CIVIL ENGINEERING STUDIES
Illinois Center for Transportation Series No. 22-017
UILU-ENG-2022-2017
ISSN: 0197-9191

3D Image Analysis Using Deep Learning for Size and Shape Characterization of Stockpile Riprap Aggregates—Phase 2

Prepared By
Haohang Huang, PhD
Erol Tutumluer, PhD
Jiayi Luo
Kelin Ding
Issam Qamhia, PhD
John M. Hart
University of Illinois Urbana-Champaign

Research Report No. FHWA-ICT-22-013

A report of the findings of

ICT PROJECT R27-214
3D Image Analysis Using Deep Learning for Size and Shape Characterization of Stockpile Riprap Aggregates—Phase 2

https://doi.org/10.36501/0197-9191/22-017

Illinois Center for Transportation

September 2022

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
FHWA-ICT-22-013	N/A	N/A

4. Title and Subtitle	5. Report Date
3D Image Analysis Using Deep Learning for Size and Shape Characterization of Stockpile Riprap Aggregates—Phase 2	September 2022
	6. Performing Organization Code N/A

7. Authors	8. Performing Organization Report No.
Haohang Huang (https://orcid.org/0000-0003-3237-5328), Erol Tutumluer (https://orcid.org/0000-0003-3945-167X), Jiayi Luo, Kelin Ding, Issam Qamhia (https://orcid.org/0000-0002-2020-8437), and John M. Hart	ICT-22-017 UILU-2022-2017

9. Performing Organization Name and Address	10. Work Unit No.
Illinois Center for Transportation Department of Civil and Environmental Engineering University of Illinois at Urbana-Champaign 205 North Mathews Avenue, MC-250 Urbana, IL 61801	N/A **11. Contract or Grant No.** R27-214

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Illinois Department of Transportation (SPR) Bureau of Research 126 East Ash Street Springfield, IL 62704	Final Report 7/1/20–9/30/22 **14. Sponsoring Agency Code**

15. Supplementary Notes

Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.
https://doi.org/10.36501/0197-9191/22-017

16. Abstract

Riprap rock and aggregates are extensively used in structural, transportation, geotechnical, and hydraulic engineering applications. Field determination of morphological properties of aggregates such as size and shape can greatly facilitate the quality assurance/quality control (QA/QC) process for proper aggregate material selection and engineering use. Many aggregate imaging approaches have been developed to characterize the size and morphology of individual aggregates by computer vision. However, 3D field characterization of aggregate particle morphology is challenging both during the quarry production process and at construction sites, particularly for aggregates in stockpile form. This research study presents a 3D reconstruction-segmentation-completion approach based on deep learning techniques by combining three developed research components: field 3D reconstruction procedures, 3D stockpile instance segmentation, and 3D shape completion. The approach was designed to reconstruct aggregate stockpiles from multi-view images, segment the stockpile into individual instances, and predict the unseen side of each instance (particle) based on the partial visible shapes. Based on the dataset constructed from individual aggregate models, a state-of-the-art 3D instance segmentation network and a 3D shape completion network were implemented and trained, respectively. The application of the integrated approach was demonstrated on re-engineered stockpiles and field stockpiles. The validation of results using ground-truth measurements showed satisfactory algorithm performance in capturing and predicting the unseen sides of aggregates. The algorithms are integrated into a software application with a user-friendly graphical user interface. Based on the findings of this study, this stockpile aggregate analysis approach is envisioned to provide efficient field evaluation of aggregate stockpiles by offering convenient and reliable solutions for on-site QA/QC tasks of riprap rock and aggregate stockpiles.

17. Key Words	18. Distribution Statement
Riprap, Aggregate Stockpile, Field Imaging, Computer Vision, Deep Learning, 3D Reconstruction, 3D Segmentation, Size and Shape Analyses	No restrictions. This document is available through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	66	N/A

ACKNOWLEDGMENT, DISCLAIMER, MANUFACTURERS' NAMES

This publication is based on the results of **ICT-R27-214: 3D Image Analysis Using Deep Learning for Size and Shape Characterization of Stockpile Riprap Aggregates—Phase 2**. ICT-R27-214 was conducted in cooperation with the Illinois Center for Transportation; the Illinois Department of Transportation; and the U.S. Department of Transportation, Federal Highway Administration.

Members of the Technical Review Panel (TRP) were the following:

- Andrew Stolba, TRP Chair, Illinois Department of Transportation
- Chad Nelson, Illinois Department of Transportation—District 3
- Del Reeves, Illinois Department of Transportation—District 6
- Kevin Tressel, Illinois Department of Transportation
- Megan Swanson, Illinois Department of Transportation
- Dennis Bachman, Federal Highway Administration
- Sheila Beshears, Riverstone Group, Inc.
- Andrew Buck, Vulcan Materials Company
- Jim Sergent, Hanson Material Service

The authors would like to extend their appreciation to the TRP Chair and all TRP members for their time and significant contributions to the findings of this research project. The help of Andrew Stolba, Sheila Beshears, Chad Nelson, Andrew Buck, and Dan Barnstable in coordinating quarry field visits to sample and image individual riprap rocks and stockpiles is highly appreciated.

The authors would like to express their sincere gratitude to the University of Illinois Urbana–Champaign (UIUC) and Illinois Center for Transportation (ICT) students who helped on several tasks dealing with materials characterization, lab testing, image acquisition, and field visits: Han Wang, Zhongyi Liu, and Wenjing Li. Sincere acknowledgement is due to Ren Liu for the insightful advice and technical help during the software development phase.

The authors would also like to acknowledge the technical support received from the Computer Vision and Robotics Laboratory in the Beckman Institute for Advanced Science and Technology at UIUC. Sincere acknowledgments are also due to ICT research engineer Greg Renshaw for his support throughout the project.

The authors are particularly grateful for the help with material procurement and technical support given by Timothy Prunkard, Megan Everette, and Suzanne Lowry from the Department of Civil and Environmental Engineering at UIUC. A great thank-you is extended to Kristi Anderson, Audrey Donoho, McCall Macomber, and Brian Lorbiecki for ICT project management.

policies of the Illinois Center for Transportation, the Illinois Department of Transportation, or the Federal Highway Administration. This report does not constitute a standard, specification, or regulation.

Trademark or manufacturers' names appear in this report only because they are considered essential to the object of this document and do not constitute an endorsement of product by the Federal Highway Administration, the Illinois Department of Transportation, or the Illinois Center for Transportation.

EXECUTIVE SUMMARY

This research project developed a 3D imaging approach based on deep learning techniques for automating the process of size and shape characterization for riprap aggregate stockpiles during on-site quality assurance/quality control (QA/QC) evaluations. Riprap rock and large-sized aggregates have been used extensively in geotechnical and hydraulic engineering applications to armor shorelines, streambeds, bridge abutments, pilings, and other shoreline structures against scour and water or ice erosion. They are made from a variety of rock types, commonly granite or limestone and occasionally concrete rubble from building and paving demolition. Their main engineering functions include erosion and sediment control as well as scour protection. At both quarry production sites and construction sites, the sustainable and reliable use of riprap materials demands efficient and accurate evaluation of their large particle sizes, shapes, and gradation information. However, determination of particle size distribution or gradation of riprap has always been difficult. The riprap gradation or sizing of the rocks shall be well graded, as outlined in Illinois Department of Transportation (IDOT) specifications. Individual rocks are weighed or visually inspected by keystone method, which is a time-consuming and labor-intensive task. In this regard, reliable field imaging techniques are of utmost need to efficiently and safely assess riprap stockpiles for gradation and to provide informative data analytics about the size distribution.

This research effort, conducted at the Illinois Center for Transportation (ICT), focused on developing an advanced nonintrusive machine vision system for field evaluation of riprap aggregates. This system will allow engineers to obtain 3D size and shape information of individual particles in a riprap stockpile simply by taking videos or images with mobile devices, such as a smartphone camera. The research followed a reconstruction-segmentation-completion pipeline based on computer vision methods. The riprap surface can be reconstructed as 3D point cloud data by using structure-from-motion (SfM) techniques to capture high-definition spatial information of the scene in the field. A 3D point cloud segmentation algorithm was developed to separate and extract individual riprap rocks from the reconstructed stockpile. To supplement the partially occluded shape information of particles on the stockpile surface, 3D shape completion algorithms were developed to estimate missing shape properties based on an established database of riprap rock shape characteristics. The deep learning approaches achieved the desired user-independent performance level. Finally, trained algorithms were integrated into a user-friendly software application with a user interface suited for field use by practitioners.

The following approach was utilized to achieve the research goals. First, riprap imaging data in the form of images and videos were collected and compiled as datasets to be used throughout the entire research study. Providing continuity with the ICT-IDOT Phase I study (R27-182), the compiled raw data included: (a) image and video data collected from several quarry visits during the Phase I study, (b) image and video data of medium-sized stockpiles built in a laboratory setting, and (c) ground-truth size and shape measurements of the riprap samples for algorithm development and validation. The collected riprap samples constitute a great resource that enables efficient experimentation, development, and verification of the state-of-the-art algorithms. This data collection activity was conducted as an ongoing effort throughout the study for the continuous 3D algorithm development and its validation.

Next, to overcome the limitations of single view 2D image representations of a stockpile, 3D reconstruction algorithms were developed. Image sequences taken from multiple views or frames extracted from a moving-camera video were used as the input in the 3D reconstruction step. Advanced techniques such as SfM were used to obtain high-definition spatial information of the scene in the field, and state-of-the-art reconstruction algorithms were selected to reconstruct the riprap stockpile surface as 3D point cloud data. Point cloud data can represent the stockpile by points with spatial coordinates and color values. 3D reconstruction approaches focus on capturing the structure of the stockpile, which is less affected by environmental conditions such as shadow effects, in field imaging, and hence the performance is more robust than 2D image-based techniques.

After the 3D reconstruction development, 3D segmentation algorithms were developed to obtain useful morphological information of individual riprap particles from the stockpile's point cloud data. Due to the spatial sparsity of the point cloud representation, conventional computer vision methods are either less robust or highly user dependent. In this regard, this effort developed state-of-the-art 3D segmentation algorithms based on deep learning techniques. As a result, the segmentation algorithms work in an automated fashion for robust processing under a wide variety of riprap and large-sized aggregate stockpile field conditions, e.g., strong sunlight or shadowy conditions.

3D segmentation algorithms were followed by the development of 3D morphological analysis algorithms applied on individual extracted particles. Although each segmented particle is an incomplete surface due to overlap and occlusion, its true shape can be better estimated by utilizing the 3D geometric information that is visible based on the multi-view observation. This effort developed state-of-the-art 3D shape completion algorithms based on deep learning techniques, which learned the representative aggregate shapes to achieve a reasonable estimation of the unseen or missing part of the segmented particle to aid in the prediction of a correct volume and weight.

As the key deliverable, a user-independent and user-friendly software application was developed in this research study. The software application was developed to process riprap stockpile data for evaluation of the size and shape of individual aggregate particles captured by the 3D reconstruction approach. Professional software engineering guidelines were reviewed and followed during the developmental stages. All developed riprap stockpile analysis algorithms were programmed and integrated into a standalone software application with a simple graphical user interface (GUI). A user manual is provided along with this final report for the developed software program.

In conclusion, 3D algorithms for reconstruction, segmentation, and shape completion of aggregate stockpiles were developed, validated, and integrated into a software application with a GUI. The product of this research project is a computer vision approach consisting of field imaging methods and advanced 3D computer vision algorithms to provide field analysis of riprap aggregate properties at a quarry or construction site. This approach helps to accurately determine if a riprap material meets gradation as well as volume and weight specifications. As a result, more efficient property characterization of riprap stockpiles and optimized material selection can be achieved in practice to improve designs through effective quality control, reduced costs, increased life cycle, lower environmental impacts, and minimized labor utilization and energy consumption.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: INTRODUCTION

BACKGROUND AND MOTIVATION

Riprap rock and large-sized aggregates have been used extensively in geotechnical and hydraulic engineering to armor shorelines, streambeds, bridge abutments, pilings, and other shoreline structures against scour and water or ice erosion. They provide erosion and sediment control as well as scour protection. The sustainable and reliable use of riprap materials demands efficient and accurate evaluation of their large particle sizes, shape, and gradation information at both quarry production lines and construction sites.

Determination of particle size distribution (gradation) of riprap has always been challenging. Despite the ongoing development of guidelines for size selection of riprap, the practical procedures for characterizing riprap size and shape properties in the field are still subjective and qualitative, primarily due to difficulties associated with measuring sizes of these large particles. As compared to coarse aggregates used in transportation engineering, the sizes of which typically range from 0.187 in. (4.75 mm) to 5 in. (12.7 cm) (ASTM D448 2017; ASTM D2940 2015), an individual riprap rock can weigh up to 1,150 lb (522 kg) with nominal top sizes greater than 24 in. (61.0 cm) (IDOT 2022a; ASTM D6092 2021). Laboratory sieve analysis is usually conducted to determine the gradation of small- to medium-sized aggregates, but the large sizes of riprap particles make this task impractical. Due to the lack of a practical and standard method to define the sizes and dimensions of individual rocks, standards or guidelines usually specify riprap gradation requirements in terms of weight.

For the current state of the practice, a nationwide AASHTO survey of many transportation agencies in the United States and Canada has indicated that riprap characterization is mostly based on visual inspection and hand measurements (Sillick 2017). Visual inspection relies greatly on the experience and expertise of practitioners. In this method, certain gauge or key stones and sample stockpiles are often used as a reference to assist the judgement (Lippert 2012). To better estimate the size distribution, the Wolman count method may be applied by statistically sampling and measuring rocks within a stockpile (Wolman 1954; Lagasse et al. 2006; Bartelt 2018). For hand measurement, transportation agencies either weigh individual particles directly or use size-mass conversion after measuring particle dimensions. The piece size-mass conversion proposed in ASTM D5519 (2015) requires measurement of the midway dimension or circumference from three orthogonal axes and estimates the volume based on a cuboid assumption or averaged sphere-cube assumption. However, visual inspection and hand measurement are both rough estimations that do not represent realistic riprap properties, and an objective and efficient approach of quantitatively characterizing the size and shape of riprap must be established.

Over the past two decades, imaging-based analysis techniques have been widely developed and adopted to characterize particle size and shape properties of aggregates from their scanned photos. These techniques have been significantly enhanced to provide high accuracy and efficiency in the characterization (Rao et al. 2002; Al-Rousan et al. 2005; Pan et al. 2006; Wang et al. 2013; Moaveni et al. 2013; Hryciw et al. 2014). Current laboratory image-based analysis approaches mostly focus on the shape analyses using camera(s) attached to a table-sized equipment, on which a set of prepared

particles are placed for photo capture. For example, the Aggregate Imaging System (AIMS) uses a setup consisting of one slide-mounted camera and two lighting sources. The Enhanced University of Illinois Aggregate Image Analyzer (E-UIAIA) takes three 2D images of particles on a conveyor system from three orthogonal axes where cameras are installed. Among the most widely used imaging-based indices for aggregate shape characterization are the flat and elongated ratio (FER), angularity index (AI), and surface texture index (STI). These were developed as key indices from national studies, NCHRP 4-30 and 4-34 and the TPF-5(023) pool fund, using E-UIAIA (Tutumluer et al. 2000; Rao et al. 2002; Pan et al. 2006; Moaveni et al. 2013).

The laboratory imaging systems generally use a limited number of 2D images, typically up to three, because of the fixed camera positions. As a result, the generated 3D particles may lack shape details such as local concave curvatures on the surface. Furthermore, the setups are usually not portable and are unwieldy for field inspection. While laboratory imaging equipment has been mostly used for shape characterization, in situ imaging alternatives, e.g., WipFrag (Maerz et al. 1996), have been developed to obtain the grain size distribution data in the field. IDOT used the WipFrag system for riprap imaging in the field with limited success because image quality was adversely affected by shadows and overlaps between particles in the image. Therefore, the existing laboratory or in situ methods do not provide a complete set of analysis data for both particle size and shape characterization.

Because the size and shape properties of riprap and large-sized aggregates cannot be obtained with standard sieve analysis or laboratory-based image analyzer setups, a newly developed field imaging technique was first implemented and used in the ICT-IDOT study R27-124 (Kazmee and Tutumluer 2015). For this effort, representative particles of aggregate subgrade materials (RR01, CS01, and CS02 per IDOT specifications) with particles greater than 3 in. (7.6 cm) in size were spread out on a blue tarp in the field. High-resolution (e.g., 1292 × 964 pixels) images of these large aggregate particles were captured with a digital single-lens reflex camera. A white-colored calibration ball placed beside the aggregate particles was introduced to scale the spatial resolution properly. The final outcome of the segmentation is a simplified black and white binary image of each particle, where the black pixels are the background and the white pixels are the object. The morphological indices FER, AI, and STI were computed from the extracted binary images. The dimensions of the particles were determined in reference to the spatial resolution of the calibration ball.

In view of recent advances in the area of computer vision and machine learning, machine vision–based inspection systems trained using sample images have the potential for directly identifying individual particles for shape and size analysis from stockpile or trench views of the constructed layer showing aggregate assemblies. For example, image segmentation has been implemented at UIUC to study individual railroad ballast particles collected in the field in terms of their shape and size variability versus depth of sampling (Moaveni et al. 2013). Through preliminary image analysis and algorithm development, individual particles can be segmented from images with overlapping particles, and their size and shape properties can be quantified from the cut sections imaged from field ballast trenches. Nevertheless, the image analysis algorithm entails heavy interaction with the user for fine-tuning segmentation parameters. The quality of the segmentation results for this initial effort also needed to be further enhanced.

In this regard, the preceding Phase I research study (ICT-IDOT R27-182) addressed the main challenges and provided a user-independent approach for 2D aggregate image analysis. A single-particle field imaging system was developed in the ICT-IDOT project R27-182 (Huang et al. 2020a). A field imaging system was designed and built as a portable and versatile toolkit for the convenience of efficient and reliable image acquisition needs. Image segmentation and 3D volumetric reconstruction algorithms were then developed for single rocks with the capabilities of segmenting object information under uncontrolled field lighting conditions and reconstructing the object three-dimensionally with necessary calibration and corrections (Huang et al. 2019). The robustness and accuracy of the developed algorithms were studied through field imaging tasks at two aggregate production sites. Good agreements between the ground-truth volume and weight measurements as well as the image-based volumetric reconstruction results were achieved. Comparisons were also made between the image-based volumetric reconstruction results and the state-of-the-practice hand measurements. Significant improvements could be achieved using the developed field imaging system (Huang et al. 2019).

