



ARTICLE

A Point Cloud Segmentation Method for Adhering Coal Piles in Industrial Coal Sheds

Ping Zhang¹, Chungang Li², Jian Song¹, Jianjun Tan¹, Han Tian¹, Le Chen³ and Yongkun Song^{3,*}

¹Fuel Business Division, Hunan Datang Xianyi Technology Co., Ltd., Changsha, China

²Business Management Department, China Datang Corporation Ltd., Beijing, China

³School of Computer Science and Technology, Changsha University of Science and Technology, Changsha, China

*Corresponding Author: Yongkun Song, Email: songyk1118@163.com

Received: 10 February 2026; Accepted: 27 March 2026; Published: 02 July 2026

ABSTRACT: Accurate segmentation of coal pile point clouds is essential for automated inventory management and safety supervision in industrial coal sheds. With the widespread adoption of LiDAR technology, massive point cloud data are available, yet traditional clustering algorithms often fail to distinguish physically adhering coal piles in high-density storage scenarios, resulting in severe under-segmentation. To address this challenge, this paper proposes an automated point cloud segmentation framework tailored for complex industrial environments. The proposed framework consists of three core stages: data preprocessing, ground segmentation, and instance segmentation. First, relative elevation mapping and selective statistical filtering are employed to normalize coordinate baselines and suppress noise. Then, a Normal-Constrained Random Sample Consensus (NC-RANSAC) algorithm is introduced to separate horizontal ground surfaces from pile slopes using geometric angular constraints. Finally, an unsupervised strategy based on Principal Component Analysis (PCA) and sliding-window concavity detection is developed to identify saddle points and decouple adhering instances. Experimental results on real-world industrial datasets demonstrate that the NC-RANSAC algorithm exhibits high parameter stability and effectively resists interference from non-ground planes. Furthermore, the proposed instance segmentation strategy successfully decouples adhering coal piles with varying morphologies without manual intervention or pre-trained models.

KEYWORDS: Point cloud segmentation; industrial coal sheds; adhering coal piles

1 Introduction

Coal, as a key non-renewable energy resource, plays a vital role in global energy consumption, especially in power generation and industry [1]. During production, storage, and transport, coal piles are the most common accumulation form, and their volume, distribution, and geometry directly impact inventory management, production scheduling, and safety supervision. With advances in LiDAR and UAV surveying, point clouds have become essential for representing coal pile 3D geometry [2]. Efficient and accurate segmentation of coal pile point clouds is therefore critical for improving yard management and safety oversight.

Traditional coal stockpile inventory practices mainly rely on manual total station measurements or belt scale weighing, which are often inefficient, cover limited areas, have restricted accuracy, and involve significant safety risks. Alsayed and Nabawy [3] conducted a systematic review of pile measurement and data acquisition technologies, noting that manual methods struggle to meet the demands of modern coal

yards, where stockpiles are irregular and dynamically changing. To overcome these limitations, UAV-based photogrammetry has been widely adopted due to its non-contact operation and wide coverage. However, in fully or partially enclosed coal sheds, poor lighting and high dust concentrations often degrade vision-based reconstruction, making it difficult to obtain high-accuracy 3D models. In contrast, LiDAR, with its active sensing and partial penetration capability, is largely insensitive to ambient lighting and has become the primary technology for point cloud acquisition in complex indoor environments [4]. Despite the high quality of LiDAR data, automatically extracting and segmenting individual coal pile instances from massive, unorganized point clouds remains a challenging problem. Existing approaches can generally be divided into deep learning-based methods and traditional geometry-based methods.

Deep learning-based point cloud segmentation methods typically use supervised learning, training neural networks on labeled point clouds for point- or region-level classification. Methods such as PointCNN [5], PointNet [6], and PointNet++ [7] achieve high accuracy on complex pile geometries, boundaries, and occlusions, enabling end-to-end automatic segmentation. However, they rely on large manually annotated datasets, which are costly and time-consuming to obtain in industrial scenarios where coal pile geometries frequently change. Model generalization is often limited to specific scenes, restricting rapid “plug-and-play” deployment. Liu et al. [8] further emphasized that, despite advances in pose estimation, defect inspection, and volume measurement, efficiently processing large, complex point clouds in real industrial production remains a major challenge.

