impractical for point-based modeling. For this, we augment the sparsely labeled points using neighbor points ($r = 0.3$ and $k = 30$) referring to the actual change detection and analysis years, i.e. 2007 and 2009. Then we follow the same process to create the dataset and train the networks, as described in Sec. 4.2. We compare the obtained results against using spatiotemporal augmentation in all the scans captured through the years 2007-2015.

In Table 1 we show the performance metrics in the test sets. In terms of precision, recall, and F1 score, we report the macro average, i.e. classes equally contribute to the average, and the weighted average, i.e. each class's contribution to the average is weighted by its number of samples. In rockfall detection, we mainly seek a model with high recall values, because recall measures the ability of a model to correctly identify all relevant instances of the positive class. Recall is particularly important when the cost of false negatives (miss-classifying a rockfall candidate) is high.

Method	**Acc$_b$**	**Precision**		**Recall**		**F1 Score**	
		M	**W**	**M**	**W**	**M**	**W**
$PointNet_{2007-2009}$	0.84	0.51	0.99	0.84	0.75	0.45	0.85
$PointNet_{2007-2015}$	0.94	0.53	0.99	0.94	0.90	0.53	0.94
$PointNet++(SSG)_{2007-2009}$	0.94	0.52	0.99	0.94	0.90	0.53	0.94
$PointNet++(SSG)_{2007-2015}$	0.96	0.54	0.99	0.96	0.92	0.56	0.95
$PointNet++(MSG)_{2007-2009}$	0.95	0.53	0.99	0.95	0.92	0.54	0.95
$PointNet++(MSG)_{2007-2015}$	**0.97**	**0.55**	0.99	**0.97**	**0.93**	**0.57**	**0.96**

Table 1: Performance Evaluation. **M** and **W** denote macro and weighted average respectively. Bold font corresponds to the best model.

Utilizing the spatiotemporal augmentation module improves the performance of all tested models. Among the models, Pointnet++ SSG and MSG had the best balanced accuracy and recall values, with the PointNet++ MSG model performing slightly better mainly due to its ability to capture multiscale patterns and learn multi-scale features.

6 Discussion and Future Work

This paper introduces an approach to intelligent rockfall detection, focusing on identifying areas prone to rockfall rather than predicting specific events. Leveraging spatiotemporal neighborhoods of points, the study demonstrates promising performance in detecting rockfall events. Nonetheless, there are certain limitations. First, our approach for detecting rockfall candidate areas is dependent on the availability of high-quality labeled data (initial set). Second, our approach assumes that rockfall candidate areas can be detected using only their coordinate information, i.e. x, y, and z. Also, we inherently assume that areas susceptible to rockfall remain consistent over time. To address these, future research directions include exploring unsupervised learning approaches, such as Xie et al. (2020); Zhang et al. (2023); Long et al. (2023), for regions lacking labeled data, considering the dynamic nature of geological processes in detection methods, and incorporating additional data from various sources such as Synthetic Aperture Radar (SAR) data and weather-related data.

Moreover, in this study, we evaluated three neural networks that employ point-wise operations. It is worth noting, however, that there are other architectures, such as voxel-based (Liu et al., 2019) or graph-based (Wang et al., 2019), and thus, further experimentation is recommended. Additionally, in our spatiotemporal augmentation module different values of radius, r, and neighbors, k, should be investigated to better understand the typical size and distribution of rockfall patches. Also, instead of using random sampling to create multiple instances, alternative sampling techniques like furthest point sampling or spatially organized sampling could be utilized. Finally, a potential next step is to learn models with further per-point features such as the intensity from the captured 3D scan. Including intensity information from the 3D scans can provide richer information (Reymann & Lacroix, 2015). This step could lead to more comprehensive models capable of capturing further details and potentially improving overall performance and robustness.