Besides the single-particle approach, an advanced deep learning–based stockpile imaging system was also developed in R27-182 (Huang et al. 2020a), which demonstrated the ability to facilitate the riprap characterization task. In accordance, significant research efforts have been made to achieve the automated segmentation of 2D stockpile images under complex field conditions. The project utilized state-of-the-art deep learning techniques to develop a convenient, nonintrusive field evaluation system based on imaging, whereby an engineer or inspector can take multiple photos of riprap and large-sized aggregates to determine the gradation (size distribution) and shape properties, using a calibration ball in each image to aid analysis (Huang et al. 2020b). A neural network was designed and trained to accomplish the challenging task of aggregate stockpile image segmentation. By establishing an image dataset of labeled aggregate images, the trained segmentation model achieved good performance for extracting individual aggregate particles in an automated manner. Morphological analyses were conducted on the segmented aggregate particles to produce size and shape distribution curves. Analysis results were verified with ground-truth labeling to measure the robustness and accuracy of the segmentation approach. Completeness and precision analyses were conducted between the labeled images in the validation set and their segmentation results. The segmentation model exhibited satisfactory precision and completeness in the stockpile image-segmentation task. As the main deliverable, all developed algorithms were integrated into a user-independent and user-friendly software application, named I-RIPRAP, for image-based quantifications of riprap size and shape characteristics.

Nevertheless, there are certain limitations of 2D imaging approaches because a significant amount of useful spatial information is lost when projecting a 3D scene onto a 2D image plane. 3D size and shape information, on the other hand, offers more comprehensive geometric features as well as more accurate characterization of riprap and large-sized material at quarries and construction sites. Reliable and efficient 3D imaging techniques that can facilitate QA/QC checks are still in high demand for accurately evaluating riprap stockpiles. Additionally, the majority of the particles seen in a riprap stockpile are not fully visible. Thus, shape information associated with the unseen part of the particles is unavailable, as well as the particles hidden underneath the particles on the surface of the stockpile. Using 3D imaging approaches is expected to provide supplemental information from

advanced stockpile reconstruction and segmentation techniques. Furthermore, deep learning–based shape completion algorithms could be leveraged to estimate and predict the occluded size and shape information, allowing a more comprehensive characterization of the aggregates comprising the riprap stockpile.

The main benefit of this project is the ability to conduct riprap and large-sized aggregate size and shape property testing in the field, i.e., at a quarry or construction site, with improved accuracy in determining the size and shape properties from imaging without having to weigh rock pieces. This is a pressing need because engineers and inspectors at quarries and government agencies have been spending considerable time and effort in riprap sizing at job sites.

Implementation of the advanced machine vision system that is being fully developed in this project will help to accurately determine if a riprap material meets gradation and 3D volume-based weight specifications. The proposed approach is sustainable for field application by utilizing state-of-the-art machine vision and stockpile segmentation and shape completion techniques. Ultimately, better property characterization and optimized material selection can be achieved to improve designs through effective quality control, increasing the life cycle and reducing costs, environmental impacts, labor needs, and energy consumption. Major cost savings in terms of personnel time, transportation, laboratory equipment, and facility use can be realized.

RESEARCH OBJECTIVE

The objective of this research project is to develop an advanced nonintrusive machine vision system for field evaluation of riprap aggregates, whereby engineers can obtain 3D size and shape information of individual particles in a riprap stockpile simply by taking videos and images with mobile devices such as a smartphone camera. The research follows a reconstruction-segmentation-completion pipeline based on computer vision methods. The riprap surface is reconstructed as 3D point cloud data by using structure-from-motion (SfM) techniques to capture high-definition spatial information of the scene in the field. A 3D point cloud segmentation algorithm is developed to separate and extract individual riprap rocks from the reconstructed stockpile. To supplement the partially occluded shape information of particles on the stockpile surface, 3D shape completion algorithms are established to estimate missing shape properties based on an established database of riprap rock shape characteristics. The deep learning approaches being adopted are intended to achieve the desired user-independent performance in the above tasks. Finally, the resulting algorithms are integrated into a user-friendly software application with a user interface suited for practical field use.

RESEARCH METHODOLOGY

To fulfill the above-stated research objective, the following tasks and methodologies are considered in this study:

- Collecting riprap imaging data in the form of images and videos and compiling as datasets. The raw data include: (a) image and video data collected from several quarry field visits during the Phase I study (Huang et al. 2020a), (b) image and video data of medium-sized stockpiles built in a laboratory setting, and (c) ground-truth size and shape measurements of the riprap

samples for algorithm development and validation. The collected riprap samples constitute a resource that enables efficient experimentation, development, and verification of state-of-the-art algorithms. This data collection activity was an ongoing effort throughout the study for continuous 3D algorithm development and validation.

- Developing an integrated framework that implements automated 3D point cloud reconstruction, segmentation, completion, and morphological analyses for aggregate stockpiles and constructed layers. To obtain more comprehensive information of aggregate stockpiles and field constructed layers, a 3D point cloud reconstruction approach is developed. State-of-the-art deep learning architecture for 3D object detection and instance segmentation is implemented and trained to enable automated segmentation of stockpile and field constructed aggregate clouds. Further, a 3D particle shape completion approach as well as 3D morphological analysis algorithms were developed to characterize the meaningful 3D size, shape, and volumetric properties of the segmented aggregates.

- Validating the accuracy and robustness of the developed algorithms by comparing results with measured ground-truth data in field applications. Comparisons are made to ensure that the algorithms' performance and reliability are in good agreement with engineering practice.

- Designing and developing a software application as the integrated riprap aggregate evaluation module for size and shape analyses of individual particles captured in stockpile images. A graphical user interface (GUI) is designed to facilitate user interaction.

REPORT ORGANIZATION

This report consists of seven chapters, including this introductory chapter.

Chapter 2, titled "Literature Review," provides a review of riprap standards and specifications, past aggregate studies and equipment that leverage imaging techniques, and the applications of artificial intelligence and deep learning techniques.

Chapter 3, titled "3D Reconstruction for Field Data of Individual Aggregates and Aggregate Stockpiles," provides an overview of the 3D reconstruction approach developed for laboratory and field samples. This chapter includes riprap source information from the riprap production sites visited by the research team, material selection and image acquisition criteria, and laboratory tests for measuring the ground-truth data of collected samples.

Chapter 4, titled "Automated 3D Segmentation of Aggregate Stockpiles," provides the algorithmic details of the 3D instance segmentation network. This chapter introduces the development of a computer vision–based approach for the instance segmentation of individual riprap particles from a stockpile point cloud, which is user-independent, automated, and powered by deep learning techniques.

Chapter 5, titled "Automated 3D Aggregate Shape Completion and Field Validation," provides the algorithmic details of the 3D shape completion and morphological analyses as well as the verification

results with ground-truth data. This chapter introduces a shape completion approach that is developed and further evaluated on several unseen aggregate shapes for robustness and reliability.

Chapter 6, titled “Software Development for 3D Riprap Stockpile Characterization,” provides a detailed description of the software development stages undertaken in this project, which includes the software engineering concepts adopted, the graphical design and control logic of the software application, and the back-end implementation details of the software.

Chapter 7, titled “Conclusions and Recommendations,” provides a summary of research findings and the main recommendations and conclusions from this study. This chapter also discusses proposed next steps to further extend this study, and recommendations for future research.

CHAPTER 2: LITERATURE REVIEW

INTRODUCTION

This chapter presents a review of the riprap quality assurance/quality control (QA/QC) requirements as well as previous research studies related to the imaging techniques and developed systems used in aggregate research and the emerging artificial intelligence–based approaches for image segmentation. An overview of the current practices for characterizing riprap and large-sized aggregate materials is also discussed.

APPLICATIONS AND SPECIFICATIONS OF RIPRAP AND LARGE-SIZED AGGREGATES

Riprap rock and large-sized aggregates are quarried from undisturbed, consolidated rock deposits or recycled from demolished concrete construction. As a natural material, the reliable and sustainable use of riprap as an integrated system requires quality control throughout the design, production, transport, installation, inspection, and maintenance stages (Lagasse et al. 2006). Case studies on riprap failure in stream channels and bridge piers indicate that undersized particles and open-graded riprap often cause insufficient resistance to hydraulic shear stress (Blodgett and McConaughy 1986; Chiew 1995; Richardson and Davis 2001; Lagasse et al. 2001).

During the material selection and QA/QC process, characterizing particle size and shape properties has become a focal point for aggregate studies. Particle size and morphological/shape properties of aggregates primarily influence the macroscopic behavior and performance of aggregate skeleton assemblies of constructed layers in transportation infrastructure, e.g., hot-mix asphalt (HMA) and Portland cement concrete (PCC) (Quiroga and Fowler 2004; Polat et al. 2013), unbound/bound layers in highway and airfield pavements (Liu et al. 2019; Tutumluer and Pan 2008; Bessa et al. 2015), the ballast layer in railway tracks (Huang 2010; Wnek et al. 2013), and riprap materials for erosion control and hydraulic applications (Lutton et al. 1981; Lagasse et al. 2006). Across all size ranges, aggregate shape properties in terms of form (e.g., flatness and elongation), angularity, and texture have been used to characterize their morphology (Barrett 1980). The information on aggregate morphology greatly facilitates the quality control process and in-depth understanding of aggregate layer behavior linked to particle composition and packing.

For producers and practitioners, the size and shape of aggregate sources are important for QA/QC requirements throughout production, mix design (if applicable), and other operations (ASTM D448 2017; ASTM D2940 2015; ASTM D6092 2021). Different quarrying processes and rock mineralogy introduce randomness to the quality of produced aggregates. Therefore, consistent and continuous monitoring of quarry products is important for efficient material selection and construction. On the other hand, particulate mechanics dealing with discrete aggregate particle interactions and realistic modeling of assembly behavior of granular materials require properly characterizing the morphological properties of aggregates. Through recently focused research efforts on modeling the aggregate layer behavior using the finite element and discrete element methods, aggregate morphological properties and grain size distribution have gained increased importance for establishing an accurate composition and packing that enable the modeling scheme to capture

complex behaviors of granular materials. This is especially challenging for most stone skeleton layers in constructed road pavements, e.g., surface course mixtures such as HMA, PCC, and unbound aggregate base/subbase, which are subjected to vehicular dynamic loading conditions (Huang 2010; Chen 2011; Ghauch 2014; Qian 2014). Therefore, uniform guidelines, specifications, or techniques that ensure reliable and efficient characterization of weight, size, shape, and gradation of different aggregate categories are critical in quarry production lines and construction sites (Lagasse et al. 2006).

Despite the ongoing development of guidelines for size selection of riprap in design, the practical procedures for characterizing riprap grain size distribution and shape properties in the field are still subjective and qualitative, primarily due to difficulties associated with measuring sizes of these large particles. The riprap gradation or sizing of the pieces shall be well graded, as outlined in Article 1005.01(c) of IDOT's (2022a) *Standard Specifications for Road and Bridge Construction*. Individual rock pieces are measured for weight, which is a time-consuming and labor-intensive task. Sample grid systems have been used for gradations of RR3 to RR7 with less success. Minnesota DOT (2018), Nevada DOT (2014), and U.S. Army Corps of Engineers (1990) have specifications on riprap gradation by individually weighing or measuring as well. Individual riprap rocks can weigh up to 1,150 lb (522 kg) with nominal top sizes greater than 24 in. (61.0 cm) (IDOT 2022a; ASTM D6092 2021). Laboratory sieve analysis is thus impractical. Because there is no uniform way to define the size dimension of individual rocks, standards or guidelines usually specify riprap gradation requirements in terms of weight.

The current IDOT specification for riprap classification into different "RR" categories is based on the grain size distribution, which is determined by the weight distribution of riprap rocks. IDOT published a policy memorandum in 2012 for the classification of riprap based on weight (Policy Memorandum No. 14-08.1). The latest revision 14.08.3 of this memorandum on April 14, 2022, requires a visual inspection of the riprap stockpiles, including inspections for flat and elongated rocks (IDOT 2022b). A collection of riprap keystones shall be maintained by the producers for all produced riprap gradations, as outlined in Table 1, to assist with visual "sizing." IDOT requires that the set of keystones shall be representative of the stockpile gradation and be replaced with a new set if they become nonrepresentative.

If the gradations by visual inspection were disputed by the producer, a second visual inspection is conducted by IDOT's Central Bureau of Materials (CBM). If the second visual inspection is again disputed by the producer, a representative sample is excavated from the working face of the stockpile and spread over the length of a marked grid to a one-rock thickness and weighed piece by piece for riprap categories RR3–RR7. The rock spalls and fines below the minimum specified weight are collected and included in the calculations for each size range. The grid size for each riprap category is outlined in Table 2. The grid length is broken into 5 ft (1.524 m) long blocks. Note that for riprap categories RR1 and RR2, the grain size distribution is performed by conventional sieve analysis in accordance with AASHTO T 27 (Illinois Modified), outlined in IDOT's (2021) *Manual of Test Procedures for Materials*.

Table 1. Keystone Requirements for Different Riprap Size and Weight Categories

Gradation	Keystone #1 (lb)	Keystone #2 (lb)	Keystone #3 (lb)
RR3	50 (±5)	10 (±1)	1 (±0.1)
RR4	150 (±15)	40 (±4)	1 (±0.1)
RR5	400 (±40)	90 (±13)	3 (±0.1)
RR6	600 (±60)	170 (±17)	6 (±0.5)
RR7	1000 (±100)	300 (±30)	12 (±1)

Note: 1 lb = 453.6 g

Table 2. Grid Size Requirements for Sampling Different Riprap Gradation Categories

Gradation	Grid Size (ft)	Sample Size (Min. Number of Tested Blocks)
RR3	2 by 25	2
RR4	3 by 25	2
RR5	4 by 25	3
RR6	5 by 30	3
RR7	5 by 35	3

Note: 1 ft = 30.48 cm

Based on IDOT (2022b), the procedure for riprap size characterization entails using three keystone particles (used as control points) to identify gradations. Figure 1 shows the upper, lower, and midpoint gradation lines for IDOT's riprap categories RR3–RR7, from left to right. This plot assumes a flat and elongated ratio of 2:1, a specific gravity of 2.5, and an ellipsoidal particle shape for a standardized weight to volume/size conversion. Note that the maximum dimension of a particle is used to indicate size on the horizontal axis.

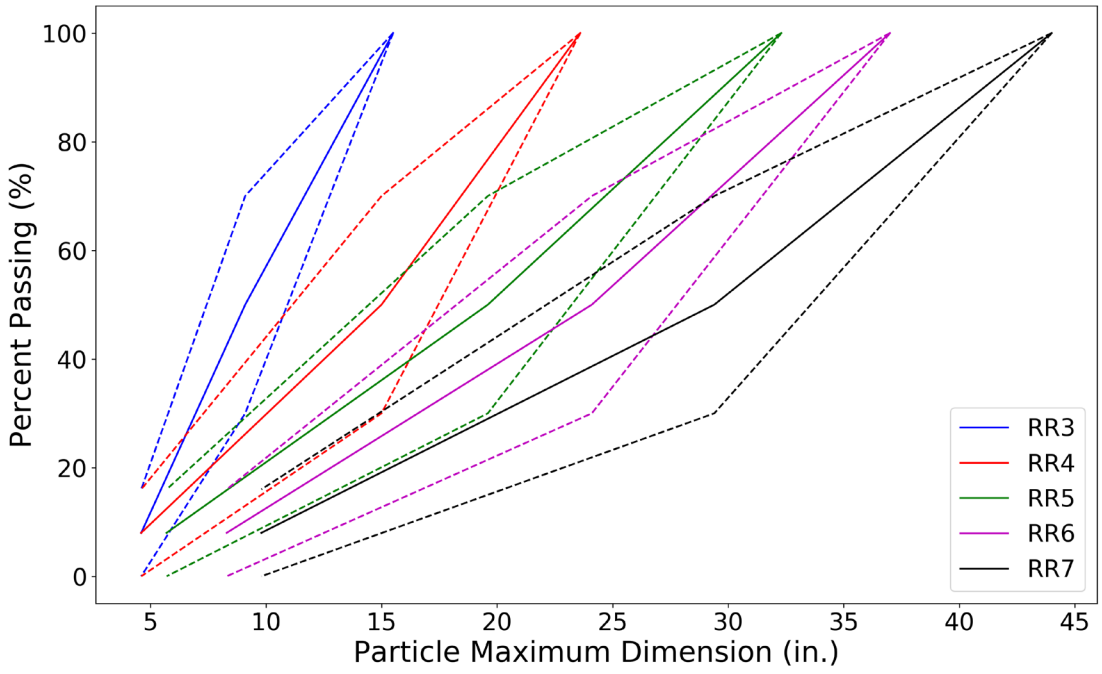

Figure 1. Graph. Converted particle size distribution of IDOT riprap categories RR3–RR7.

A nationwide AASHTO survey of transportation agencies in the United States and Canada indicates that the current state of practice for riprap characterization is mostly based on visual inspection and manual measurements (Sillick and AASHTO 2017). Visual inspection depends on the experience and expertise of practitioners. In this method, certain gauge or keystones and sample stockpiles are usually used as a reference to assist the judgement (Lippert 2012). To better estimate the size distribution, the Wolman count method is applied by statistically sampling and measuring rocks within a stockpile (Lagasse et al. 2006). For manual measurements, transportation agencies either weigh individual particles directly or use size-to-mass conversion after measuring particle dimensions. The size-to-mass conversion proposed in ASTM D5519 (2015) requires measurement of the midway dimension or circumference from three orthogonal axes and estimates the volume based on a cuboid assumption or averaged sphere-cube assumption. However, visual inspection and manual measurements provide rough estimations that do not necessarily represent realistic riprap properties. A more objective and efficient approach for quantitatively characterizing the size and shape of riprap has yet to be established.

ADVANCED AGGREGATE IMAGING SYSTEMS

Over the past two decades, imaging-based analysis has been widely developed and adopted to characterize particle size and shape properties from digital images of aggregates. Related image-segmentation techniques have been significantly enhanced to provide high accuracy and efficiency in practice (Rao et al. 2002; Al-Rousan et al. 2005; Pan et al. 2006; Wang et al. 2013; Hryciw et al. 2014). Most imaging analysis techniques are developed and applied to aggregates with maximum particle sizes less than 6 in. (15.2 cm) using a fixed-position camera setup for acquiring images in a laboratory. Masad (2003) and Gates et al. (2011) developed the Aggregate Imaging System (AIMS), which consists of one slide-mounted camera and two lighting sources visualizing a maximum particle size up to 6 in. (15.2 cm). Tutumluer et al. (2000) originally developed and later Moaveni et al. (2013) improved the Enhanced University of Illinois Aggregate Image Analyzer (E-UIAIA). This system takes three 2D images of particles on a conveyor system, where cameras are installed along three orthogonal axes, and processes the set of orthogonal views for each particle, with a maximum particle size up to 3 in. (7.6 cm).