To reduce reliance on labeled data, many studies have adopted unsupervised geometric methods. Wang et al. [9] proposed an automated extraction pipeline combining cloth simulation filtering with Euclidean clustering using flight trajectories, while Yang et al. [10] used multi-scale directional curvature to detect pile boundaries. However, in high-density scenarios where piles are physically connected at their bases, standard Euclidean clustering often merges multiple piles into a single component, causing under-segmentation. Curvature-based methods can partially mitigate this, but their reliance on complex local differential computations makes them inefficient for million-level point clouds, limiting rapid on-site processing.

In summary, current mining and coal yard operations urgently require an efficient algorithm that can segment physically connected coal piles without relying on pre-trained models. To address this, we propose a coal pile point cloud segmentation framework for industrial coal sheds. Unlike Wang et al. [9], which relies on trajectory-intersection filtering, our method uses geometric concavity as a topological criterion to exploit the natural physical characteristics of coal piles, enabling fine-grained decoupling of adhered piles and accurate volume estimation under unsupervised conditions. Experiments on real-world coal yard datasets show that the framework can robustly generate independent pile-level point clouds, supporting pile-wise volume computation and refined inventory management in practice.

2 Related Works

Regarding coal pile segmentation, extensive research has been conducted by various researchers. In early studies, the segmentation of piled objects primarily relied on the geometric and topological characteristics of point clouds, among which ground separation constituted one of the fundamental preprocessing steps for target segmentation. Cheng et al. [11] proposed a simple ground segmentation method for LiDAR-based three-dimensional point clouds, in which ground and non-ground points are distinguished using statistical and topological features, providing a reliable preprocessing basis for subsequent pile boundary detection and object separation. In addition to ground removal, some studies employed morphological filtering [12] for noise suppression and local smoothing of point clouds, thereby enhancing the stability and accuracy of subsequent segmentation algorithms.

After ground removal, region growing methods are widely used for precise pile segmentation. Luo et al. [13] proposed a supervoxel-based approach, voxelizing point clouds and growing regions based on fitting residuals and local geometric similarity to achieve smooth, accurate boundaries. Chen et al. [14] introduced a DBSCAN-based 3D boundary detection method, enabling unsupervised segmentation without prior models. To improve DBSCAN's efficiency on large-scale point clouds, Zhang et al. [15] proposed class-constrained DBSCAN, incorporating category constraints for faster and more stable clustering. In complex natural environments, variations in point density and noise can degrade clustering. Ma et al. [16] enhanced DBSCAN for landslide detection in LiDAR point clouds, improving robustness and segmentation accuracy. These studies show that density-based unsupervised clustering is effective for extracting complex piles or irregular point cloud objects.

In engineering applications, point cloud segmentation is typically a prerequisite for 3D reconstruction and dimensional measurement of piles, including ore, coal, aggregate, and other bulk material stockpiles. Yang et al. [10] proposed an automated pile extraction and 3D volume measurement method, which separates piles from 3D point clouds using multi-scale directional curvature analysis and competitive region growing, then computes volumes via mesh reconstruction and triangular prism calculation, achieving high accuracy on both real and synthetic datasets. Ling et al. [17] introduced the DIVESPOT method, which performs pile separation through point cloud pose correction and hierarchical density clustering, followed by volume estimation. This unsupervised approach requires no labeled data and maintains low volume error in complex outdoor environments, demonstrating its practicality for real engineering scenarios.

In recent years, end-to-end deep learning methods have become a major focus for point cloud segmentation. Unlike traditional approaches relying on manually designed geometric features or clustering rules, these methods automatically learn spatial structures of point clouds, enabling direct mapping from raw data to segmentation results. PointNet [6] first introduced a network capable of processing unordered point clouds for end-to-end semantic segmentation, while PointNet++ [7] incorporated hierarchical feature learning and multi-scale neighborhood modeling to capture local geometry, improving accuracy in complex pile scenarios. DGCNN [18] further enhanced local structure and boundary representation via dynamic graph convolutions. Although these methods achieve high segmentation accuracy, they require large labeled datasets and incur high computational costs. In industrial environments, frequent changes in coal pile geometry make labeling costly and model generalization limited, hindering rapid deployment. In contrast, clustering-based unsupervised methods can quickly segment piles without supervision, remaining valuable for efficient processing of large-scale coal pile point clouds and supporting practical volume estimation and management.

3 Proposed Framework

As shown in Fig. 1, the proposed framework automates point cloud processing for industrial coal sheds through three core stages. First, data preprocessing unifies coordinates and reduces noise via relative elevation calculation, selective statistical filtering, and voxel grid downsampling. Second, ground segmentation employs Normal-Constrained RANSAC to separate horizontal ground from pile slopes using normal vector angular constraints. Finally, adhering pile instance segmentation addresses physical adhesion using PCA principal direction projection and sliding window concavity detection, automatically decoupling connected domains into independent instances by constructing segmentation planes at geometric saddle points.