7 Conclusion

In conclusion, our spatiotemporal point-based approach offers a promising solution for rockfall detection in cases where labeled data instances are limited and have a significant class imbalance. In addition, point-based neural networks show high potential in understanding the highly complex geometry of cliffs and mountainous areas facilitating downstream tasks such as rockfall detection. Finally, improving the performance and reliability of rockfall detection and prediction mechanisms enhances the safety of infrastructure and communities in mountainous regions.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860843.

REFERENCES

Antonio Abellán, Thierry Oppikofer, Michel Jaboyedoff, Nicholas J. Rosser, Michael Lim, and Matthew J. Lato. Terrestrial laser scanning of rock slope instabilities. Earth Surface Processes and Landforms, 39(1):80–97, 2014. doi: 10.1002/esp.3493.

Federico Agliardi, Giovanni B. Crosta, and Paolo Frattini. Integrating rockfall risk assessment and countermeasure design by 3d modelling techniques. Natural Hazards and Earth System Sciences, 9(4):1059–1073, 2009. doi: 10.5194/nhess-9-1059-2009.

Massimiliano Alvioli, Michele Santangelo, Federica Fiorucci, Mauro Cardinali, Ivan Marchesini, Paola Reichenbach, Mauro Rossi, Fausto Guzzetti, and Silvia Peruccacci. Rockfall susceptibility and network-ranked susceptibility along the italian railway. Engineering Geology, 293:106301, 2021. doi: 10.1016/j.enggeo.2021.106301.

Yuchen Bai, Jean-Baptiste Durand, Grégoire Laurent Vincent, and Florence Forbes. Semantic segmentation of sparse irregular point clouds for leaf/wood discrimination. In NeurIPS, pp. 48293–48313, 2023.

Rushikesh Battulwar, Ebrahim Emami, Masoud Zare Naghadehi, and Javad Sattarvand. Automatic extraction of joint orientations in rock mass using pointnet and dbscan. In Advances in Visual Computing, pp. 718–727, 2020. doi: 10.1007/978-3-030-64559-5_57.

Laura Blanco, David García-Sellés, Marta Guinau, Thanasis Zoumpekas, Anna Puig, Maria Salamó, Oscar Gratacós, Josep Anton Muñoz, Marc Janeras, and Oriol Pedraza. Machine learning-based rockfalls detection with 3d point clouds, example in the montserrat massif (spain). Remote Sensing, 14(17):4306, 2022. doi: 10.3390/rs14174306.

Franck Bourrier, Luuk Dorren, François Nicot, Frédéric Berger, and Félix Darve. Toward objective rockfall trajectory simulation using a stochastic impact model. Geomorphology, 110(3):68–79, 2009. doi: 10.1016/j.geomorph.2009.03.017.

Ioannis Farmakis, Paul-Mark DiFrancesco, D. Jean Hutchinson, and Nicholas Vlachopoulos. Rockfall detection using lidar and deep learning. Engineering Geology, 309:106836, 2022. doi: 10.1016/j.enggeo.2022.106836.

Ioannis Farmakis, D. Jean Hutchinson, Nicholas Vlachopoulos, Matthew Westoby, and Michael Lim. Slope-scale rockfall susceptibility modeling as a 3d computer vision problem. Remote Sensing, 15(11):2712, 2023. doi: 10.3390/rs15112712.

Ilenia G. Gallo, Mónica Martínez-Corbella, Roberto Sarro, Giulio Iovine, Juan López-Vinielles, Mario Hérnandez, Gaetano Robustelli, Rosa María Mateos, and Juan Carlos García-Davalillo. An integration of uav-based photogrammetry and 3d modelling for rockfall hazard assessment: The cárcavos case in 2018 (spain). Remote Sensing, 13(17):3450, 2021. doi: 10.3390/rs13173450.

Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. TPAMI, 43(12):4338–4364, 2021. doi: 10.1109/TPAMI.2020.3005434.

Michel Jaboyedoff, Thierry Oppikofer, Antonio Abellán, Marc-Henri Derron, Alex Loye, Richard Metzger, and Andrea Pedrazzini. Use of lidar in landslide investigations: a review. Natural Hazards, 61:5–28, 2012. doi: 10.1007/s11069-010-9634-2.