Komba et al. (2013) established a 3D laser-based aggregate system that scans an individual particle and analyzes the generated 3D mesh model, with a maximum detected sample size of 0.75 in. (1.9 cm). Hryciw et al. (2014) used a translucent segregation table and evaluated the effect of particle size and morphology on the shear strength mechanical properties of the material. Zheng and Hryciw (2014) introduced an algorithm to accurately determine the thickness as well as the size characterization of particles ranging from sand to gravel (maximum size of 1.2 in. [3.0 cm]) simultaneously. Obaidat et al. (2017) established a smartphone-based scheme that utilizes image-processing techniques and ImageJ commercial software to extract coarse aggregate properties, with a maximum particle size up to 2 in. (5.1 cm). Jin et al. (2018) developed an algorithm to perform aggregate shape characterization and volume estimation based on a 3D solid model constructed from X-ray CT images.

Although these imaging systems were developed and validated with ground-truth measurements under laboratory conditions, their capabilities for field application have not been verified. First, these systems are designed with a laboratory-scale setup for object and camera positioning. Thus, they may not be easily transported, assembled, and deployed for field applications, especially for those involving advanced devices such as a 3D laser scanner or an X-ray CT scanner. Moreover, most of the systems have a maximum particle size restriction, limiting their application for handling large-sized aggregates. Further, the lighting conditions for these systems are controlled using backlighting or multiple light sources to minimize the shadow and reflection effects. In addition, the image-segmentation algorithm originally developed for laboratory conditions may not provide accuracy and robustness under field lighting conditions. Consequently, laboratory imaging systems are not readily applicable or adaptable for field inspection.

Furthermore, most of the laboratory imaging systems focus on 2D particle size and shape analysis in lieu of 3D volumetric information, although the weight of individual rocks is needed for determining size distribution of riprap material. Note that the WipFrag software developed by Maerz et al. (1996, 1999) and commercialized by WipWare, Inc. is the only imaging-based system found in the literature that was used to provide riprap characterization based on field images. It was integrated with mobile devices to allow for convenient on-site use to roughly estimate the particle sizes and gradation in a stockpile image. Nevertheless, the image-segmentation procedure used in this software is highly user dependent and its gradation property estimation is based on a single-view stockpile image of riprap, without accounting for detailed 3D particle shape or volumetric information. To achieve user-independent image-segmentation analysis for relatively large-sized aggregates, Huang et al. (2019, 2020b, 2021) recently designed a field imaging system for the convenient volumetric estimation of aggregates under field conditions and developed a morphological analysis software, I-RIPRAP, for stockpile aggregate image analysis.

Among these aggregate imaging systems developed to date, the 3D aggregate stockpile analysis remains an extremely challenging task, with existing systems either focused on 2D image analysis or 3D analysis of separated aggregate particles. An efficient approach with robust and efficient algorithms for obtaining comprehensive 3D geometric information of aggregate particles from stockpiles has not yet been developed.

ANALYSIS APPROACHES FOR AGGREGATE EVALUATION

Traditional methods for aggregate evaluation include visual inspection, geometry measurements, and sieve analysis. More recently, computer vision techniques have been widely adopted to characterize particle size and shape properties, and a variety of aggregate imaging systems have been developed.

Aggregate image analysis typically consists of an image-segmentation module followed by a morphological analysis module based on computational geometry (Al-Rousan et al. 2007). Image segmentation extracts the region of interest (the individual particles) from the image background, which is a key step for extracting useful information after image acquisition. Aggregate imaging systems are usually configured to ensure spacing among particles such that the effort required to separate overlapping or touching particles during the image segmentation is minimized. The AIMS

system (Masad et al. 2007; Gates et al. 2011) can capture multiple aggregates manually spread onto a tray. Post-processing is required by conducting the convex hull test. The E-UIAIA system (Tutumluer et al. 2000; Moaveni et al. 2013) acquires aggregate photos from orthogonal views of individual particles only. Huang et al. (2019) recently developed a field imaging system for large-sized aggregates following a similar design concept. These variations of UIAIA systems deal with single-particle imaging with no touching or overlapping involved. Other imaging systems such as a 3D laser-based system by Anochie-Boateng et al. (2013) and a stereophotography-based system by Zheng and Hryciw (2017) mainly focused on aggregate particles with minimal contact or overlap.

The above aggregate imaging systems manually control the arrangement of particles and achieve high-precision measurements of separated or nonoverlapping aggregates. However, when aggregates are in a densely stacked or stockpile form, which are the more practical scenarios, their capability to simultaneously characterize a large quantity of aggregates may not be sufficient. First, these systems manually separate the particles and provide a constant background to simplify the image-segmentation task. This condition can no longer be satisfied when aggregates are in a stockpile background or other field scenes. Second, manually feeding many particles into these systems is inefficient for time-sensitive applications. Moreover, the application of these advanced imaging systems is further limited when only in-place or field evaluation is available at quarry and construction sites or when characterization of large-sized aggregates is needed.

To overcome the challenges of analyzing stockpile aggregate images, more advanced image-segmentation techniques are required. Traditional 2D image-segmentation methods include three major types—region based, edge detection based, and watershed—among which the variations of edge-based and watershed segmentation algorithms have been shown to perform better in the presence of mutually touching particles in dense images such as stockpile aggregate views (Wani and Batchelor 1994; Senthilkumaran and Rajesh 2009; Vincent and Soille 1991). In this connection, several research software systems and industrial applications have been developed. For example, Tutumluer et al. (2017) and Huang et al. (2018) applied watershed segmentation to characterize the degradation level in trench-view images of railway ballast by classifying the size distributions of image segments. Similarly, the commercial cross-platform software WipFrag, developed by Maerz et al. (1996, 1999), uses edge-based segmentation to partition rock fragments and estimate the particle size distribution in a stockpile image. Nevertheless, both image-segmentation algorithms used in the software programs are user-dependent. Considerable user interaction for fine-tuning parameters and interactive editing is required to achieve an acceptable segmented image. To this end, the state-of-the-art implementation in stockpile aggregate image analysis is the deep learning–based approach developed by Huang et al. (2020b), which trains an image-segmentation network called Mask R-CNN on a manually labeled stockpile image dataset.

However, the available approaches for 3D aggregate analysis are very limited and mostly focused on single-particle analysis. To fully reconstruct the aggregates as 3D models, many 3D scanning-based approaches have been developed in the past decade. Anochie-Boateng et al. (2013) and Komba et al. (2013) designed and developed a 3D laser scanning device to obtain 3D aggregate models by a spot-beam triangulation scanning method. Jin et al. (2018) constructed 3D solid models of nine aggregates by merging X-ray CT slices from the cross sections of the specimens. Complicated searching and

merging algorithms were developed to orient the CT slices to form valid 3D shapes. Thilakarathna et al. (2021) used a structured light 3D scanner to reconstruct 3D models by projecting preset light patterns onto the aggregate surface. These 3D scanning-based approaches usually utilize expensive scanning devices and require external lighting sources. Alternatively, more convenient and cost-effective photogrammetry approaches were investigated and demonstrated a comparable reconstruction quality when compared to the approaches requiring expensive imaging devices.

Paixão et al. (2018) reconstructed 18 ballast particles by fixing the individual aggregates with a support pedestal and obtaining all-around views at three elevations. The particle sizes were below 3 in. (7.6 cm) to ensure stable support from the pedestal. The photogrammetry results were compared with the results from 3D laser scanning, and both methods demonstrated very close results. Ozturk et al. (2020) followed a similar photogrammetry procedure that captures all-around views from different viewing angles when the aggregate particle is glued to a stick and elevated in the air. The particle sizes were around 0.5 in. (1.3 cm) to be stably fixed using glue. Both researchers used a support system to elevate the aggregate in the air so that all-around views are accessible. The size range of aggregates that can be reconstructed by the procedure is greatly limited by the design of the support system. More importantly, an effective 3D analysis approach that can apply to aggregate stockpiles is lacking in the aggregate research domain.

DEEP LEARNING FACILITATED COMPUTER-VISION TECHNIQUES

Over the last decade, machine learning–based methods have enabled significant advances in many complex vision tasks benefitting from the development of artificial intelligence and computer vision techniques (Prince 2012; Goodfellow et al. 2016). Dense image-segmentation tasks, along with many object classification and detection tasks, are difficult in the sense that the features in the image are usually implicit and thus cannot be easily extracted and represented by human means. While traditional segmentation methods are not effectively applicable to identifying these features, machine learning methods may better handle such tasks by capturing the underlying features based on data-driven mechanisms. Zheng and Hryciw (2017) proposed an approach using pattern recognition to identify particles in images of sand assemblies. The approach followed the Voila-Jones object detection framework, which is a combination of traditional feature extraction techniques and machine learning algorithms. However, this approach still requires careful design of a feature extractor that transforms the raw image into effective internal representation.

During recent developments in the computer vision domain, a deep learning framework proposed by LeCun et al. (2015) exhibits several advantages over conventional machine learning techniques because of its better capability of discovering intricate structure in large datasets with minimal human-guided interaction. With multiple levels of abstraction in the neural network, deep learning has dramatically improved the state of the art in many complicated tasks in computer vision, such as image classification, object detection, semantic segmentation, etc. The power and benefits of deep learning techniques in 2D aggregate image analysis have been demonstrated in Huang et al. (2020a), and the similar mechanism is expected to apply to 3D aggregate analysis. Considering this fact, applying deep learning techniques to 3D stockpile aggregate analysis tasks is promising.

SUMMARY

This chapter provided a review of riprap standards and specifications, findings from previous aggregate studies, relevant equipment that leveraged imaging techniques, and the applications of machine learning–based technology in machine vision.

Traditional methods for assessing riprap geometric properties involve subjective visual inspection and time-consuming manual measurements. As such, achieving the comprehensive in situ characterization of riprap materials remains challenging for practitioners and engineers. In this regard, several advanced aggregate imaging systems developed over the years utilized computer vision techniques to approach this task in a quantitative, objective, and efficient manner.

In terms of imaging-based approaches, aggregate imaging systems developed to date for size and shape characterization have primarily focused on 2D analysis of separated or nonoverlapping aggregate particles. The development of efficient computer vision algorithms based on emerging deep learning techniques is urgently needed for 3D evaluations of aggregates in densely stacked and stockpile forms.

CHAPTER 3: 3D RECONSTRUCTION FOR FIELD DATA OF INDIVIDUAL AGGREGATES AND AGGREGATE STOCKPILES

INTRODUCTION

Given the challenges presented in the previous chapter, a 3D aggregate particle database/library would serve as the cornerstone for any development related to 3D aggregate segmentation research. This chapter presents the development of a photogrammetry-based 3D reconstruction approach for obtaining models for individual 3D aggregates and aggregate stockpiles. Based on this approach, the researchers will establish a 3D aggregate particle database that is used for the development of segmentation and completion algorithms.

MARKER-BASED 3D RECONSTRUCTION APPROACH

First, a convenient and cost-effective procedure for the 3D reconstruction of individual aggregate particles from multi-view images was developed. The proposed photogrammetry approach follows a marker-based design that enables background suppression, point cloud stitching, and scale referencing to obtain high-quality aggregate models. The approach allows reconstruction across flexible size ranges (especially for relatively large-sized aggregates) and is potentially extensible to work under field conditions as well. The equipment setup, reconstruction mechanism, and key designs of the reconstruction approach are detailed herein.

Equipment Setup

The equipment of the reconstruction system includes a digital camera, a camera tripod, a 12 in. (30.5 cm) diameter turntable, and a white cardboard background, as shown in Figure 2. The digital camera used in this study was a smartphone camera (Model: iPhone XR) with a 4032x3024 pixel resolution, but other types of digital cameras can also be used if the collected images are of sufficient quality and resolution. The camera was mounted on the tripod at a viewing angle of 30 to 45 degrees with respect to the horizontal plane. A proper viewing angle ensures the top and side surfaces of the inspected aggregate particle are visible to the camera. During reconstruction, the camera was at a fixed position, and the multi-view images of the aggregate were obtained by manually rotating the turntable. The smartphone camera was programmed with an automatic shutter (with a beeping sound) to take an image every two seconds. In between two shutters, the operator rotates the turntable around 30 degrees and switches to the next view. Note that the use of a turntable and a white background with a fixed-position camera is one of many possible setups to collect multiple views. The approach is flexible and designed to accommodate different configurations. For example, when applying this approach to larger aggregates that cannot easily fit onto a turntable, or a turntable is not available for field inspection, it is recommended to acquire multi-view images by moving the camera around the static object.

Figure 2. Photo. Equipment setup for 3D reconstruction of aggregates.

3D Reconstruction by Structure from Motion

In the computer vision domain, the structure-from-motion (SfM) technique is a powerful photogrammetry method for 3D reconstruction of static scenes. The previous photogrammetry-based methods used by aggregate researchers (Paixão et al. 2018; Ozturk and Rashidzade 2020) also belong to the SfM category. SfM solves the problem of recovering a 3D stationary structure from a collection of multi-view 2D images. A typical SfM pipeline involves three main stages: (i) extracting local features from 2D views and matching the common features across views; (ii) estimating the motion of cameras and obtaining relative camera positions and orientations; and (iii) recovering the 3D structure by jointly minimizing the total re-projection error (Longuet-Higgins 1981; Andrew 2001). The fundamentals and implementation of SfM are omitted from this discussion, but the key steps, i.e., steps (ii) and (iii), are discussed herein with necessary details. The process of simultaneously estimating the camera parameters and the 3D structure is also called bundle adjustment, which is essentially an optimization problem, as shown in Figure 3:

$$\underset{\{\mathbf{P},\mathbf{X}\}}{\text{minimize}} \sum_{i=1}^{m}\sum_{j=1}^{n}\|\mathbf{x}_{ij}-\mathbf{P}_i\mathbf{X}_j\|^2$$

Figure 3. Equation. Re-projection error during bundle adjustment.

where P_i is the projection matrix of the i^{th} camera, X_j is the coordinates of the j^{th} feature point in the 3D structure, and x_{ij} is the projected pixel location of X_j in the i^{th} camera view. The total re-projection

error, the objective function in Figure 3, is the squared pixel distance of all feature points across all camera views. The bundle adjustment process then iteratively finds the best estimates of the camera parameters and the point coordinates by minimizing the objective. After convergence, the reconstructed structure is available as a sparse 3D point cloud and can be further processed to generate a dense point cloud.

Background Suppression by Masking for Noise Reduction

The standard SfM procedure extracts features from the whole 2D images and attempts to reconstruct the entire scene, as shown in Figure 5(a). This usually results in a 3D model that requires manual cleaning to remove unrelated background information (noise) and to obtain a clean model of the aggregate sample. Depending on how much of the background is reconstructed, the manual cleaning process could become considerably time-consuming, especially in regions where the aggregate is touching the background surface, as illustrated in Figure 6(a). Note that this manual cleaning requirement is not only limited to the SfM procedure. During 3D reconstruction with costly devices (i.e., laser scanner, structured light scanner, etc.), manual cleaning is also a necessary step. This is because the scanning mechanism does not distinguish the foreground from the background, because their relative definition will vary from one application to another. To reduce noise from unrelated background regions, the proposed approach improves the standard SfM approach by generating a foreground object mask, *M*, for each image. During bundle adjustment, the object mask is applied as an additional constraint in the original objective function, as shown in Figure 4:

$$\underset{\{\mathbf{P},\mathbf{X}\}}{\text{minimize}} \sum_{i=1}^{m}\sum_{j=1}^{n} M_{ij} \cdot \|\mathbf{x}_{ij} - \mathbf{P}_i\mathbf{X}_j\|^2$$

Figure 4. Equation. Re-projection error during bundle adjustment with mask constraints.

where M_{ij} is the object mask indicating the inclusion or suppression of feature X_j in the i^{th} camera view.

The generation of this type of foreground object mask is an image-segmentation problem. Although traditional segmentation methods can be applied using color and edge information, the proposed approach adopts a deep learning–based segmentation method. The neural network architecture used is called U^2-Net, which is a successful design for the salient object detection task (Qin et al. 2020). Salient object detection is utilized to detect and extract the potential region of interest (RoI) of objects that may be salient in the image. The network uses deep-nested U-shaped convolutional-deconvolutional blocks to capture multiscale contextual information without significantly increasing the computation cost. The training dataset was image-mask pairs prepared by both manual labeling and 3D to 2D projection of several manually cleaned 3D models. Based on experiments, around 100 image-mask pairs yield very robust and accurate foreground extraction for a given background environment. Note that for a given background environment, the network is trained only once, and no further training is involved in the reconstruction workflow. The raw images and generated foreground masks of an example aggregate are illustrated in Figure 5.

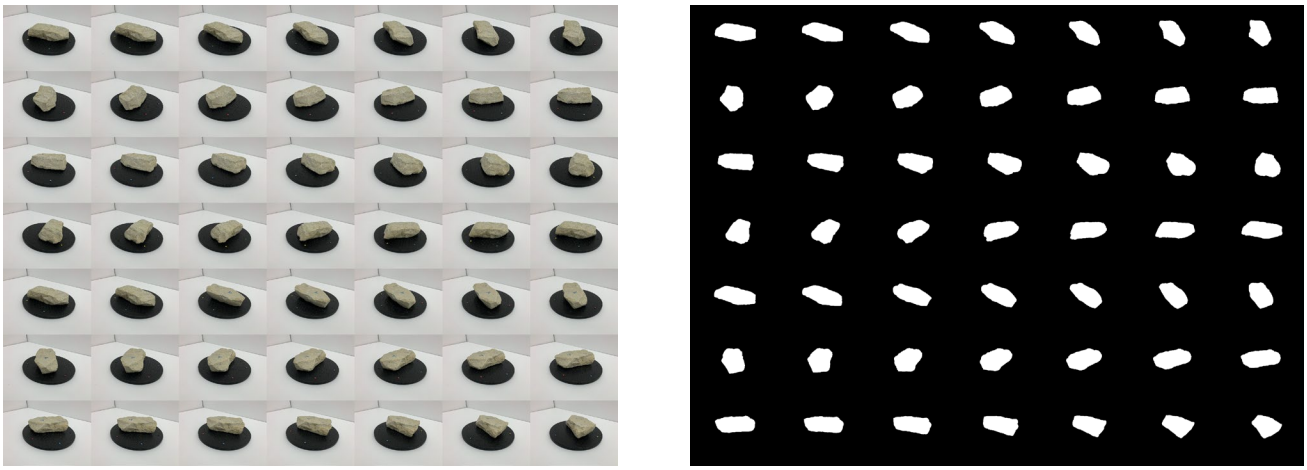

(a) (b)

Figure 5. Illustration. (a) Multiple view images of an aggregate particle and (b) salient object masks for each view.