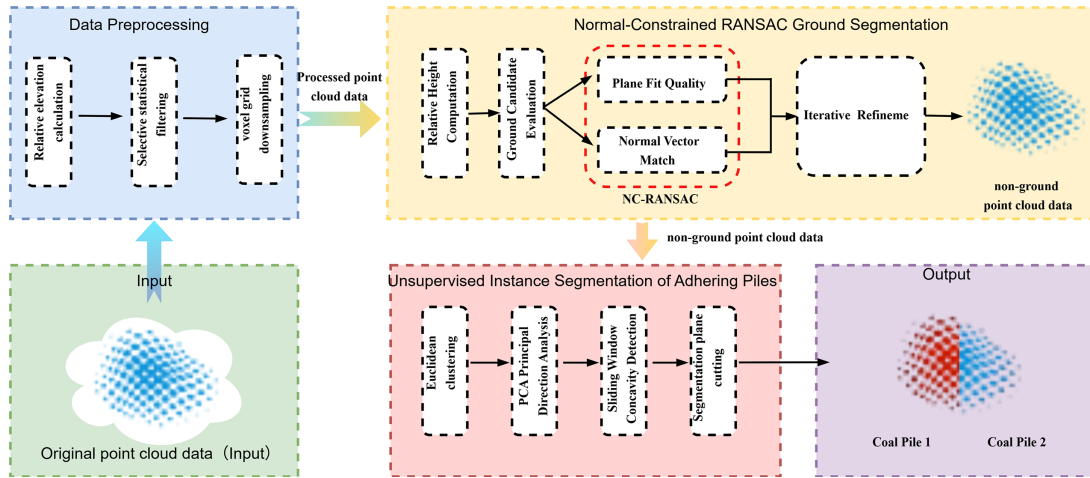


Figure 1: Coal pile point cloud segmentation framework for industrial coal shed scenarios.

3.1 Data Preprocessing

3.1.1 Relative Elevation Calculation and Selective Statistical Filtering

The absolute coordinates of the raw point cloud data $P = \{p_1, p_2, \dots, p_n\}$ are affected by the installation position of the acquisition equipment, making it inconvenient to set a unified threshold. First, the point cloud is traversed to calculate the global minimum point z_{\min} , and the Z -axis coordinates of all points are mapped to the relative ground height h_i :

$$h_i = p_i.z - z_{\min} \quad (1)$$

Considering that the main noise in the coal mine site originates from ground dust or underground reflection noise, and that the sparse points at the top of the coal pile would affect volume calculation if over-smoothed, this paper adopts a selective filtering strategy. A height threshold $H_{th} = 0.15$ is set to divide the point cloud into a near-ground set P_{low} and a pile set P_{high} . This threshold was empirically determined to balance effective ground noise removal and preservation of sparse pile points. Statistical Outlier Removal (SOR) is performed only on P_{low} . For any point p_i in P_{low} , the average distance \bar{d}_i to its k nearest neighbors is calculated. If the following formula is satisfied, it is determined to be a noise point and removed:

$$\bar{d}_i > \mu_k + \alpha \cdot \sigma_k \quad (2)$$

where μ_k and σ_k are the mean and standard deviation of the average distances, respectively, and α is the standard deviation multiplier threshold. The processed P'_{low} is merged with the unprocessed P_{high} , which not only purifies the ground noise but also fully preserves the geometric details of the coal piles.

3.1.2 Voxel Grid Downsampling

To address the issue of slow processing of million-scale point cloud data, an improved voxel filtering method is employed to sparsify the original point cloud. This approach reduces the number of points while preserving the key geometric features of the coal pile surface. The initial voxel size n was determined based on the spatial density ρ of the point cloud to ensure that the downsampling scale matches the data resolution. Specifically, the voxel size is calculated as:

$$\mathbf{n} = \frac{1}{\sqrt[3]{\rho}} \quad (3)$$

For each voxel grid, instead of retaining all points, the centroid \bar{p} of all points p_j within the grid is calculated as the unique representative point of that voxel:

$$\bar{p} = \frac{1}{N} \sum_{j=1}^N p_j \quad (4)$$

This step significantly reduces the data volume while also smoothing out minute sensor noise, tending to make the point cloud density distribution uniform. It is worth noting that the downsampling process changes the topological structure and index order of the point cloud. To ensure the accuracy of subsequent ground segmentation, geometric attribute updates must be performed on the downsampled point cloud. The algorithm traverses the processed point cloud again to recalculate the relative elevation information h'_i for all points. This mechanism can eliminate potential biases in elevation judgment caused by downsampling, providing a precise data foundation for subsequent ground segmentation.