Dimitri Lague, Nicolas Brodu, and Jérôme Leroux. Accurate 3D comparison of complex topography with terrestrial laser scanner: Application to the Rangitikei canyon (N-Z). ISPRS Journal of Photogrammetry and Remote Sensing, 82:10–26, 2013. doi: 10.1016/j.isprsjprs.2013.04.009.

Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. In NeurIPS, 2019. doi: 10.5555/3454287.3454374.

Fuchen Long, Ting Yao, Zhaofan Qiu, Lusong Li, and Tao Mei. Pointclustering: Unsupervised point cloud pre-training using transformation invariance in clustering. In CVPR, pp. 21824–21834, 2023. doi: 10.1109/CVPR52729.2023.02090.

Charles R. Qi, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In CVPR, pp. 77–85, 2017a. doi: 10.1109/CVPR.2017.16.

Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In NeurIPS, pp. 5105–5114, 2017b.

Christophe Reymann and Simon Lacroix. Improving lidar point cloud classification using intensities and multiple echoes. In IROS, pp. 5122–5128, 2015. doi: 10.1109/IROS.2015.7354098.

Guruh Samodra, Guangqi Chen, Junun Sartohadi, Danang Sri Hadmoko, Kiyonobu Kasama, and Muhammad Anggri Setiawan. Rockfall susceptibility zoning based on back analysis of rockfall deposit inventory in Gunung Kelir, Java. Landslides, 13(4), 2016. doi: 10.1007/s10346-016-0713-7.

Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. TOG, 38(5), 2019. doi: 10.1145/3326362.

Luke Weidner, Gabriel Walton, and Ryan Kromer. Classification methods for point clouds in rock slope monitoring: A novel machine learning approach and comparative analysis. Engineering Geology, 263:105326, 2019. doi: 10.1016/j.enggeo.2019.105326.

Matthew Westoby, James Brasington, Neil F. Glasser, Michael J. Hambrey, and J.M. Reynolds. Structure-from-motion photogrammetry: A low-cost, effective tool for geoscience applications. Geomorphology, 179:300–314, 2012. doi: 10.1016/j.geomorph.2012.08.021.

Lukas Winiwarter, Katharina Anders, and Bernhard Höfle. M3C2-EP: Pushing the limits of 3D topographic point cloud change detection by error propagation. ISPRS Journal of Photogrammetry and Remote Sensing, 178, 2021. doi: 10.1016/j.isprsjprs.2021.06.011.

Saining Xie, Jiatao Gu, Demi Guo, Charles R. Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In ECCV, 2020. doi: 10.1007/978-3-030-58580-8_34.

Wei Zeng and Theo Gevers. 3dcontextnet: K-d tree guided hierarchical learning of point clouds using local and global contextual cues. In ECCV Workshops, 2019. doi: 10.1007/978-3-030-11015-4_24.

Zihui Zhang, Bo Yang, Bing Wang, and Bo Li. Growsp: Unsupervised semantic segmentation of 3d point clouds. In CVPR, June 2023. doi: 10.1109/CVPR52729.2023.01690.

Thanasis Zoumpekas, Anna Puig, Maria Salamó, David Garcia-Sellés, Laura Blanco Nuñez, and Marta Guinau. An intelligent framework for end-to-end rockfall detection. International Journal of Intelligent Systems, 36(11):6471–6502, 2021. doi: 10.1002/int.22557.

SUPPLEMENTARY MATERIAL FOR SPATIOTEMPORAL ROCKFALL DETECTION USING POINT-BASED NEURAL NETWORKS

Thanasis Zoumpekas, Anna Puig, Maria Salamó
Department of Mathematics and Computer Science
University of Barcelona, Spain
{thanasis.zoumpekas,annapuig,maria.salamo}@ub.edu

ABSTRACT

In this supplementary document, we first present additional details regarding the utilized loss function for the learning of the selected point-based neural networks (Sec. 1). Then, we provide further results and visualizations (Sec. 2).