Adopting a deep learning–based method improves the flexibility of the proposed approach. Although the experiments conducted in this study were set up with a fixed background, the approach is designed to work in different environments, such as using different colors for the turntable and background, or under field conditions with natural lighting conditions. In such cases, a traditional segmentation method may not generate masks robustly, while the method based on deep learning only requires a few image-mask pairs to tune its behavior. The robustness of detection in natural backgrounds has been validated in the original U^2-Net development.

By applying the foreground masks, the unrelated background is suppressed, and the reconstructed model is noise-free and does not require any further manual cleaning. The resulting background suppression effect is illustrated in Figure 6.

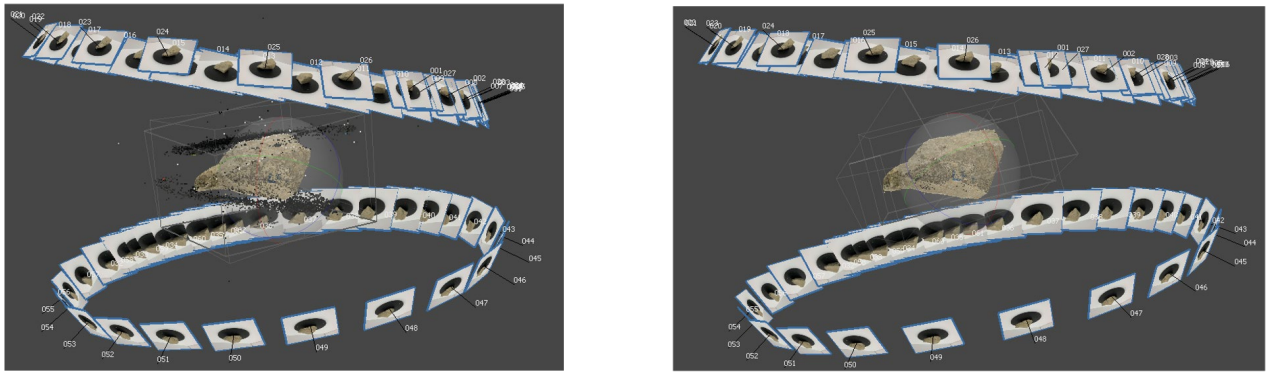

(a) (b)

Figure 6. Illustration. Reconstructed sparse point cloud (a) without background suppression and (b) with background suppression.

Object Markers for Robust Point Cloud Stitching

Unlike small-sized aggregates that can be easily elevated by a support pedestal, medium- and large-sized aggregates usually need to sit on a flat surface during reconstruction or scanning. This limits the possibility of obtaining all-around views of the aggregate and reconstructing with one run of SfM. Two or more rounds of reconstruction are thus required on different parts of the aggregate by adjusting its pose in between, and the partial point clouds must be stitched into a complete 3D model. The most common method to stitch multiple point clouds is to use point set registration algorithms (Choi et al. 2015). However, based on experiments, automatic registration algorithms are not always robust and may fail for certain aggregate samples with less distinct surface features.

In this regard, a set of object markers was designed to provide robust feature-matching during point cloud stitching. Two markers were drawn with colored pencils on the side of each aggregate. The markers were designed to have a head-tail pattern with purple and red colors, as shown in Figure 7(a) and Figure 7(b). Note that the selected colors are not fixed and can be adjusted based on the color of the aggregate for better contrast. The head and tail of each marker are the ends of short and long line segments, respectively. Such pattern is invariant to different viewing angles and can, thus, be identified robustly. After the sparse reconstruction is completed, manual labeling of the markers is required on few views (typically three views) to obtain a consistent localization of the markers in 3D coordinates. When the marker localization is completed for each partial point cloud, the stitching process can be conducted successfully, and a complete 3D model is obtained for the aggregate.

Background Markers for Scale Reference

The reconstructed 3D model from previous steps is in a local coordinate system. To bring the model into a true physical scale and a global coordinate system, a set of background markers was designed to provide a scale reference. The design follows the same concept of ground control points (GCP) in land surveying (Bernhardsen 2002). Color-coding labels in red, green, blue, and yellow were placed at four corners of the turntable, as illustrated in Figure 7(c). The distances between the markers were measured in advance and given as the scale factor. As discussed previously, when the proposed approach is applied without a turntable, the background markers could take other forms such as GCPs.

(a) (b) (c)

Figure 7. Photo. (a) Purple-colored and (b) red-colored object markers for robust point cloud stitching and (c) background markers for scale reference.

Reconstruction Workflow

The reconstruction workflow can be summarized by the following steps:

- Step 1: Preparation (executed only once for each environment). This step involves setting up the equipment, tuning the foreground detection network, and placing the background markers.

- Step 2: Placing the aggregate sample. The sample is placed in the camera view, and object markers are labeled on the side surface.

- Step 3: Capturing visible sides (two or more) of the sample. By rotating the turntable (or moving the camera), multiple view images are taken. The same procedure is repeated for each side. In our experiments, 30 views were taken for each side with a two-second shutter interval, resulting in two minutes per sample for a two-side inspection.

- Step 4: Reconstruction. First, foreground masks are generated from the foreground detection network. Second, SfM is executed using the raw multi-view images and the associated foreground masks. Next, object markers and background markers are labeled on a subset of images (usually three images from each side). Finally, a complete 3D point cloud model is obtained by stitching the partial point clouds together, and an associated 3D mesh model is reconstructed from the complete dense cloud using the screened Poisson surface reconstruction method (Kazhdan and Hoppe 2013).

- Steps 2 to 4 are repeated for each aggregate sample.

The reconstructed results presented in this study were generated by extending the Agisoft Metashape (Agisoft 2021) software program. Note that the implementation of the reconstruction step is not limited to certain software tools. Commercial software programs such as Agisoft Metashape (Agisoft 2021), free software available such as VisualSFM (Wu 2011), or open-sourced software available such as Meshroom (Griwodz et al. 2021) can all be extended to implement the proposed approach. Also note that even though this research study focused on relatively large-sized aggregates, the setup previously shown in Figure 2 is expected to work for smaller sizes such as base course aggregates or ballast without further adjustments.

MATERIAL INFORMATION AND PROPERTIES OF THE 3D AGGREGATES LIBRARY

The outlined reconstruction procedure was used to inspect a set of 46 RR3 aggregate particles and 36 RR4 aggregate particles collected from field site visits to aggregate producers in Illinois. The samples conform to the “RR3” and “RR4” categories based on IDOT specification, which typically refers to aggregates that have weights above 10 lb (4.54 kg). In the specification, the RR1 and RR2 categories refer to small-sized riprap aggregates having the same size ranges as aggregate subgrade material in pavement engineering and ballast material in railway engineering, and the RR3 to RR7 categories are medium- to large-sized aggregates or rocks that are more common in riprap applications.

Example reconstruction results are visualized in Figure 8. The reconstructed models are available in different formats, such as the textured model that preserves the surface color information in Figure 8(a), the mesh model that shows the wireframe of vertex connectivity in Figure 8(b), and the point cloud model with discrete point coordinates in Figure 8(c). An image collage of 40 RR3 aggregate samples reconstructed in this study is presented in Figure 8(d). In terms of geological classification, these aggregate samples are dolomite rocks with white to yellowish colors, as shown in Figure 8(d). Based on the aggregate particle library, necessary datasets were generated to develop a deep learning–based algorithm for 3D segmentation and 3D shape completion.

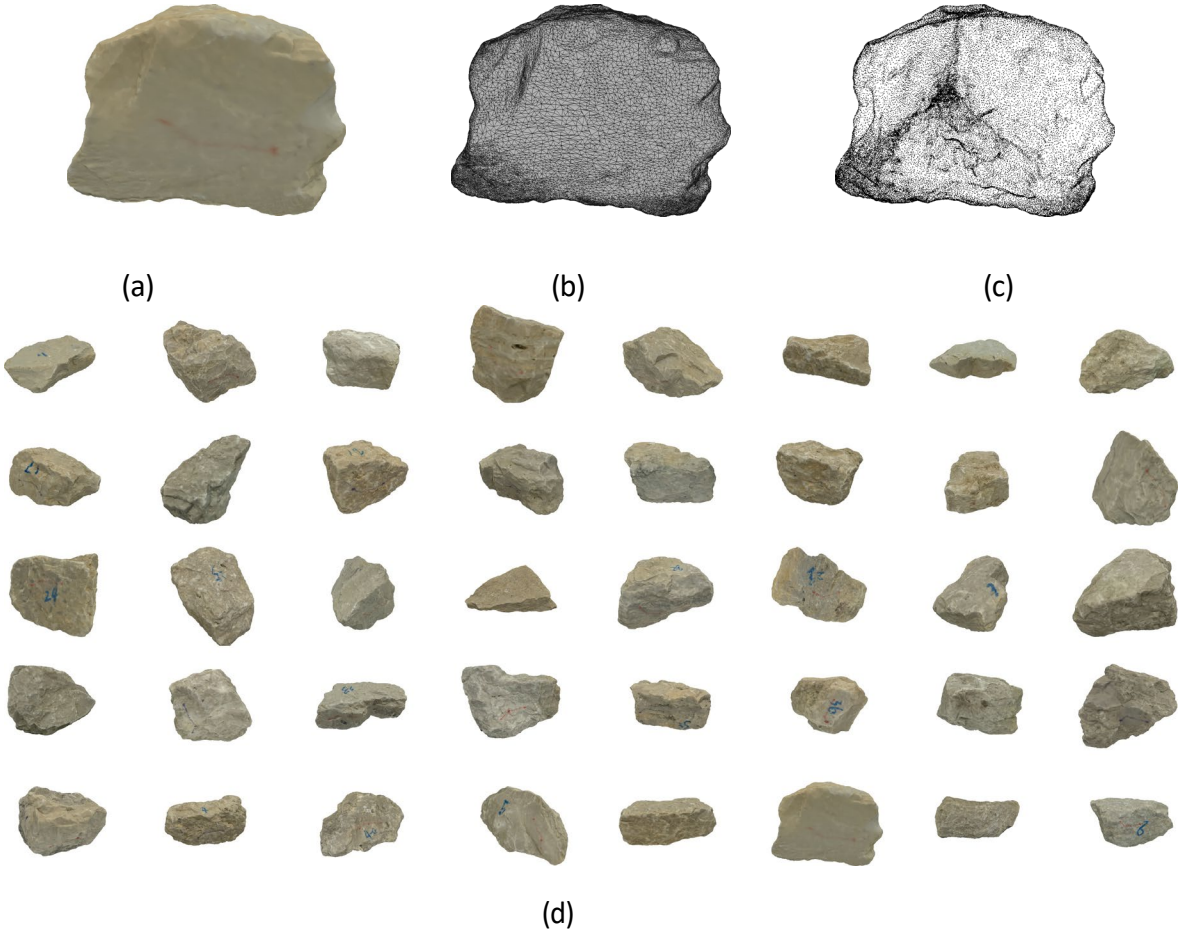

Figure 8. Illustration. (a) Textured model, (b) mesh model, (c) point cloud model of an aggregate particle and (d) collage of 40 reconstructed aggregate particles.

The quality and fidelity of the reconstruction results were assessed from visual effects and quantitative methods. Qualitatively, the reconstructed aggregate models are of high quality and fidelity, as shown in Figure 8. The aggregate models reproduce the geometric features and texture features of the original aggregate samples. Quantitatively, the surface resolution (or point density) of the reconstructed results is considerably high in aggregate research. On average, each sample was exported at a resolution of around 100,000 vertices and 200,000 faces. The surface resolution and point density of 10 example RR3 aggregate particles are listed in Table 3. The resolution is calculated based on the ratio between the number of points in the point cloud model and the surface area of

the reconstructed mesh model. The average resolutions for all 46 RR3 aggregates and all 36 RR4 aggregates are 1.66 points/mm^2 and 0.93 points/mm^2, respectively. The resolution statistics indicate that the aggregate models were reconstructed at a resolution of approximately 1 point/mm^2, i.e., the average distance between adjacent points is around 0.04 in. (1 mm).

For each reconstructed aggregate particle, the basic 3D properties can be calculated from the 3D mesh model, including volume, surface area, and the shortest, intermediate, and longest dimensions in the three principal axes. The 3D properties of the 10 selected RR3 aggregate particles are listed in Table 4. If the intermediate dimension is denoted as the nominal size of an aggregate, then the sizes of these aggregate samples ranged from 3 in. (7.6 cm) to 6 in. (15.2 cm). For the ground truth, the submerged volume of each aggregate sample was measured by a water displacement method following ASTM D6473 (2015), listed as the measured volume in the second column in Table 4.

Table 3. Surface Resolution and Point Density of 10 Example RR3 Particles

Rock ID	Surface Area (cm^2)	No. of Vertices	No. of Faces	Resolution (points/mm^2)
1	1308.69	99680	199356	0.76
2	2201.8	209440	418872	0.95
3	2586.81	261948	523884	1.01
4	2108.77	297392	594760	1.41
5	2257.61	151599	303190	0.67
6	1397.78	93369	186734	0.67
7	1664.52	86056	172108	0.52
8	1836.54	134359	268714	0.73
9	2154.91	223307	446594	1.04
10	1607.77	80549	161094	0.50

Note: 1 cm^2 = 0.155 in.^{2}, 1 mm = 0.04 in.

Table 4. Measured Volume, Reconstructed Volume, Area, and Principal Dimensions of 10 Selected RR3 Particles

Rock ID	Measured Volume (cm^3)	Reconstructed Volume (cm^3)	Surface Area (cm^2)	Shortest Dimension (cm)	Intermediate Dimension (cm)	Longest Dimension (cm)
1	1014.9	1042.3	685.32	7.682	13.142	22.695
2	763.5	786.33	537.87	9.308	12.519	17.412
3	601.8	605.04	418.69	9.477	10.075	14.572
4	791.4	795.69	558.41	9.118	10.133	19.925
5	727.6	744.83	503.13	9.803	10.649	18.842
6	688.1	691.96	478.72	7.497	9.987	15.925
7	644	662.47	465.96	11.614	13.867	14.041
8	1140.5	1165.03	704.29	10.617	12.213	21.923
9	592.7	601.1	435.01	8.068	11.517	17.851
10	890.8	920.92	590.14	10.374	14.513	17.37

Note: 1 cm = 0.4 in., 1 cm^2 = 0.16 in.^{2}, 1 cm^3 = 0.06 in.^{3}

To validate the accuracy of the 3D reconstruction procedure, the reconstructed volume is compared against the measured ground-truth volume, as presented in Figure 10. A 45-degree line is plotted as a reference for the comparison. As the quantitative measure of accuracy, a statistical indicator—the mean percentage error (MPE)—is calculated using the equation in Figure 9. Note that unlike the mean absolute percentage error (MAPE), MPE can have a positive or a negative sign, indicating a systematic overestimate or underestimate behavior, respectively.

$$MPE(\%) = \frac{\sum_{i=1}^{N} \frac{R_i - M_i}{M_i}}{N}$$

Figure 9. Equation. Mean percentage error (MPE).

where R_i is the reconstructed result of the i^{th} sample, M_i is the ground-truth measurement of the i^{th} particle, and N is the total number of particles.

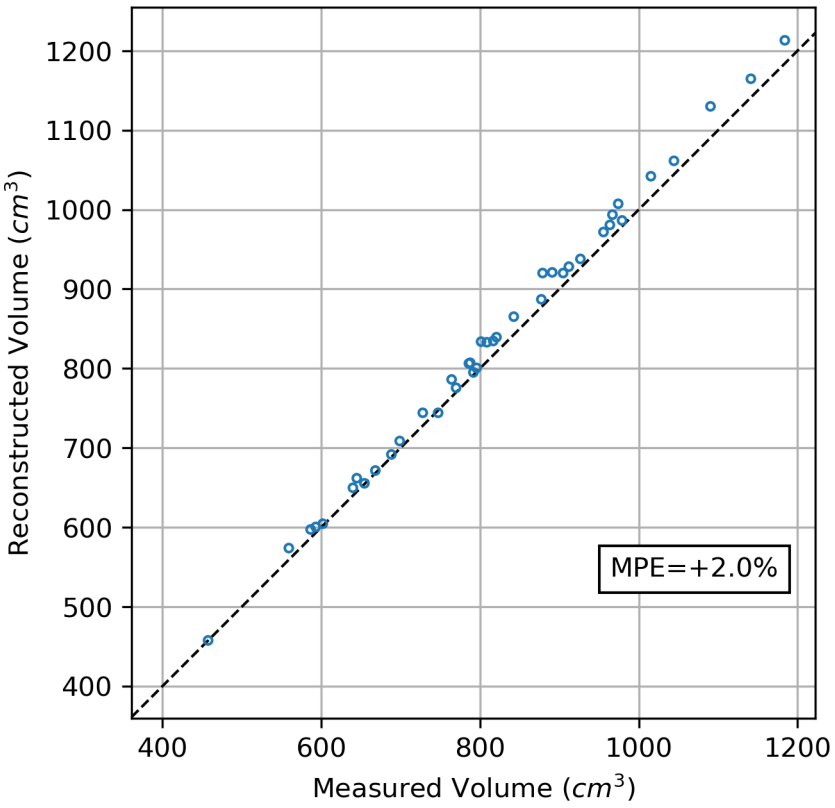

Figure 10. Graph. Comparison of reconstructed and measured volumes of aggregate samples.

Figure 10 shows a very good agreement between the reconstructed volume from the marker-based reconstruction approach and the ground-truth measured volume, with an MPE of +2.0%. The positive MPE also indicates a consistent, systematic overestimate of the reconstructed volumes. There are three potential reasons for this overestimation. First, the pixel locations of the background markers are used to localize the marker in 3D coordinates. Therefore, pixel deviation when labeling the background markers may lead to a slight change of the scale reference. Second, a porous surface

condition was observed on these dolomite aggregate particles, and the micro-texture areas that are filled with water during the measurement of the submerged volume may be reconstructed as flat faces. This could also lead to a systematic overestimate of the true submerged volume. Third, since SfM-based photogrammetry methods entail an optimization approach to jointly approximate the true object geometry, and cameras provide sparser representation (pixels) than laser scanning devices, it is reasonable to assume that certain systematic deviation may exist within acceptable accuracy. Also, the mesh reconstruction from point cloud is an approximation algorithm that may bring systematic deviation near the true surface of the aggregates.