3.2 Normal-Constrained RANSAC Ground Segmentation

Targeting the differences in spatial geometric features between the ground and pile side slopes in coal pile scenarios, this paper adopts the Normal-Constrained RANSAC (NC-RANSAC) algorithm for ground segmentation, operating on the preprocessed relative-height coordinates obtained in [Section 3.1](#). Traditional RANSAC mainly relies on the Euclidean distance from points to the plane as the criterion for inlier determination, while NC-RANSAC adds a normal vector angle criterion on this basis, distinguishing horizontal ground from gentle slopes by limiting the inclination of the fitted plane.

The algorithm assumes the ground follows the plane model $ax + by + cz + d = 0$. In each iteration, a subset is randomly selected to fit the plane model, and the normal vector $\vec{n} = (a, b, c)$ of the model is calculated. To distinguish between the ground and coal pile slopes, the angle θ between \vec{n} and the standard vertical axis $\vec{v} = (0, 0, 1)$ is calculated:

$$\theta = \arccos \left(\frac{|\vec{n} \cdot \vec{v}|}{|\vec{n}| \cdot |\vec{v}|} \right) \quad (5)$$

Only when $\theta < \theta_{th}$ is the plane considered a potential ground candidate. Furthermore, to thoroughly clear uneven ground, the algorithm employs a multi-round iterative stripping strategy: Once a compliant ground plane is identified, its inliers are removed, and RANSAC detection is repeated on the remaining point cloud until the number of remaining points is insufficient or no plane satisfying the angular constraint is detected.

3.3 Unsupervised Instance Segmentation of Adhering Piles

3.3.1 PCA-Based Principal Direction Projection Analysis

In actual scenarios, adjacent coal piles may be identified as the same connected domain in Euclidean clustering due to close proximity, resulting in under-segmentation phenomena. To address this issue, this paper proposes a projection strategy based on Principal Component Analysis (PCA).

The point cloud after ground removal is mainly composed of discrete coal piles. First, a k-d tree is utilized to establish a spatial index, and the Euclidean clustering extraction algorithm is adopted. A distance threshold $d_{cluster}$ is set; if the Euclidean distance between point p_a and p_b , $|p_a - p_b|_2 < d_{cluster}$, they are merged into

the same class. This step outputs the initial cluster set $C = \{C_1, C_2, \dots, C_k\}$. Next, PCA is performed on each cluster C_k . The covariance matrix is calculated and decomposed to obtain eigenvectors, and the eigenvector $\overrightarrow{e_{main}}$ corresponding to the largest eigenvalue is selected as the principal axis direction of the coal pile. This vector indicates the direction of longest extension of the coal pile, providing a key geometric benchmark for subsequent slice analysis along the coal pile “ridge line”. Subsequently, all points p_i within the cluster are projected onto the principal axis direction to obtain a one-dimensional projection scalar value s_i :

$$s_i = (p_i - \overline{p_{centroid}}) \cdot \overrightarrow{e_{main}} \quad (6)$$

where $\overline{p_{centroid}}$ is the geometric center point of the current cluster. By reducing the three-dimensional morphology to a one-dimensional distribution along the principal axis direction, a computational basis is provided for subsequently identifying the geometric concavity features at the adhesion points.

3.3.2 Sliding Window Concavity Detection and Cutting

To find the adhesion point between two coal piles, the algorithm first projects the point cloud cluster onto its PCA principal axis. The algorithm divides the value range of the projection scalar s_i into N bins (typically 50–300 depending on the cluster length), and calculates the average height $\overline{z_{bin}}$ of all points within each bin, where $\overline{z_{bin}}$ refers to the original z-coordinate of the points. Then we constructs a height profile histogram along the principal axis direction. A sliding window of width W bins is applied along the histogram to detect local troughs. For the current window W_i , the difference between its height and the average height of the adjacent front and rear windows is calculated to define the concavity D :

$$D = \frac{\overline{z_{front}} + \overline{z_{back}}}{2} - \overline{z_{current}} \quad (7)$$

If D is greater than the set tolerance threshold $T_{concavity}$ and the absolute depth satisfies physical constraints, the location is determined to be a concave region caused by adhesion. The algorithm selects the position with the maximum concavity to construct a vertical segmentation plane and construct a vertical segmentation plane:

$$n * (p - p_0) = 0 \quad (8)$$

where p_0 is the center of the bin with maximum concavity, n is perpendicular to the PCA principal axis, and p represents any point in the cluster. The plane cuts the connected cluster into two independent coal pile instances, achieving refined instance-level segmentation.