1 LOSS FUNCTION

To address the class imbalance in the utilized data, we use a weighted binary cross entropy loss function, where the weight assigned to each class is proportional to its inverse frequency. Specifically, the assigned weight to the rockfall class is set to 0.995, and the weight to the non-rockfall class is set to 0.005.

Formally, the weighted binary cross entropy loss (negative log-likelihood) function can be expressed in Eq. 1, as follows:

$$\mathcal{L} = -\frac{1}{N}\sum_{i=1}^{N}\Big[w_i y_i \log(p_i) + (1-y_i)\log(1-p_i)\Big] \tag{1}$$

where N is the total number of samples, y_i is the ground truth label of the i^{th} point, p_i is the predicted probability that the i^{th} point belongs to the rockfall class, and w_i is the weight assigned to the i^{th} point.

The weight for each data point is given by Eq. 2. This approach allows us to give more importance to the rare class, which can lead to improved performance in cases of class imbalance.

$$w_i = \begin{cases} 0.995 & y_i = 1 \\ 0.005 & y_i = 0 \end{cases} \tag{2}$$

2 RESULTS & VISUALIZATIONS

This section presents further results and visualizations facilitating the understanding and interpretation of rockfall detection as a 3D scene segmentation problem. In Figs. 1, 2, 3, we show visualization results using PointNet, PointNet++ Single-Scale Grouping (SSG) and Multi-Scale Grouping (MSG) respectively. Please note that the *Ground Truth Scan* refers to the test set acquired after the creation of the densely labeled 3D scene using our spatiotemporal augmentation module in all temporal 3D point clouds, i.e. referring to years 2007-2015. Following, we observe the segmentation result and we provide additional visualizations of the per-point classification errors and per-point classification errors concerning only rockfall candidates.