Other morphological indices that can be calculated based on the 3D reconstructed models include the 3D flat and elongated ratio (FER) and 3D sphericity. As the 3D counterpart of the aspect ratio indicator, 3D FER can be calculated after determining the minimum volume bounding box of the particle. O'Rourke (1985) developed algorithms to find the minimal enclosing box of a point set. First, for each possible direction that originated from the particle centroid, a 3D local coordinate frame is formed in the orthogonal searching directions. Next, for each orthogonal pair, the three edge-to-edge distances (3D Feret diameters) within the point set are calculated. The volume of the bounding box can then be computed, and the Feret diameters (Feret 1930) of the minimum volume bounding box are denoted as the shortest dimension "a," intermediate dimension "b," and longest dimension "c." Accordingly, the orthogonal pair associated with the minimum volume bounding box represents the three principal axes of the particle. The 3D FER can then be defined based on the principal dimensions (Figure 11):

$$FER_{3D} = \frac{c}{a}, (c \geq b \geq a)$$

Figure 11. Equation. 3D flat and elongated ratio (FER).

For 3D sphericity, Wadell (1932) defined the sphericity as the ratio between the surface area of an equivalent sphere having the same volume as the particle, S_e, and the measured surface area of the particle, S. This is often called the true sphericity. Given the surface area (A) and the volume (V) of a 3D model, the 3D sphericity can be computed using Figure 12. As a reference, a tetrahedron has a sphericity of 0.67 and a cube has a sphericity of 0.81, again with higher values indicating the 3D shape is closer to a perfect sphere.

$$Sphericity_{3D} = \frac{S_e}{S} = \frac{\sqrt[3]{36\pi V^2}}{A}$$

Figure 12. Equation. 3D sphericity.

3D RECONSTRUCTION OF FIELD AGGREGATE STOCKPILES WITH SCALE REFERENCE

Using a similar marker-based 3D reconstruction approach that works well for individual aggregates, the field aggregate stockpiles can also be reconstructed based on multi-view stereo photography. Different from the previous approach, which is specially designed to obtain a complete model (i.e., two-side reconstruction) of aggregates, the reconstruction of an aggregate stockpile only requires

one pass of 3D reconstruction from the multi-view images collected by walking around the stockpile. Object markers are not needed for this one-side reconstruction either. Background markers, though, are still necessary for providing the scale reference as GCPs.

Under field conditions where the aggregate stockpile could vary in size, the previous design of a fixed-distance marker system no longer applies. To address this issue, a new marker system was designed to provide a flexible scale reference in the field. The marker system consists of three colored blocks (red, blue, and yellow), as shown in Figure 13(a). The top surface of each colored block was marked with a cross sign intersecting at the center, which can be conveniently identified in an image. Figure 13(b) demonstrates the use of the marker system during the 3D reconstruction approach. Before taking the multi-view images, the marker system is placed near the stockpile to form an angle. Note that the principle of using the marker system as GCPs is to form a plane that can be localized in the reconstruction coordinate space; therefore, the three markers should not be co-linear (i.e., approximately lie on the same line). Next, the distances between the markers are measured. In the field experiments, the blue block was used as a pivot marker, and the blue-red and the blue-yellow distances were measured. During the 3D reconstruction approach, by identifying the marker pixel location on a subset of multi-view images and taking the two measured distances as inputs, the reconstructed point cloud of the stockpile can be accurately resized to match the real-world scale.

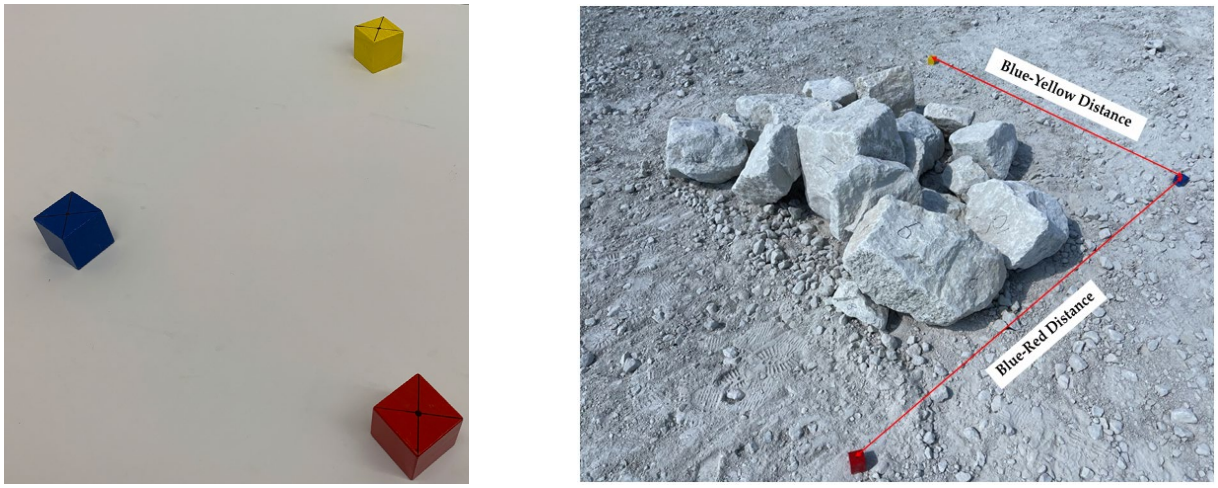

(a) (b)

Figure 13. Photo. Field marker system for scale reference.

Two different types of field stockpiles were inspected. First, stockpiles of different size categories were created at ICT in Rantoul, IL. At each time, all 46 RR3 aggregate samples were used to build a RR3 stockpile, and multi-view images were taken for the 3D reconstruction. After the image acquisition step was completed, the aggregate samples were randomly permuted (e.g., rocks buried inside the current stockpile were placed preferably on the surface for the next stockpile) to vary the stockpile configuration. As a result, six RR3 stockpiles were built, and the multi-view image data was acquired, denoted as stockpiles S1 to S6. The same process was repeated to build six RR4 re-engineered stockpiles based on the 36 RR4 rocks.

For the second type, field stockpile images were collected during site visits to Illinois aggregate quarries at Rantoul and Kankakee. RR3, RR4, and RR5 stockpiles were prepared at these sites manually (for RR3) and by front loader trucks (for RR4 and RR5). At the beginning, 24 RR3 rocks, 16 RR4 rocks, and 20 RR5 rocks were selected in the field. Similar to the first type of stockpile prepared at ICT, numbers were marked on many faces of each rock, and weight measurement was performed to obtain the ground-truth data. Then, the front loader truck moved the aggregate rocks to form a stockpile and permute the stockpile after the multi-view images were collected. This process was repeated three times for each RR category. To distinguish from the first type of stockpile prepared at ICT, these stockpiles are denoted as RR3R, RR4K, and RR5K, where letters "R" and "K" indicate the source locations of the field stockpiles. The information of all stockpiles is listed in Table 5.

Table 5. Information of Re-engineered and Field Stockpiles

Size Category	Number of Aggregate Samples in Stockpile	No. of Stockpiles	Ground Truth
RR3 (Re-engineered)	46	6	Morphological Properties
RR4 (Re-engineered)	36	6	Morphological Properties
RR3R (Field)	24	3	Weight
RR4K (Field)	16	3	Weight
RR5K (Field)	20	3	Weight

Examples of the 3D reconstruction results are presented in Figure 14 for S1 stockpiles in each category. The ground surface was also reconstructed along with the stockpile but was manually removed as a preprocessing operation for the 3D segmentation step. Depending on the operator, the number of multi-view images collected per stockpile ranges from 26 to 50 images. Note that the reconstructed clouds are of consistently high quality for the presented results and for other stockpiles. Based on this practice, around 36 to 50 multi-view images are considered as the recommended number of images to be collected for an all-around inspection of aggregate stockpiles.

SUMMARY

This chapter describes the developed 3D reconstruction approaches for individual aggregates and aggregate stockpiles. A marker-based 3D reconstruction approach was developed to establish the 3D aggregate particle database as a cost-effective and flexible procedure to allow full reconstruction of 3D aggregate shapes. The approach is a photogrammetry-based method with auxiliary designs to achieve background suppression, robust point cloud stitching, and scale reference. The approach was used to inspect a set of 46 RR3 aggregate particles and 36 RR4 aggregate particles collected from field site visits to aggregate producers in Illinois. Based on the aggregate particle library, necessary datasets were generated to develop a deep learning–based algorithm for 3D segmentation and 3D shape completion. Further, the approach was extended to field applications of reconstructing field aggregate stockpiles. For that purpose, an additional scale reference procedure to provide real physical scale of the reconstructed scene was developed.

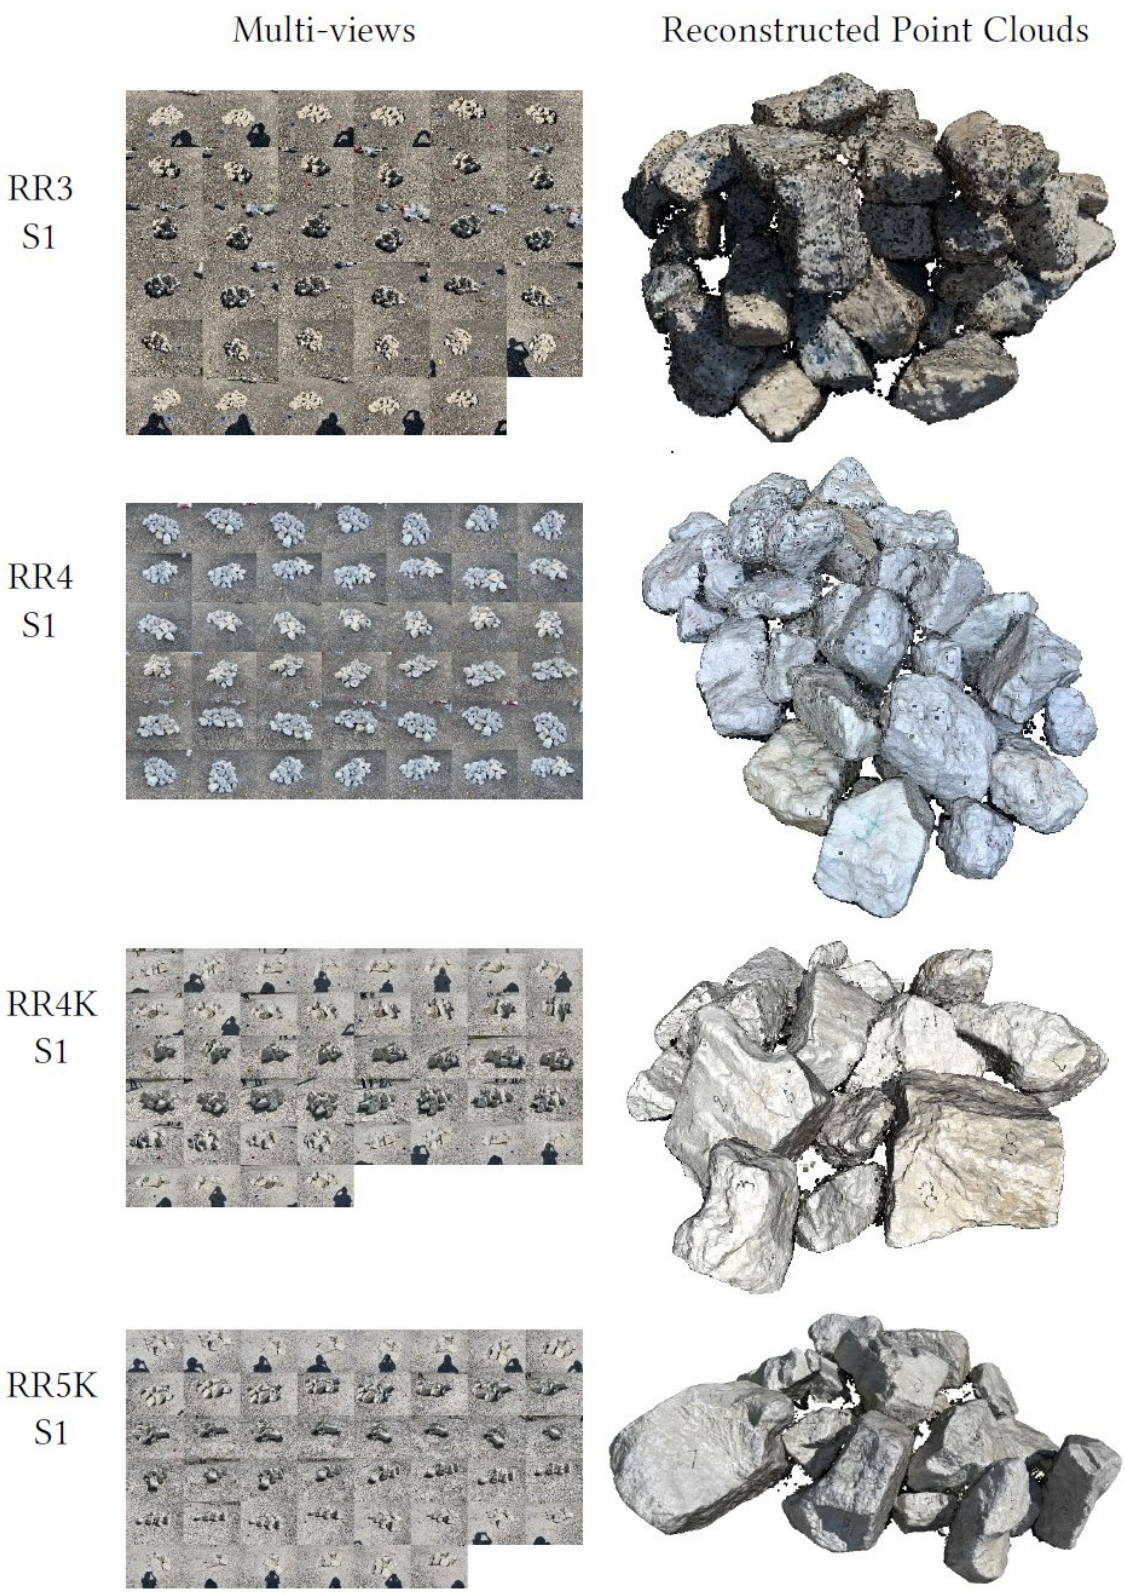

Figure 14. Illustration. 3D reconstruction results of stockpiles in different categories.

CHAPTER 4: AUTOMATED 3D SEGMENTATION OF AGGREGATE STOCKPILES

INTRODUCTION

This chapter first reviews state-of-the-art advancements in computer vision regarding the 3D instance segmentation task and then analyzes the most suitable strategy for the application of dense stockpile segmentation. A selected deep learning framework is then implemented with necessary modifications for the automated stockpile segmentation. Based on the established dataset, the framework is trained to learn the segmentation of individual aggregate instances from the stockpile.

DEEP LEARNING FRAMEWORK FOR AUTOMATED 3D STOCKPILE SEGMENTATION

Review of 3D Instance Segmentation in Computer Vision

Similar to 2D instance segmentation, 3D instance segmentation focuses on detecting and separating objects at the instance level, which is a much harder task than 3D object detection and semantic segmentation. This makes 3D instance segmentation a fundamental yet challenging topic in computer vision that facilitates various types of applications in autonomous driving, robotics, medical imaging, etc. (Guo et al. 2020; He et al. 2021). 3D data provides more comprehensive geometric and scale information than 2D images, especially in the understanding of spatial features and relations. However, unlike 2D images represented in a pixel grid that can naturally be handled by the convolutional neural network (CNN) design, the typical 3D data representation (i.e., point clouds, meshes, voxels) usually presents high disorderliness and irregularity than 2D images. To handle the challenges in 3D instance segmentation, two major categories of methods, i.e., detection-based and detection-free, were developed in the computer vision community.

Detection-based methods are essentially a two-stage approach: (1) detect object proposals and (2) refine the proposals by generating instance masks. These methods usually propose a 3D bounding box of object instances in an explicit way. Because this type of approach imitates the mechanism of human attention by refining from a high-level perception, it is usually depicted as a top-down method. Typical detection-based methods include the Generative Shape Proposal Network (GSPN) and the framework Region-based PointNet (R-PointNet) proposed by Yi et al. (2019), the Gaussian Instance Segmentation Network (GICN, Liu et al. 2020), and 3D-BoNet (Yang et al. 2019).

Different from the detection-based methods, detection-free methods often learn the pointwise features and then apply clustering (or grouping) to obtain instance information. These methods work in the reverse direction of human perception, which first focuses on fine-grained details, and are therefore called a bottom-up method. Representative methods in this category include Similarity Group Proposal Network (SGPN) proposed by Wang et al. (2018), Submanifold Sparse Convolution Network (SSCN, Graham et al. 2018), Multi-scale Affinity with Sparse Convolution (MASC) proposed by Liu and Furukawa (2019), PointGroup (Jiang et al. 2020), and OccuSeg (Han et al. 2020).

Point Grouping Framework with Shifted Coordinates

Based on the review of 3D instance segmentation research in computer vision, it was concluded that detection-free methods are more suitable for the task of aggregate stockpile segmentation. The most important reason is the salience of the aggregate stockpile structure: aggregate stockpiles are point clouds with very densely stacked instances. The 3D instance segmentation datasets in computer vision are mostly available from autonomous driving (Geiger et al. 2013) and indoor environments (Armeni et al. 2016; McCormac et al. 2017), where the separation among object instances is considerably higher than that in a stockpile. Therefore, detection-based methods that strongly rely on the precision of predicted bounding boxes are likely to fail or produce inaccurate results on the densely stacked structure. This observation also agrees with the nature of humans' top-down perception, i.e., they can easily distinguish sparsely separated objects but fail to disentangle small pieces from a pile. Detection-free methods, on the other hand, follow a bottom-up strategy that builds up high-level segmentation from fine-grained details and may better handle the stockpile structure. As a result, a state-of-the-art network, PointGroup (Jiang et al. 2020), was selected, implemented, and customized for the 3D stockpile segmentation task.