4 Experiments and Analysis

4.1 Experimental Setup

4.1.1 Dataset

To evaluate the reliability of the proposed framework in complex industrial environments and the segmentation accuracy of adhering piles, we employ a dataset collected from a large enclosed coal shed using a LiDAR scanner. This dataset is primarily used to verify the qualitative segmentation performance and environmental adaptability in practical engineering applications.

4.1.2 Baselines and Validation Strategy

To verify the effectiveness of each module, a two-dimensional validation strategy is designed:

- **Comparison and Parameter Sensitivity Validation for Ground Segmentation:** To demonstrate the advantage of the proposed NC-RANSAC over conventional plane fitting methods, a comparative experiment with standard RANSAC is conducted, evaluating both visual results. Additionally, due to the lack of pixel-level ground truth in industrial sites and the difficulty in quantifying minute errors, parameter sensitivity analysis is adopted. By adjusting the normal constraint threshold (θ) over a wide range and introducing Relative IoU and the number of extracted ground points as metrics, the study aims to reveal the stability and adaptability of the algorithm against sensor installation errors and complex environmental disturbances.
- **Feasibility Validation for Adhering Pile Segmentation:** This focuses on verifying the algorithm's ability to capture non-convex geometric features. Typical physically adhering pile samples from real scenes are selected for visual analysis of concavity features. The experiment aims to demonstrate how the concavity extraction algorithm precisely locates "saddle point" regions without manual intervention, and to verify the feasibility and accuracy of the manifold cutting strategy in decoupling continuous connected domains into independent coal pile instances.

4.2 Ground Segmentation Visualization Comparison

- Ground segmentation is a critical prerequisite for adhering pile instance segmentation, as the accurate separation of ground points from pile points directly affects the completeness and accuracy of subsequent instance segmentation. In this study, we employ the NC-RANSAC for ground segmentation, and compare it with the standard RANSAC method. The visualization results are presented in Fig. 2, where Fig. 2a depicts the original point cloud of the coal pile.
- As shown in Fig. 2b, the standard RANSAC tends to misclassify points at the bottom of adhering piles, especially in regions where the pile surface transitions smoothly to the ground, as ground points, resulting in incomplete pile structures. This is particularly evident in the connection point areas where the pile and ground meet at shallow angles. In contrast, as illustrated in Fig. 2c, the proposed NC-RANSAC outperforms the standard RANSAC. By incorporating normal vector constraints, NC-RANSAC effectively preserves the integrity of pile bottoms and connection points, even in scenes with uneven terrain or diverse pile morphologies. Ground points are accurately removed while the complete pile structures are well retained.




Figure 2: Comparison of ground segmentation visualization results. (a) Original point cloud. (b) Ground segmentation result obtained by the standard RANSAC method. (c) Ground segmentation result obtained by the proposed NC-RANSAC method.

4.3 Parameter Sensitivity and Stability Analysis of Normal Constraints

This section investigates the impact of the core parameter the normal vector angle threshold (θ) of the NC-RANSAC algorithm—on ground segmentation and assesses stability against minute environmental deviations. Due to the absence of pixel-level ground truth, we adopted relative convergence validation. If the segmentation result ceases to fluctuate with parameter changes beyond a certain range, it is considered to have reached a stable convergence state. The ground segmentation result at $\theta = 5^\circ$ is selected as the baseline reference set P_{ref} . For any given threshold θ , the segmented ground point set is denoted as P_θ . The Relative Intersection over Union (Relative IoU) is defined as:

$$Relative\ IoU_\theta = \frac{|P_\theta \cap P_{ref}|}{|P_\theta \cup P_{ref}|} \quad (9)$$

where $|P_\theta \cap P_{ref}|$ represents the number of correctly segmented baseline ground points under the current parameter (the intersection), and $|P_\theta \cup P_{ref}|$ represents the total number of points combining the current and baseline results (the union). By calculating the set difference and degree of overlap between P_θ and P_{ref} , this metric objectively reflects the performance evolution of the algorithm under varying constraint conditions.