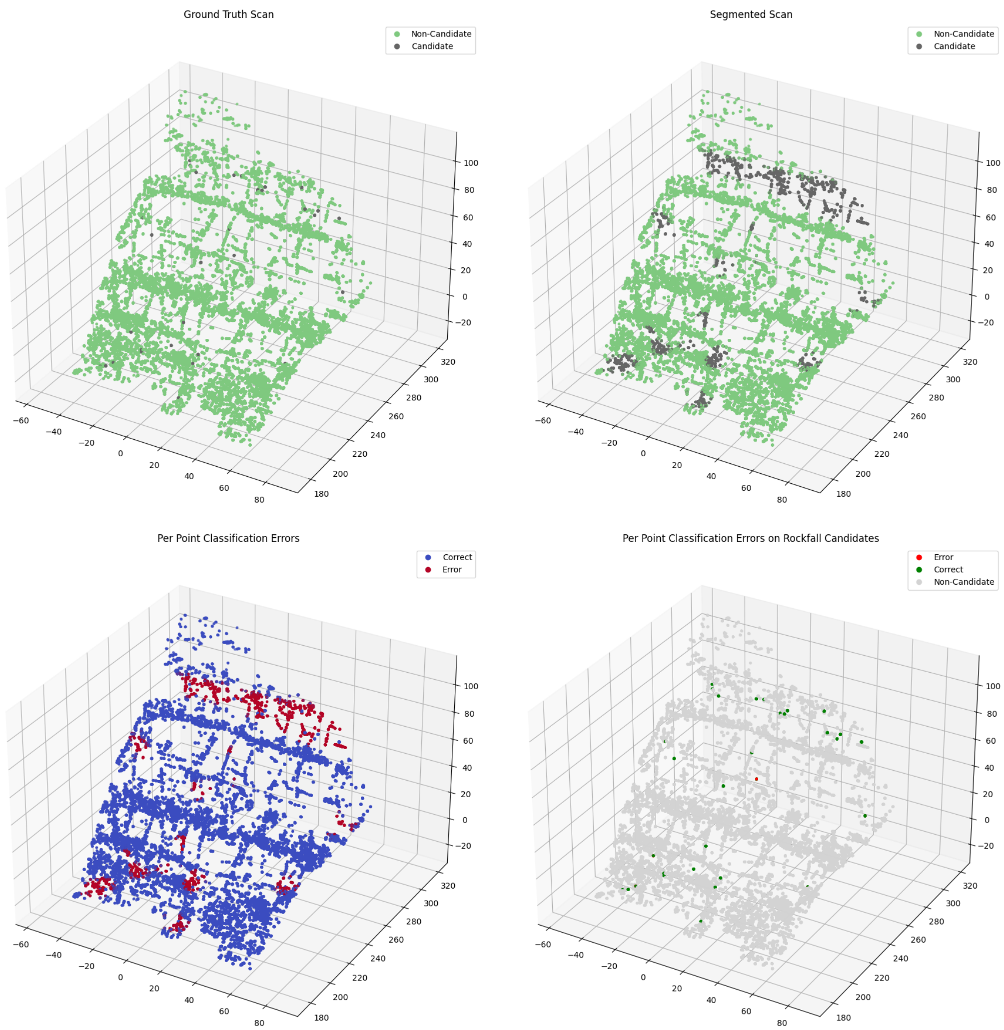

Figure 1: PointNet Results.

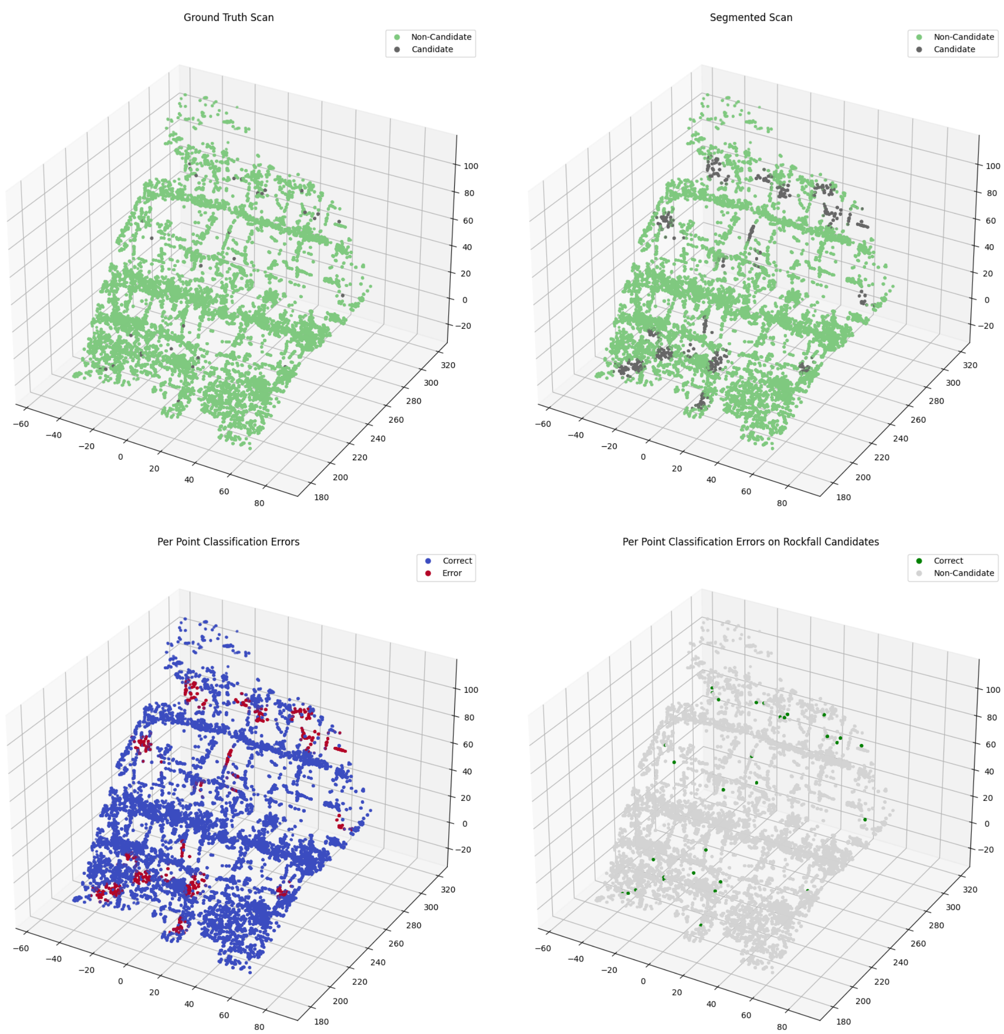

Figure 2: Pointnet++ SSG Results.