The overall architecture of the PointGroup network is illustrated in Figure 15. The network consists of three main components: feature extraction by a backbone network, point clustering on dual coordinate sets, and cluster scoring. The key design in the network is to learn per-point offset vectors to shift the original coordinates into a more compact coordinate space, such that the clustering process will be more robust. The backbone feature extraction network follows a U-Net structure (Ronneberger et al. 2015) with submanifold sparse convolution (SSC, Graham et al. 2018) layers. The original PointGroup architecture uses point colors as additional input features, but in the context of aggregate stockpile segmentation, learning should be based on point coordinates only. This is because unlike the datasets in autonomous driving and indoor environments, aggregates can have high in-class variation in terms of the particle color. Aggregates from different geological origins and experiencing various weathering conditions may have very distinct colors. Considering that instance segmentation of an aggregate stockpile is theoretically plausible by exploring the void space between instances, the geometry information (i.e., point coordinates) is expected to be the crucial input. Therefore, the network was customized to only take point coordinates as the input.

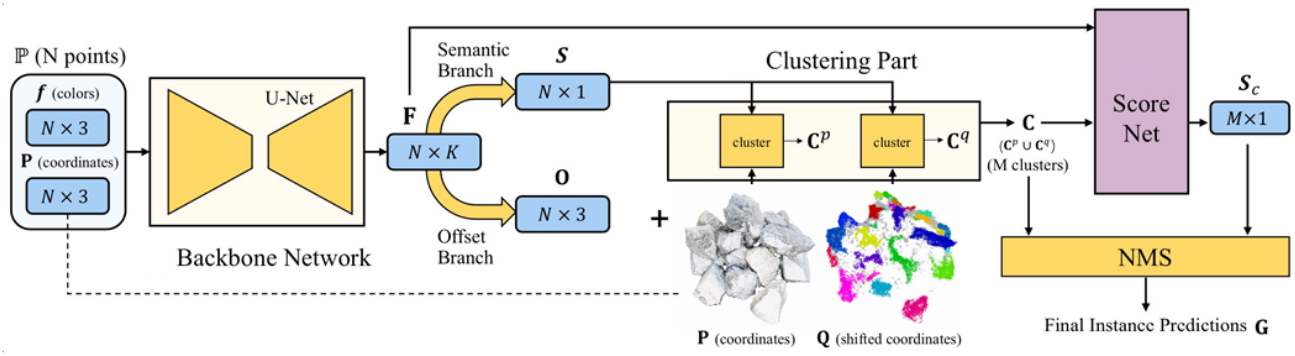

Figure 15. Illustration. PointGroup architecture for instance segmentation.

The per-point offset vectors are then predicted to shift points toward the centroid of its potential instance, as illustrated in the offset branch. Note that the original PointGroup architecture design uses two branches, one for semantic segmentation and the other for predicting the per-point offset vector. In the context of aggregate stockpile segmentation, the semantic branch was removed because the point cloud is expected to contain only single-class aggregate instances. The offset branch predicts per-point offset vector, and the shifted coordinate space can be obtained by applying the per-point offset to the original coordinates. The shifted coordinate space was found to be more efficient for clustering and grouping since the points have now been rearranged in an instance-aware pattern. Based on the original coordinates space "P" and the shifted coordinates space "Q," a clustering step is performed to generate instance proposals. During the experiments, the shifted coordinates were more effective for generating instance proposals. This may be explained by the nature of a dense structure such as an aggregate stockpile, where segmentation on a more compact shifted space is easier than segmentation on the uniformly spaced original representation.

The raw instance proposals may contain many overlapped prediction duplicates as well as low-confidence predictions; therefore, the ScoreNet module is used to rank the clusters. ScoreNet is a sub-network that applies another U-Net structure on the per-point coordinates and feature vectors. As a final step, a 3D version of non-maximum suppression (NMS) is applied to condense highly overlapped instance proposals by selecting the proposal with the highest confidence score among overlapping proposals.

EVALUATION OF STOCKPILE SEGMENTATION PERFORMANCE

The network was trained, and the performance of the instance segmentation was evaluated on the "test set" of the dataset. Qualitative results are presented in Figure 16 and Figure 17. First, the space of the original and shifted coordinates are visualized to indicate the effectiveness of learning the per-point offset. One example is given for each stockpile scene type in the dataset. As shown in Figure 16, the network successfully learned the per-point offset prediction by showing reasonable clustering of the points in the shifted coordinates. Note that each different color in the shifted coordinates "Q" represents the clustered points belonging to individual instances. With the more compact clustered coordinates, the generation of instance proposals is expected to be more robust and reasonable. Across different stockpile scene types in the test set, the network also demonstrates consistent effectiveness and performance in predicting the per-point offset.

Next, the segmentation results were compared with the ground-truth labels in the test set to qualitatively evaluate the segmentation effect, as shown in Figure 17. The final instance proposals are visualized with enclosing bounding boxes to better show the location of the segmented instances. The segmentation results are reasonably satisfactory compared to the ground-truth instances, with most of the aggregate particles identified and successfully segmented. Although some over-segmentation and under-segmentation effects can be observed from the segmentation results, it is considered an efficient and high-quality segmentation surpassing human vision's capability of handling such dense structures.

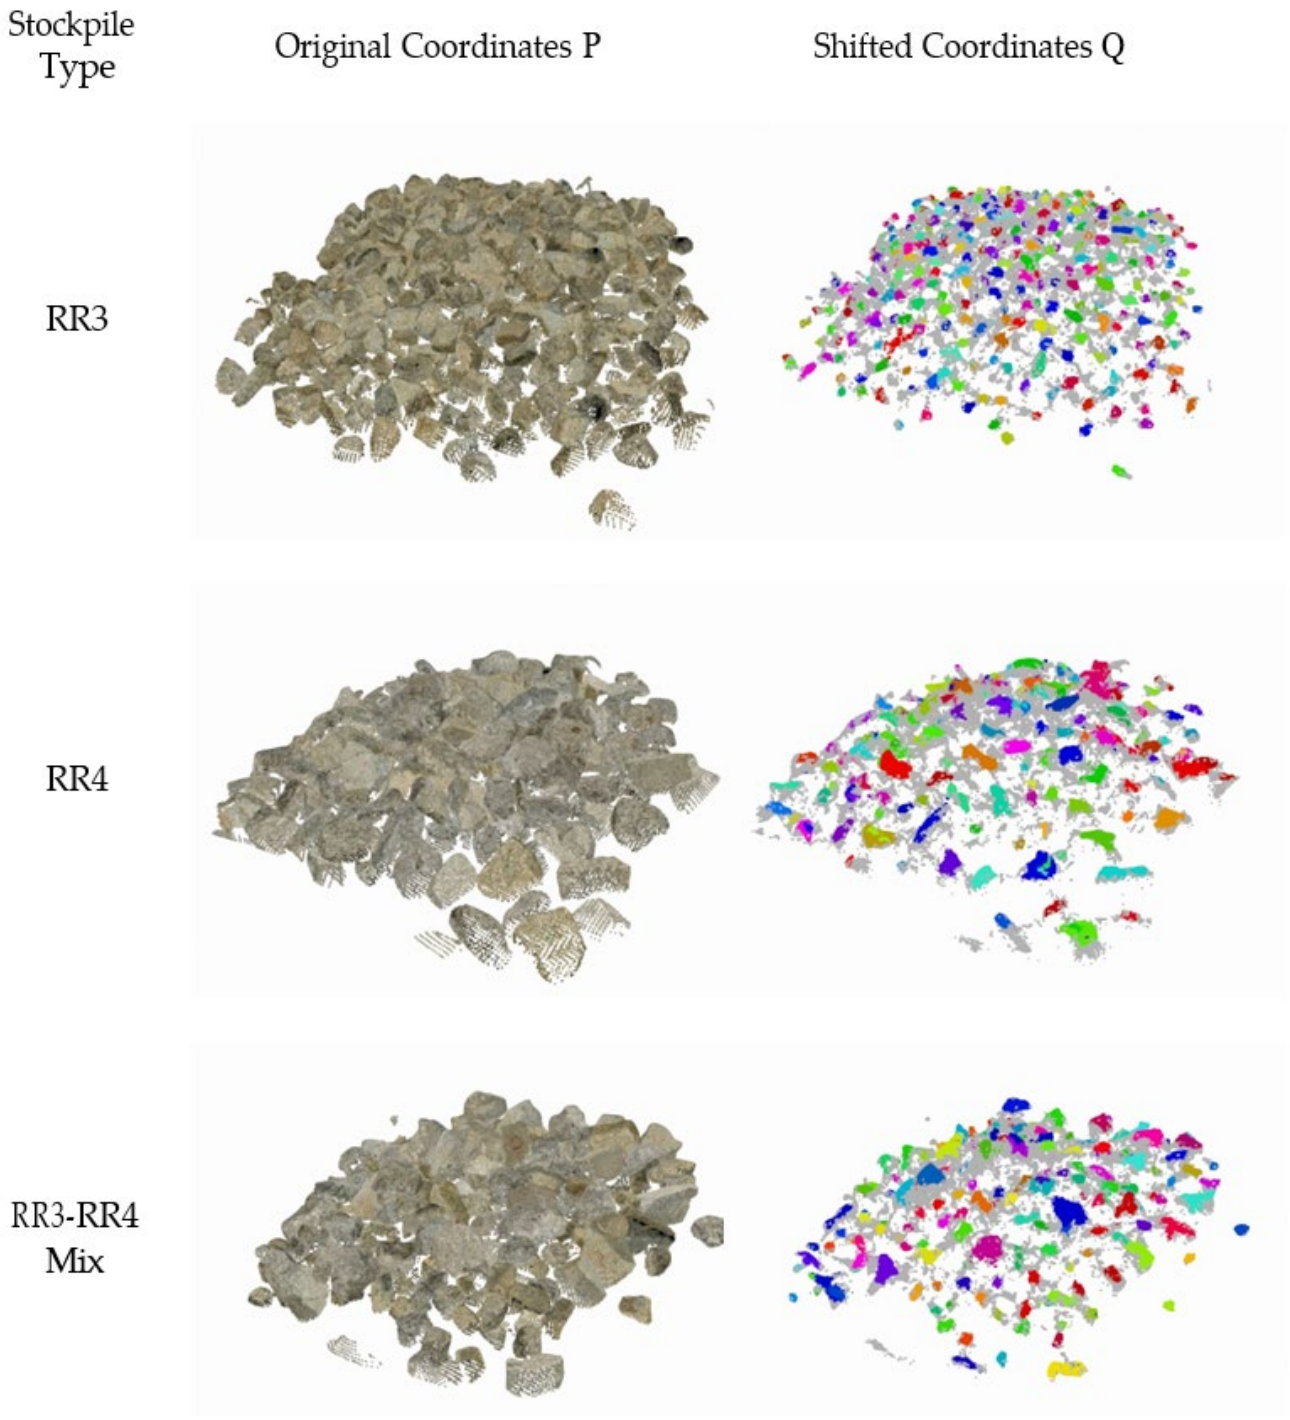

Figure 16. Illustration. Original coordinates (P) and shifted coordinates (Q) by applying the per-point offset.

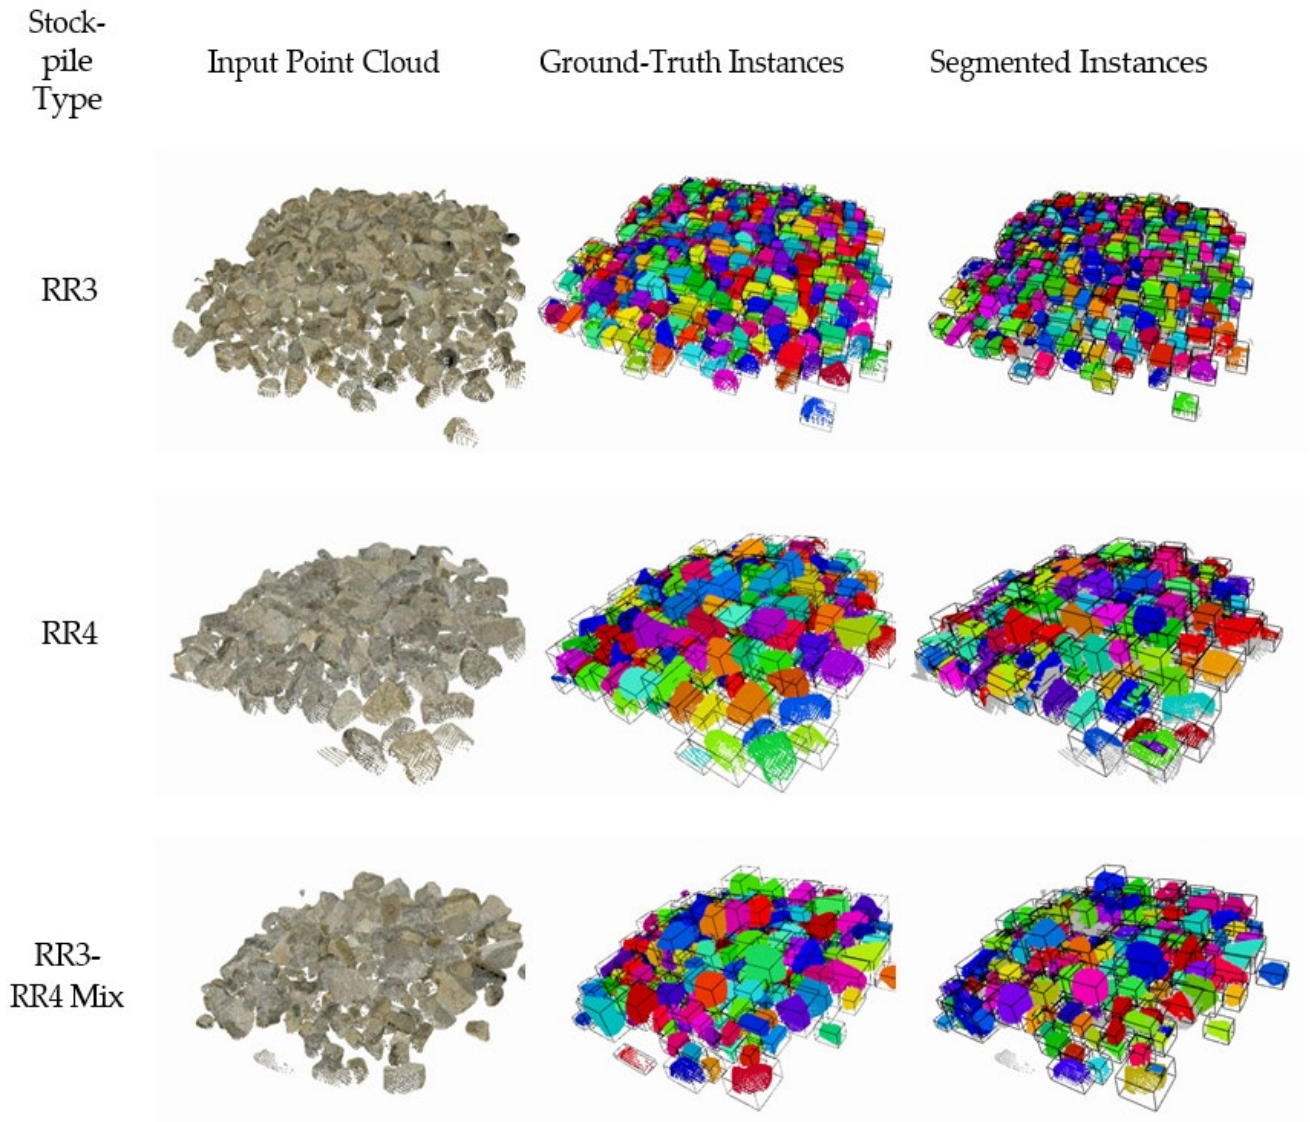

Figure 17. Illustration. Comparisons of segmentation results and ground-truth instances.

Quantitative measurement was also conducted on the quality of segmentation. Before introducing the metrics used for stockpile segmentation, a brief overview is presented for the popular evaluation standard used in machine learning and computer vision research (Powers 2011). By comparing the predictions of a machine learning model and the ground truth, the results can be categorized into four groups: True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN). Positive and negative represent the prediction results, while true and false indicate the correctness of the predictions when compared with the ground truth. For example, TP means a sample is predicted as positive and the prediction is true, i.e., the prediction is consistent with the ground truth. For 2D and 3D instance segmentation tasks that do not have clear true/false correspondences, the definition of the "match" between a prediction and a ground truth commonly follows the Intersection over Union (IoU) concept. 2D IoU for instance segmentation is the number of common pixels between the segmented and ground truth masks divided by the total number of pixels present across both masks.

Similar to 2D IoU, 3D IoU for point cloud data is commonly defined by the intersection and union volumes (*V*) between two axis-aligned bounding boxes of the instance (Figure 18):

$$\mathbf{IoU_{3D}}(\%) = \frac{V_{\text{Segmented}} \cap V_{\text{Ground-Truth}}}{V_{\text{Segmented}} \cup V_{\text{Ground-Truth}}}$$

Figure 18. Equation. 3D intersection over union (IoU).

Therefore, by setting an IoU threshold, the correspondence between prediction and ground truth can be determined, and TP, FP, TN, FN can be defined. Typically, precision and recall metrics are used to measure the performance of the model, as shown in Figure 19. Next, to capture the precision-recall behavior at different threshold IoU values, a precision-recall curve is usually generated by varying IoU thresholds, and an average precision (*AP*) is defined as the area integral under the precision-recall curve.

$$Precision = \frac{TP}{TP + FP}, \; Recall = \frac{TP}{TP + FN}$$

Figure 19. Equation. Precision and recall.

In the context of the aggregate stockpile segmentation task, however, the metrics are customized based on the standard metrics to better indicate the most relevant performance. First, the IoU threshold is fixed at 0.5 to determine the prediction and ground-truth correspondence. At this threshold, "completeness" is defined as the ratio between the number of segmented instances (TP) and the number of ground-truth instances (TP+FN). This ratio describes the percentage of aggregate instances correctly detected as compared to the ground-truth labeling. In fact, the completeness metric herein is identical to the standard recall metric but is renamed to distinguish it in the context of stockpile segmentation. Because a fixed IoU threshold is used, the AP concept no longer applies. However, a metric is needed to further indicate how closely the segmented instances align with the ground truth, even if they all have IoUs beyond the threshold. Therefore, an IoU precision metric is defined as the per-instance 3D IoU score that calculates the percent overlap between the segmentation and the corresponding ground truth. Then, for the entire stockpile, an IoU average precision (IoU AP) metric can be calculated to measure the overall volumetric similarity between the segmented and ground-truth instances. The definition of the two metrics is given in Figure 20, and the demonstrations are shown in Figure 21.

$$Completeness = \frac{TP}{TP + FN}, \; IoUAP = \frac{\sum_{i=1,\ldots,N} IoU_{3D,i}}{N}$$

Figure 20. Equation. Completeness and IoU average precision.