Using the Set Subtraction Method, ground point clouds under different parameters are obtained via segmentation. Relative IoU and the number of segmented ground points are used as metrics to analyze performance changes within the parameter domain. Fig. 3 focuses on the parameter interval ($0^\circ \sim 5^\circ$) where critical performance changes occur, and Table 1 details the statistical data.

Experimental results reveal three significant performance stages under different constraint intensities.

- (1) The Over-constrained Failure Stage ($0^\circ \leq \theta < 1^\circ$). When $\theta = 0^\circ$, the algorithm mandates the ground normal vector to be absolutely vertical. Data shows IoU is only 0.1619, and average extracted points drop to 1930 (13.3% of normal), with a standard deviation of 0.0661. This quantitatively reflects minute mechanical errors in LiDAR installation and inherent micro-undulations of the site ground. As actual ground normals cannot satisfy the idealized mathematical condition of 0° , severe ground under-segmentation occurs.
- (2) The Performance Recovery and Phase Transition Stage ($1^\circ \leq \theta < 3^\circ$). When the threshold is fine-tuned to 1° , metrics surge significantly: IoU recovers to 0.9765, and extracted points rise to 14,139. This steep performance curve proves that actual ground inclination deviations in the experiment concentrate within the minute interval of $(0^\circ, 1^\circ]$. With just 1° of tolerance, the algorithm effectively overcomes environmental errors and recaptures most ground features.
- (3) The Performance Saturation and Stability Stage ($\theta \geq 3^\circ$). As θ increases to 3° and above, IoU remains constant at 1.0000, and standard deviation drops to 0.0000. This indicates the algorithm has entered a saturation zone, exhibiting high stability. In this interval, consistency is no longer disturbed by parameter fluctuations; regardless of threshold relaxation, the same ground region is completely and precisely segmented. Based on the above analysis, the algorithm maintains stable segmentation performance within the wide parameter domain of $\theta \geq 3^\circ$, indicating insensitivity to threshold selection and good engineering adaptability. In practical applications, the threshold can be flexibly set within the stable interval according to specific scenarios, effectively tolerating terrain undulations while maintaining operational stability.

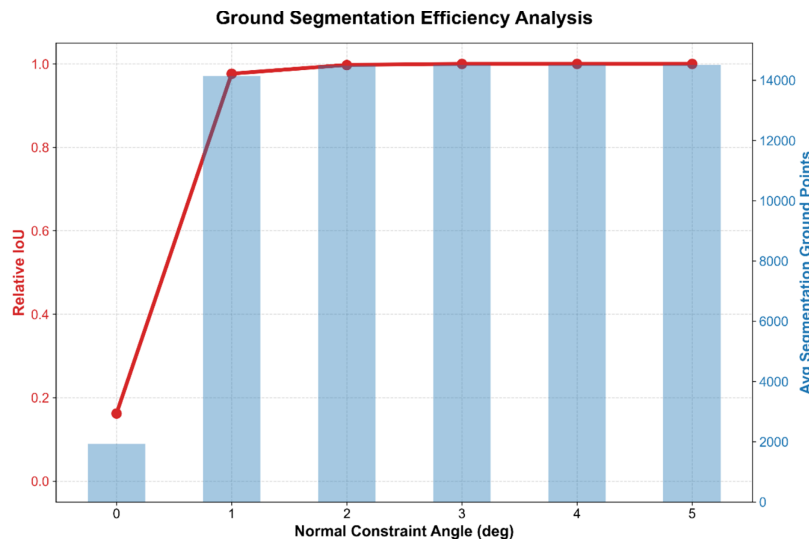


Figure 3: Performance of ground segmentation under different normal constraint angles.

Table 1: Statistical results of ground segmentation performance under different normal constraint angles.

Angle	Mean IoU	Std. Dev IoU	Median IoU	Avg. Ground Points
0	0.1619	0.0661	0.1646	1930.5603
1	0.9765	0.0712	1.0000	14,139.2931
2	0.9975	0.0193	1.0000	14,485.5345
3	1.0000	0.0000	1.0000	14,513.2759
4	1.0000	0.0000	1.0000	14,513.2759

4.4 Validation of Adhering Pile Instance Segmentation in Real Scenes

To validate the decoupling capability of the proposed framework for complex adhering piles in industrial applications, representative heavily adhering areas from real coal shed sites are selected for testing. The experiment aims to demonstrate how the method achieves automatic segmentation from continuous connected domains to independent coal pile instances based solely on geometric topological features without manual intervention.