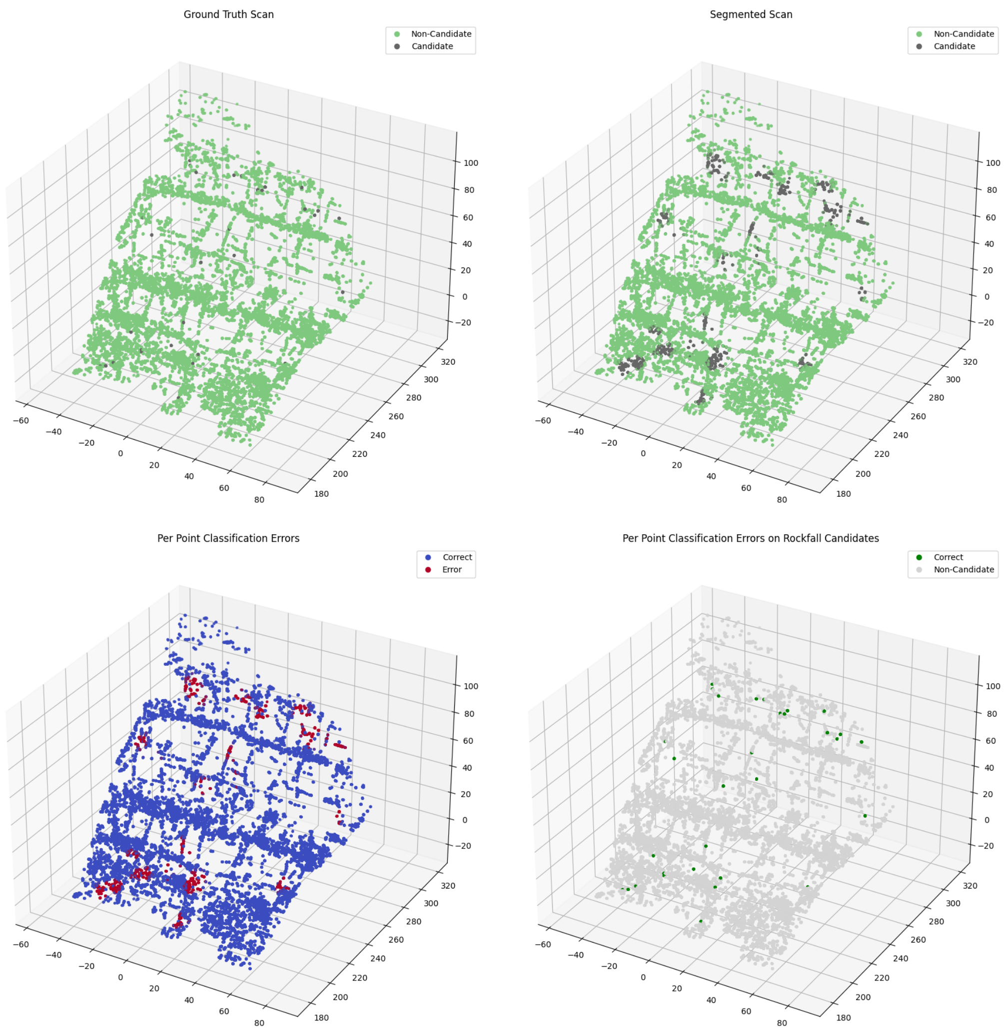

Figure 3: Pointnet++ MSG Results.