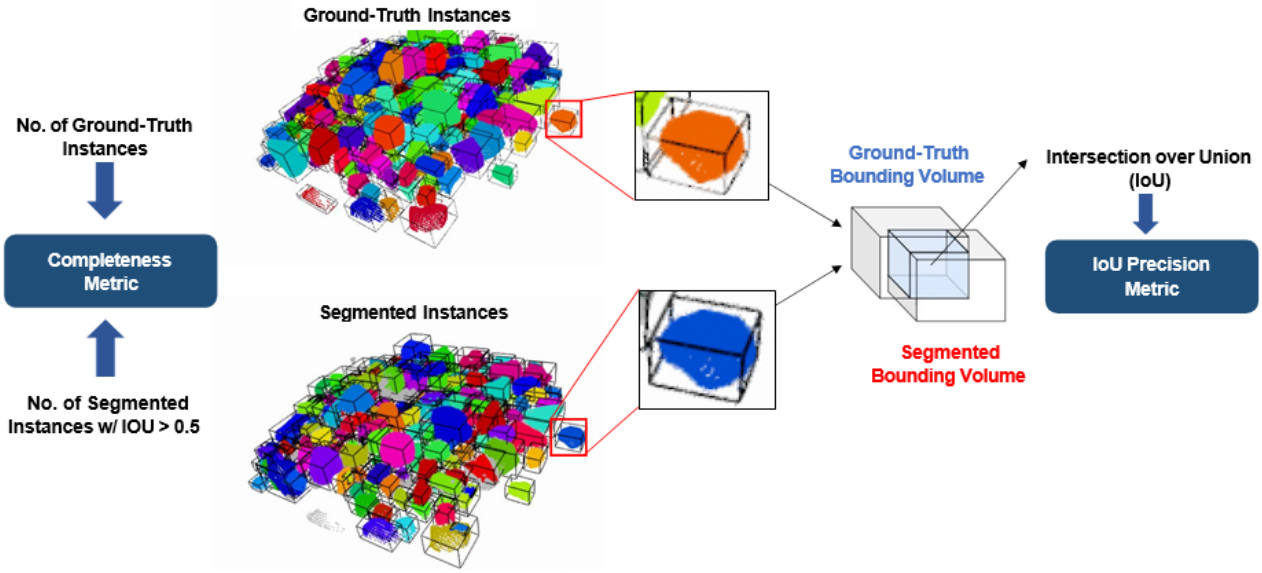

Figure 21. Illustration. Completeness and IoU precision metrics used to compare the segmentation instances with the ground-truth labels.

Following the completeness and IoU AP metrics, the network performance was evaluated on 30 stockpiles from the test set. The average completeness and IoU AP values are 78.4% and 82.2%, respectively, which are considered high for the dense stockpile segmentation task. The average completeness value shows that at least 75% of the aggregates can be successfully identified as compared to the ground truth, with individual instances segmented at a relatively good IoU AP of 82%, on average. Note that the completion metric also indicates that about 20% of aggregates were not segmented. Some of the non-segmented points are likely not recognized as a true instance by human vision because of occlusion and overlapping. Overall, the network demonstrates good performance on the stockpile segmentation task. Moreover, the standard deviation values for completeness and precision are 6.3% and 4.8%, respectively, which implies good generality and robustness of the network among different stockpile scene types.

SUMMARY

This chapter reviewed state-of-the-art advancements in computer vision regarding the 3D instance segmentation task and selected the most suitable strategy for the application of dense-stockpile segmentation. A state-of-the-art deep learning framework was implemented with necessary modifications for automated stockpile segmentation and was trained on the aggregate dataset. Based on the qualitative and quantitative evaluation, the network demonstrated satisfactory performance on segmenting individual aggregate instances from dense stockpiles with considerably high completeness and precision. A more realistic evaluation of the 3D instance segmentation network was conducted on field stockpiles, which will be presented next upon integration with the 3D shape completion component.

CHAPTER 5: AUTOMATED 3D AGGREGATE SHAPE COMPLETION AND FIELD VALIDATION

INTRODUCTION

This chapter introduces the state-of-the-art strategy applicable to learning irregular aggregate shapes. Partial and complete shape pairs are generated from the 3D aggregate particle library based on varying-visibility and varying-view ray-casting techniques. The selected deep learning framework is implemented and trained on the partial-complete shape pairs to learn the shape completion of aggregates. Finally, field validation is conducted after the segmentation and shape completion steps to check the robustness and reliability of the final analysis results.

GENERATION OF PARTIAL-COMPLETE AGGREGATE SHAPE PAIRS

To serve as the dataset for learning a 3D shape completion task, aggregate shapes from partial observations associated with their corresponding ground-truth complete views should be generated and learned in pairs. Establishing such a dataset is usually challenging because it is difficult to obtain partial and complete views for an aggregate at the same time.

A simplified approach could be randomly removing parts from the complete aggregate models, generating incomplete views of the shape. However, this approach is likely to suffer from the following issues. Point cloud is an unordered and irregular data format, so randomly removing points by index may result in an inconsistent effect of removal, i.e., the removed points could be clustering around a certain region or randomly distributed on the original surface. The former emulates the missing parts of partial observations, but the latter merely leads to a nearly uniform down-sampling of the complete shape without missing parts. This limitation could be addressed by intersecting certain shape primitives (e.g., sphere, cylinder, etc.) with the complete aggregate models. Even with this approach, the missing regions of the partial shapes are expected to have many artifacts, such as unnatural cuts, along the shape boundaries.

A more realistic approach was developed by further investigating the cause of partial observations of aggregate shapes. During the reconstruction of an aggregate stockpile, multi-view sensors (i.e., cameras, LiDARs) are commonly used to observe the stockpile surface. Individual aggregates on the stockpile surface may be visible to several sensors simultaneously from different viewing angles. However, the sensors can only occupy the open space around the stockpile, with viewing angles from the other side of the stockpile being missing observations. Based on this fact, the proposed approach was to simulate the sensing process by varying visibility and view to generate realistic partial views that are possible in a real observation.

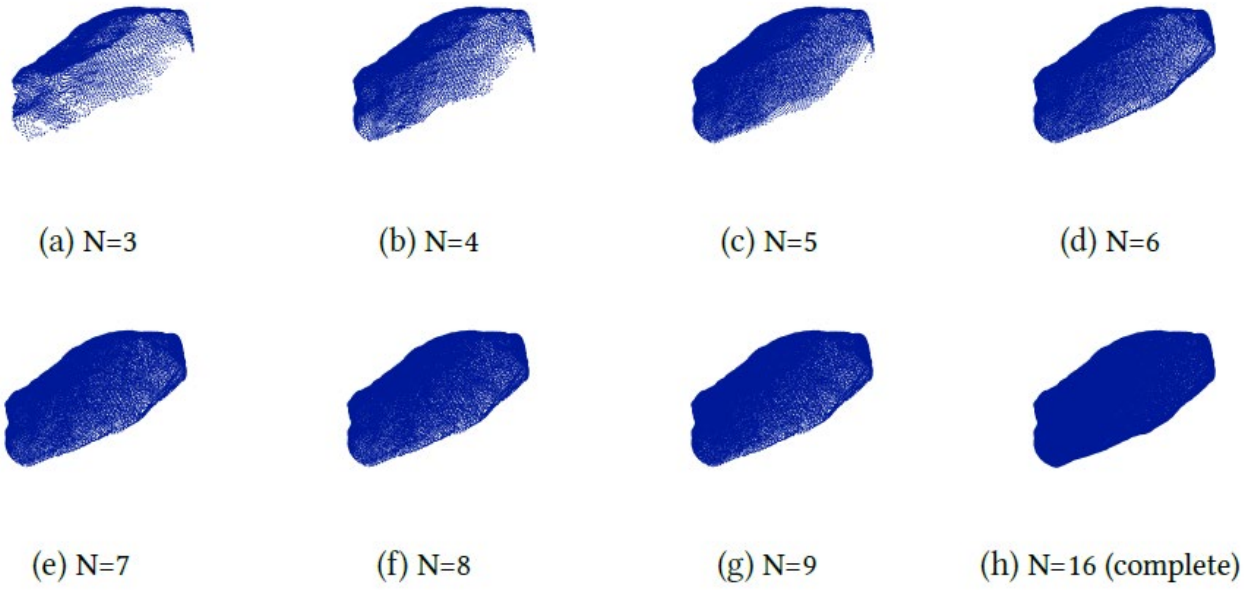

Figure 22. Illustration. Varying-visibility shapes with an increasing number of active sensor views (N).

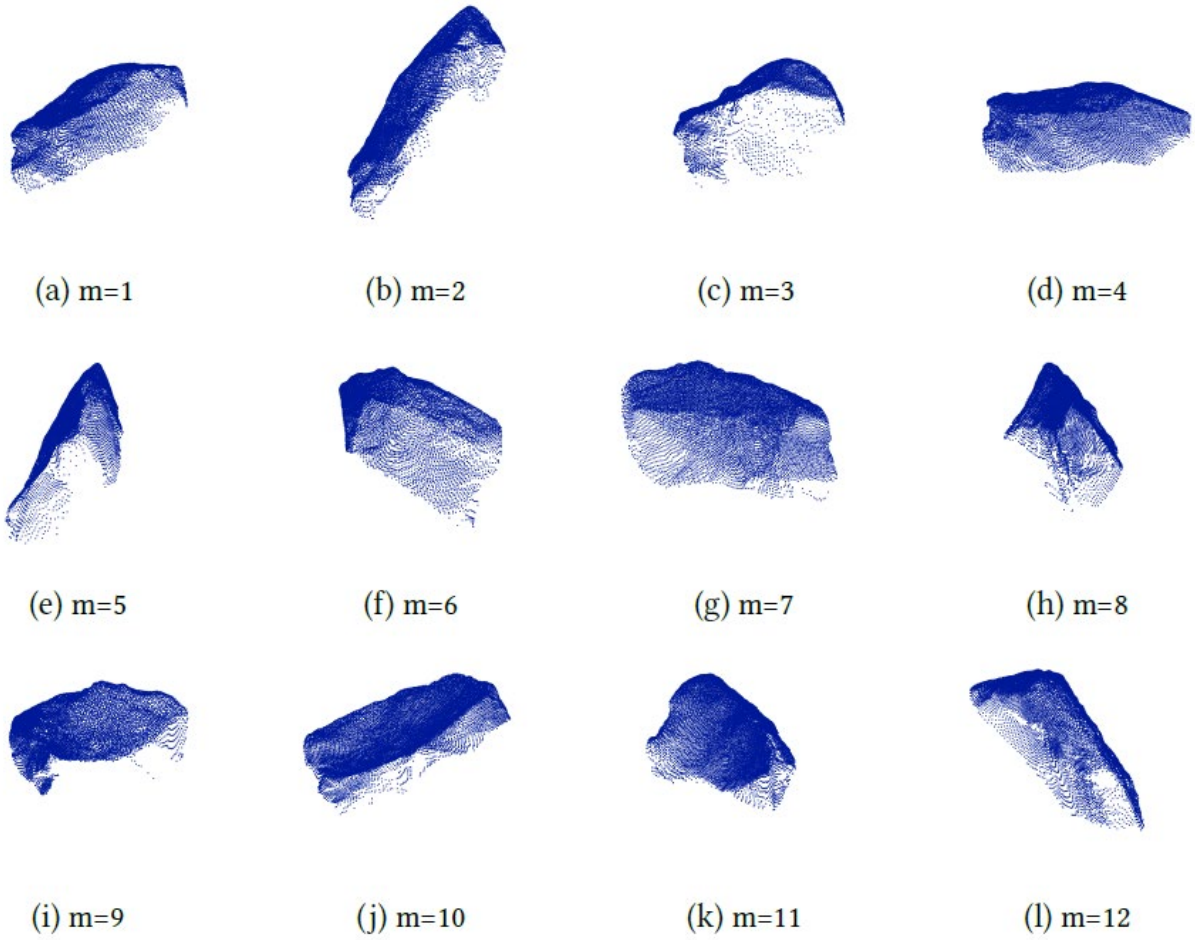

Figure 23. Illustration. Varying-view shapes at m^{th} aggregate model orientation.

As a result, the varying-visibility and varying-view ray-casting schemes together simulate the partial observation process in a comprehensive way, as illustrated in Figure 22 and Figure 23. By treating the all-around sensor view (N = 16) as the ground-truth complete shape, a dataset of partial-complete aggregate shape pairs can be efficiently established. For each of the 82 rock models (46 RR3 rocks and 36 RR4 rocks) in the 3D aggregate particle library, a total of 9,184 partial-complete shape pairs were generated, because one model has seven visibility levels (N = {3, 4, 5, 6, 7, 8, 9}) and M = 16 model orientations. The dataset was further divided into 9,000 training pairs and 184 validation pairs. The validation pairs were randomly selected and separated from the dataset. In addition, to further check the network performance on unseen aggregate shapes, six extra aggregate models were used to generate 672 partial-complete shape pairs. These are RR3 models that were not included in the 3D aggregate library and are therefore considered as an independent "test set" of the shape completion network. Note that the dataset was regularized by uniform down-sampling to 2,048 points per partial shape and 16,384 points per complete shape, which is common fixed data sizes in other popular datasets such as ShapeNet (Chang et al. 2015) and Completion3D (Tchapmi et al. 2019).

DEEP LEARNING FRAMEWORK FOR LEARNING 3D SHAPE COMPLETION

Based on the review of 3D shape completion approaches in computer vision, a state-of-the-art network, SnowflakeNet (Xiang et al. 2021), was selected and implemented for learning the 3D shape completion of aggregates. The overall architecture of SnowflakeNet is presented in Figure 24. The network models the 3D shape completion process as a multi-stage snowflake-like growth of points in space, which consists of three major modules: feature extraction, seed generation, and point generation.

The input of the network is a point cloud with fixed data dimension. The overall shape completion process follows an encoder-decoder approach, where the partial input cloud is condensed into a high-dimensional feature vector by an encoder, and then the decoder generates fine-grained completion by enriching the feature space.

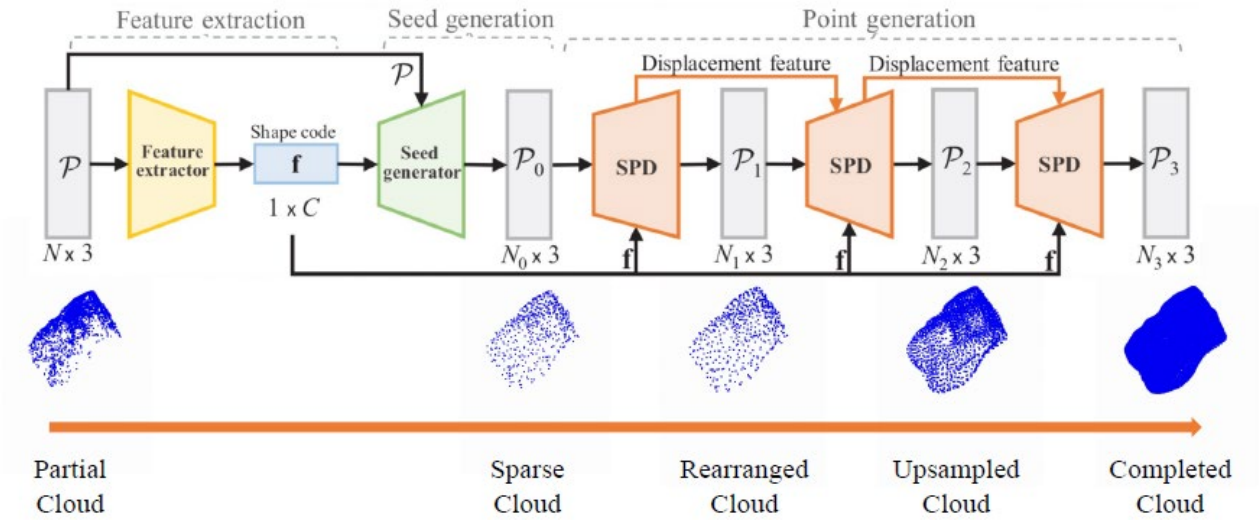

Figure 24. Illustration. SnowflakeNet architecture for 3D shape completion.

For the encoder part, the network uses set abstraction layers developed in PointNet++ (Qi et al. 2017b) and point transformer layers (Zhao et al. 2021) together to encode the global and local shape context into a linear feature vector or shape latent code. This step is denoted as the feature extraction process to obtain high-level shape characteristics with a condensed representation. Although all the training data have 2,048 points, the network can actually take any arbitrary data size since this feature extraction step will first perform regularization to sample the data size down to 512 points following the farthest point sampling (FPS) algorithm proposed in PointNet++ (Qi et al. 2017b). FPS is a shape feature-preserving technique to efficiently reduce the 3D data size while maintaining the prominent features.

After the feature extraction, a two-stage decoder in the network conducts the shape completion task. First, a coarse-grained decoder denoted as the seed generator predicts a sparse version of the complete cloud with 256 points/seeds. This decoder consists of 1D deconvolution (i.e., transposed convolution) layers and multi-layer perceptron (MLP) layers to learn the seed generation, which is referred to as the point splitting operation in the network. The point splitting operation is essentially the 1D deconvolution operation with a large receptive field such that it can capture both existing and missing shape characteristics. The generated seeds are then merged with the input partial cloud to fill the missing portions. However, the merged cloud has nonuniform point density with fewer points in the missing regions. Therefore, FPS is used to re-sample the cloud into a uniform sparse cloud of N_0 = 512 points with the complete shape, denoted as P_0. The overall design concept is similar to a seeded region growing approach, which first focuses on capturing the high-level shape characteristics with sparse representation and then enhancing the shape details as the next step.

Based on the coarse cloud with complete shape, a fine-grained decoder is designed to predict the high-quality complete cloud while preserving the shape features. This decoder uses the snowflake point deconvolution (SPD) layers to up-sample the points by splitting each parent point into multiple child points, which is achieved by first duplicating the parent points and then adding variations to the duplicates. Unlike previous methods that ignore the local shape characteristics around the parent point, SPD utilizes a point-wise splitting operation to fully leverage the local geometric information around the parent point. The key design in the SPD is the skip transformer (ST). With an up-sampling factor of the SPD r, all parent points are first duplicated with r copies. Each point is passed through the ST layer to get per-point displacement feature vectors “K.” Then, an MLP layer computes a per-point coordinate shift, which is added to the original coordinates to get the up-sampled points.