Experimental data are cropped from on-site point cloud samples, covering “conjoined” coal piles with varying heights and adhesion morphologies. These samples are physically adjacent with no obvious spatial separation due to continuous stacking operations. Notably, the “saddle point” region at the connection is affected by dust coverage and natural particle sliding, resulting in smooth and blurred geometric transitions lacking distinct edge features. Such non-convex connected structures are easily judged as single objects by traditional Euclidean clustering due to satisfying distance threshold conditions, preventing the acquisition of independent information for individual piles during inventory.

Fig. 4 presents four images labeled (a)–(d). Images (c) and (d) show the segmentation of the same coal pile sample from different viewpoints, illustrating the algorithm’s capability to handle varying adhesion morphologies. Images (a) and (b) display the original scene and the corresponding segmentation result, respectively. The visual results indicate that the proposed strategy based on PCA principal direction projection and concavity feature manifold cutting effectively identifies geometric local minima within connected

domains. The algorithm precisely locates the natural boundary between piles (the “saddle point” position with maximum concavity), and constructs a segmentation interface consistent with physical stacking logic, without obvious over-segmentation or under-segmentation. These results verify the effectiveness and universality of the method in handling complex non-convex connected domains in industrial sites.

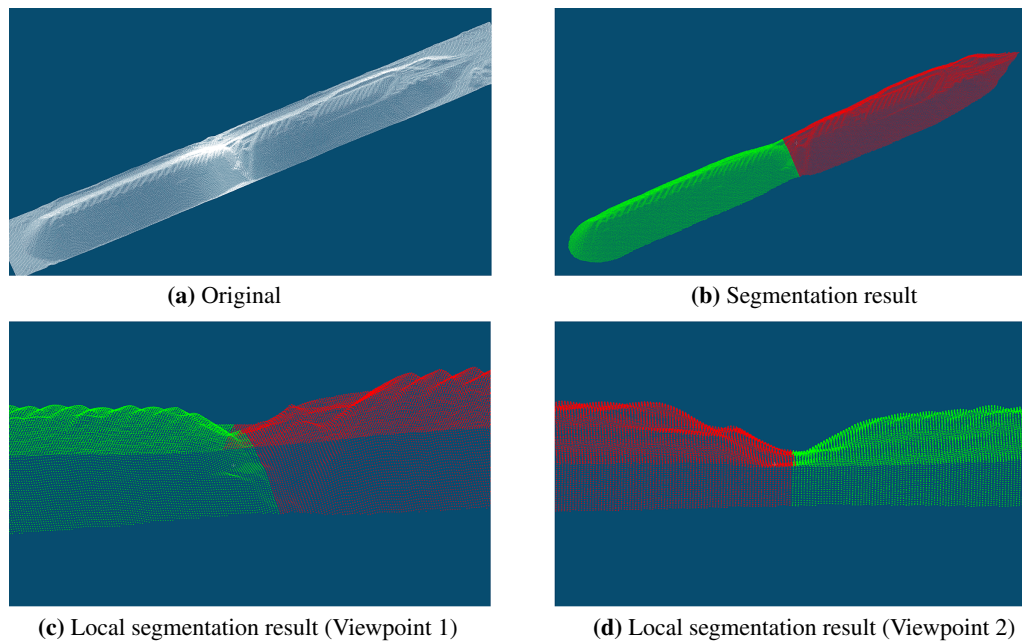


Figure 4: Examples of instance segmentation results for adhering coal piles.

5 Conclusions

This paper proposes a point cloud segmentation framework for industrial coal shed scenarios, aiming to address the challenges of ground segmentation accuracy and adhering to instance decoupling in complex pile environments. In the data preprocessing and ground segmentation stages, the framework adopts relative elevation mapping and selective statistical filtering strategies, preserving geometric details of pile tops while removing environmental noise. For the ground extraction phase, the NC-RANSAC algorithm introduces a normal vector angular geometric constraint. Experiments demonstrate that this mechanism exhibits high parameter stability within a certain parameter domain, effectively avoiding interference from non-ground planes and ensuring strict consistency of the extraction results in terms of geometric-physical properties. In the instance segmentation stage, addressing the challenge where physically adhering piles are difficult to distinguish via Euclidean clustering, this paper proposes an unsupervised segmentation strategy based on PCA principal direction projection and sliding window concavity detection. This method leverages geometric topological features to precisely locate saddle points within connected domains, achieving automatic decoupling of continuous connected domains into independent coal pile instances without manual intervention or deep learning sample training.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding.

Author Contributions: The authors confirm contribution to the paper as follows: study conception and design: Ping Zhang, Chungang Li; data collection: Jian Song; analysis and interpretation of results: Jianjun Tan, Han Tian;

draft manuscript preparation: Le Chen, Yongkun Song. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are not publicly available due to commercial confidentiality restrictions imposed by the cooperating industrial partner.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. International Energy Agency. Global Energy Review 2025. 2025 [cited 2026 Jan 24]. Available from: <https://www.iea.org/reports/global-energy-review-2025>.
2. Yu X, Shi D, Li Y, Jiao K, Chen C, Ou L. A LiDAR-based approach for real-time bulk coal extraction in cargo holds. *Ocean Eng.* 2025;342:122819. doi:10.1016/j.oceaneng.2025.122819.
3. Alsayed A, Nabawy MRA. Stockpile volume estimation in open and confined environments: a review. *Drones.* 2023;7(8):537. doi:10.3390/drones7080537.
4. Kovanič Ľ, Topitzner B, Peřovský P, Blišťan P, Gergelřov MB, Blišťanov M. Review of photogrammetric and lidar applications of UAV. *Appl Sci.* 2023;13(11):6732. doi:10.3390/app13116732.
5. Li Y, Bu R, Sun M, Wu W, Di X, Chen B. PointCNN: convolution on X-transformed points. In: *Proceedings of the Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*; 2018 Dec 3–8; Montreal, QC, Canada.
6. Charles RQ, Hao S, Mo K, Guibas LJ. PointNet: deep learning on point sets for 3D classification and segmentation. In: *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; 2017 Jul 21–26; Honolulu, HI, USA. doi:10.1109/CVPR.2017.16.
7. Qi CR, Yi L, Su H, Guibas LJ. PointNet++: deep hierarchical feature learning on point sets in a metric space. In: *Proceedings of the Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*; 2017 Dec 4–9; Long Beach, CA, USA.
8. Liu Y, Zhang C, Dong X, Ning J. Point cloud-based deep learning in industrial production: a survey. *ACM Comput Surv.* 2025;57(7):1–36. doi:10.1145/3715851.
9. Wang Y, Xie D, Ren S, Sun Z, Yi C, Xiao R, et al. Indoor coal inventory algorithm using UAV-based LiDAR SLAM. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2025;18:26219–34. doi:10.1109/JSTARS.2025.3614735.
10. Yang X, Huang Y, Zhang Q. Automatic stockpile extraction and measurement using 3D point cloud and multi-scale directional curvature. *Remote Sens.* 2020;12(6):960. doi:10.3390/rs12060960.
11. Cheng J, He D, Lee C. A simple ground segmentation method for LiDAR 3D point clouds. In: *Proceedings of the 2020 2nd International Conference on Advances in Computer Technology, Information Science and Communications (CTISC)*; 2020 Mar 20–22; Suzhou, China. doi:10.1109/CTISC49998.2020.00034.
12. Zhang K, Chen SC, Whitman D, Shyu ML, Yan J, Zhang C. A progressive morphological filter for removing nonground measurements from airborne LIDAR data. *IEEE Trans Geosci Remote Sens.* 2003;41(4):872–82. doi:10.1109/TGRS.2003.810682.
13. Luo N, Jiang Y, Wang Q. Supervoxel-based region growing segmentation for point cloud data. *Int J Patt Recogn Artif Intell.* 2021;35(3):2154007. doi:10.1142/s0218001421540070.
14. Chen H, Liang M, Liu W, Wang W, Liu PX. An approach to boundary detection for 3D point clouds based on DBSCAN clustering. *Pattern Recognit.* 2022;124:108431. doi:10.1016/j.patcog.2021.108431.
15. Zhang H, Duan Z, Zheng N, Li Y, Zeng Y, Shi W. An efficient class-constrained DBSCAN approach for large-scale point cloud clustering. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2022;15:7323–32. doi:10.1109/JSTARS.2022.3201991.
16. Ma J, Luo B, Zhao Y, Li S, Chen M, Wei J, et al. Detection and automatic identification of landslide areas from the LiDAR point clouds using improved DBSCAN. *Landslides.* 2025;22(11):3843–54. doi:10.1007/s10346-025-02569-w.

17. Ling Y, Zhao R, Shen Y, Li D, Jin J, Liu J. DIVESpot: depth integrated volume estimation of pile of things based on point cloud. arXiv:2407.05415. 2024.
18. Wang Y, Sun Y, Liu Z, Sarma SE, Bronstein MM, Solomon JM. Dynamic graph CNN for learning on point clouds. ACM Trans Graph. 2019;38(5):1–12. doi:10.1145/3326362.