ST uses the PointNet (Qi et al. 2017a) features as query “Q,” generates the shape context feature “H,” and further conducts deconvolution to get the internal displacement features as key “K.” Following the general design of a transformer, per-point query and key vectors are concatenated to form the value vector, and the attention vector is estimated based on the key and value vectors. Note that the attention vector denotes how much attention the old shape characteristics receive during the up-sampling process. The displacement features “K” are carried between SPD operations, which allows the shape context to propagate along the sequential up-sampling process.

By applying SPD with different up-sampling factors, a sequence of gradually refined point clouds can be generated. The up-sampling factors used in the network are r_1 = 1, r_2 = 4, and r_3 = 8. The first SPD

with $r_1 = 1$ generates a rearranged point cloud "P_1" as the same size of the sparse cloud ($N_1 = 512$) but with points slightly rearranged to form a more reasonable shape. The following two SPDs with $r_2 = 4$ and $r_3 = 8$ predicts the up-sampled cloud "P_2" ($N_2 = 2048$) and the final completed cloud "P_3" ($N_3 = 16,384$), respectively.

EVALUATION OF 3D SHAPE COMPLETION RESULTS

Evaluation on Novel Views of Known Shapes

For the partial shapes in the validation set, the shape completion results of three randomly selected inputs are presented in Figure 25. It can be observed that the network can generate high-quality results that agree well with the ground-truth shape for those novel views of the shapes known in the training dataset. This may indicate the network effectively learns the high-level shape representation rather than behaving similarly to a template matching-based approach.

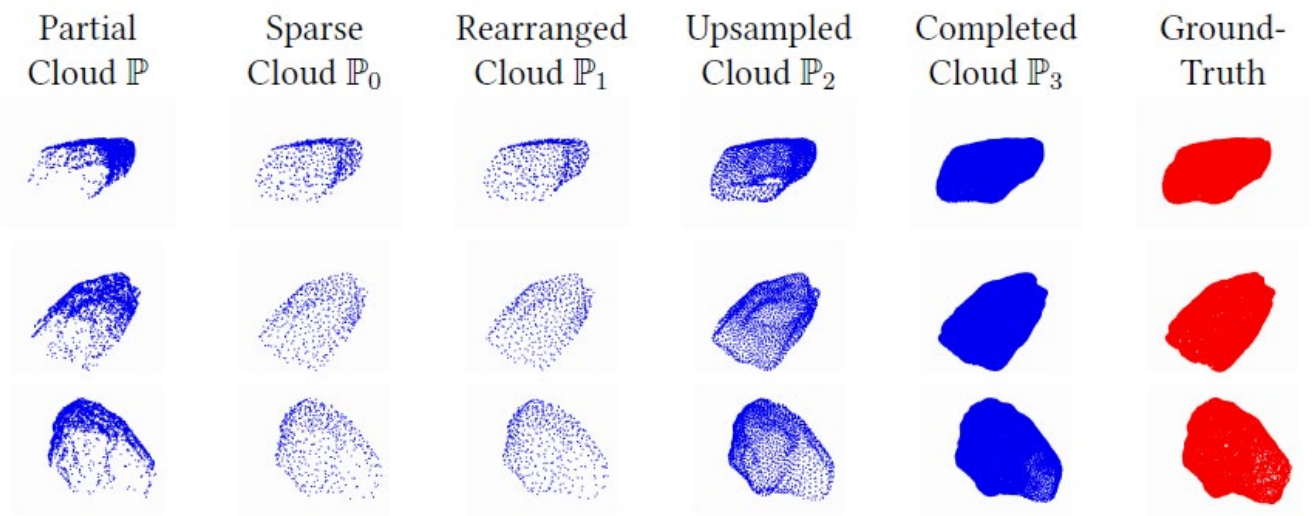

Figure 25. Illustration. Intermediate and final shape completion results for shapes in the validation set.

Macroscale metrics that describe the particle shape at the instance level were used for evaluation. The metrics include equivalent spherical diameter (ESD), shortest/intermediate/longest dimensions, 3D FER, surface area, and volume. Comparisons were made between the completed shape and the ground-truth shape for each of the metrics, as shown in Figure 26. The comparisons of macroscale metrics demonstrate that the completed shapes in the validation set achieve a good match in terms of aggregate morphological properties, where the mean average percent error (MAPE) between the prediction and the ground truth is less than 2.5% for all metrics.

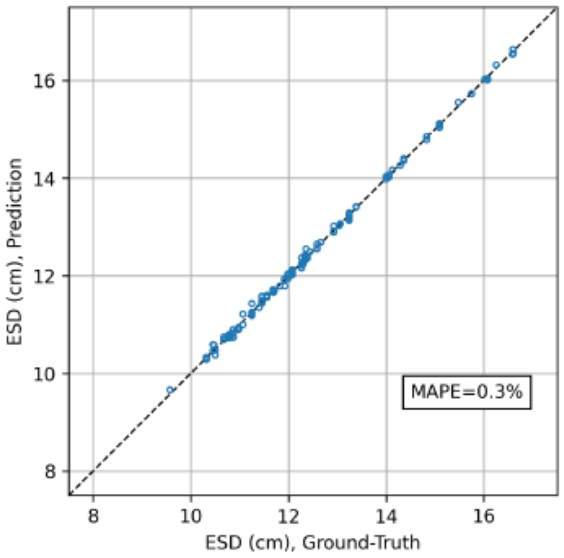

(a) ESD

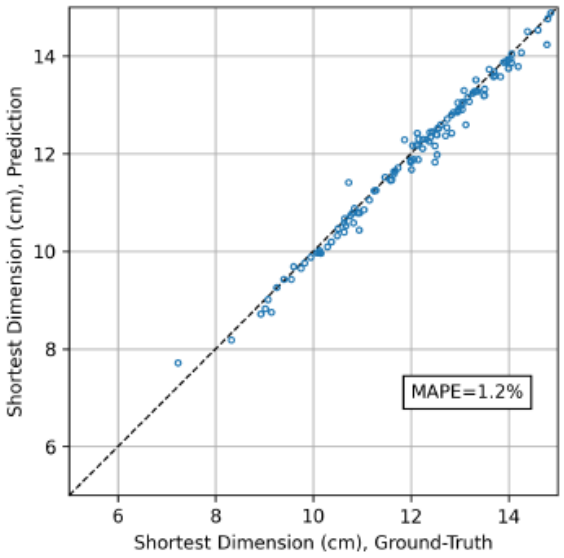

(b) Shortest Dimension

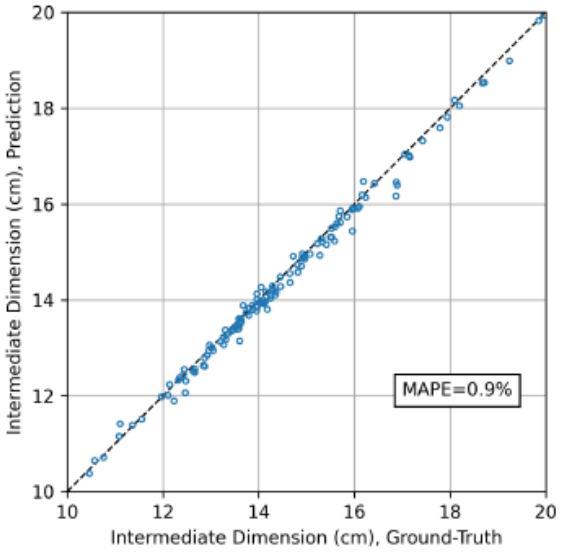

(c) Intermediate Dimension

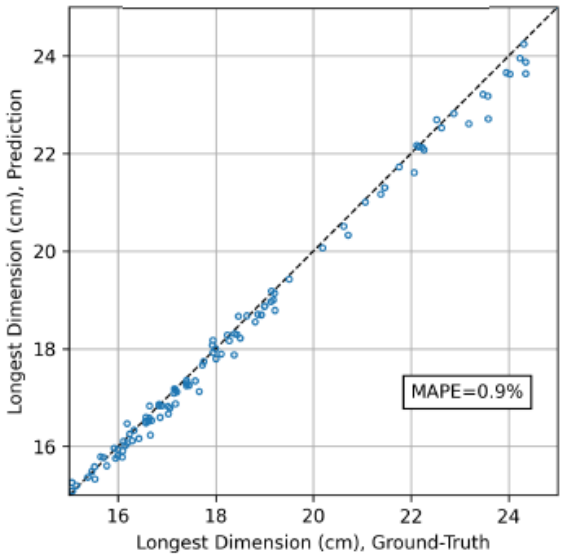

(d) Longest Dimension

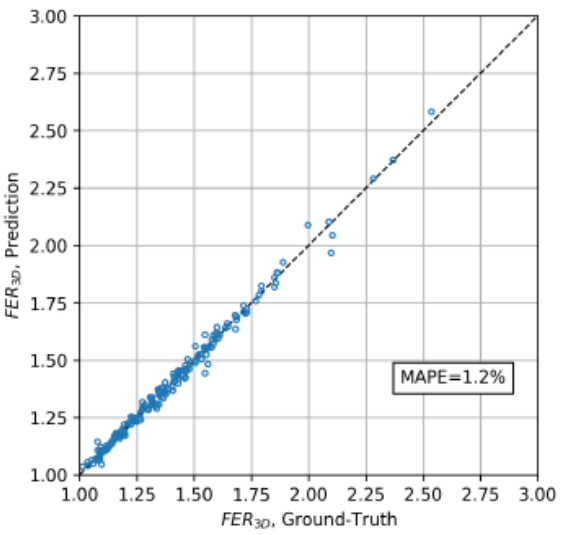

(e) 3D FER

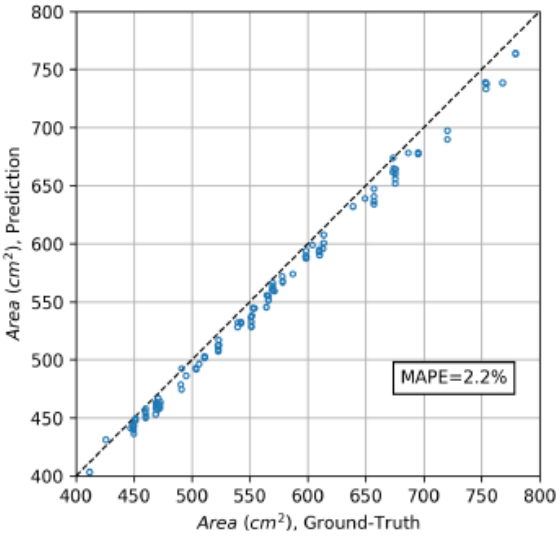

(f) Surface Area

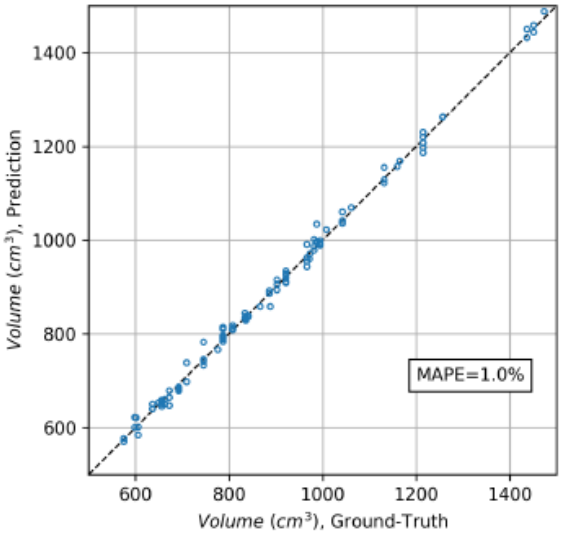

(g) Volume

Figure 26. Graph. Comparisons of macroscale metrics between the completed shapes and ground-truth shapes in the validation set.

Evaluation on Unseen Aggregate Shapes

The comparisons above demonstrate the good performance of the network in handling novel views of known shapes, but the network's ability to predict reasonable shapes for a completely unseen particle has not been verified. In this regard, the same type of comparison was made for the unseen shapes in the test set. First, the shape completion results of three randomly selected inputs from the test set are presented in Figure 27. By comparing with the shape completion performance on the validation set, two major observations were made between the validation set performance and the test set performance. First, the results from the test set show more uncertain prediction toward the

missing region of the shape. The test set results demonstrate a more scattered pattern among the predicted points that are near the missing region, meanwhile the validation set results generate sharper and more confident completion in the missing space. This is a positive sign, showing that given a completely unseen shape, the network is trying to predict the missing part in a probabilistic manner instead of force-fitting certain shape primitives.

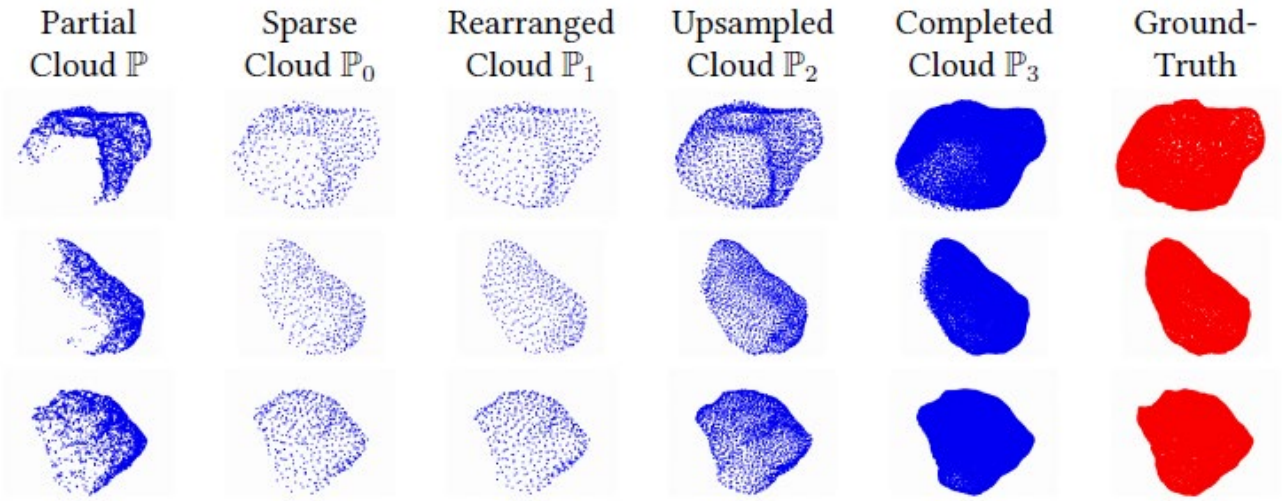

Figure 27. Illustration. Intermediate and final shape completion results for shapes in the test set.

In terms of macroscale metrics, the network is still able to predict shapes with reasonable matches regarding morphological properties, as shown in Figure 28. The MAPE errors of the results are consistently higher than the validation set results, which aligns with the fact that these are completely unseen shapes. The MAPE errors still lie within 5% but note that the MAPE error describes an error average rather than extremes. The maximum percentage error of the morphological properties could reach 10% or 20% for certain predictions. Predicting the unseen shape should always be a probabilistic approach, so the authors concluded that the network presents satisfactory performance in predicting reasonable shapes for unseen aggregates.

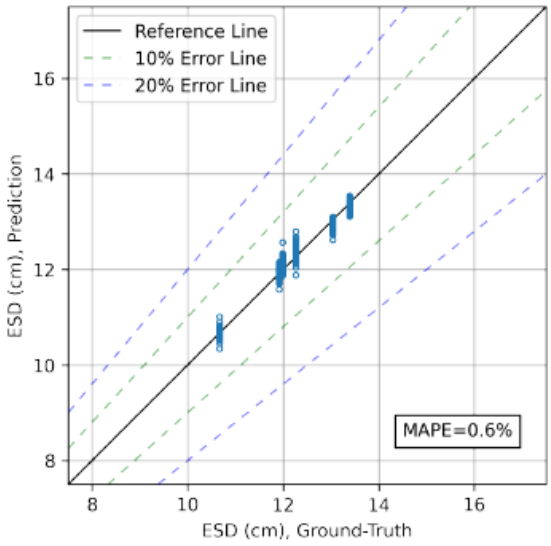

(a) ESD

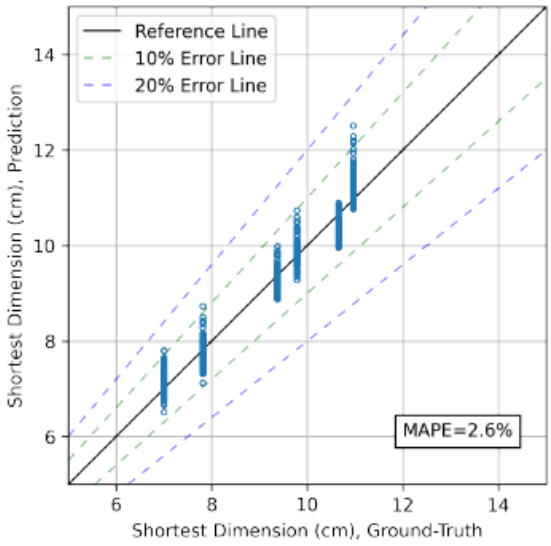

(b) Shortest Dimension

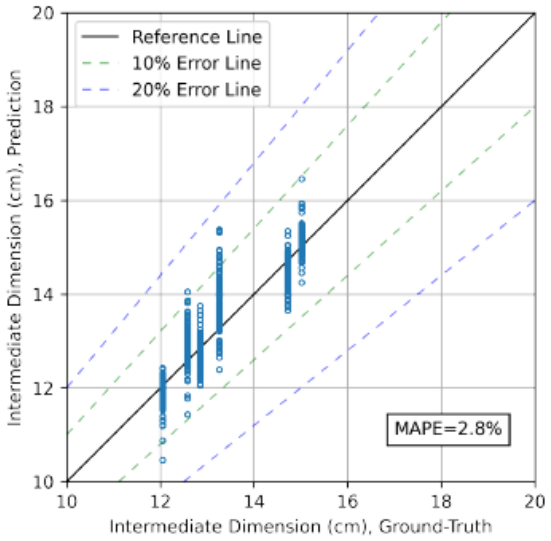

(c) Intermediate Dimension

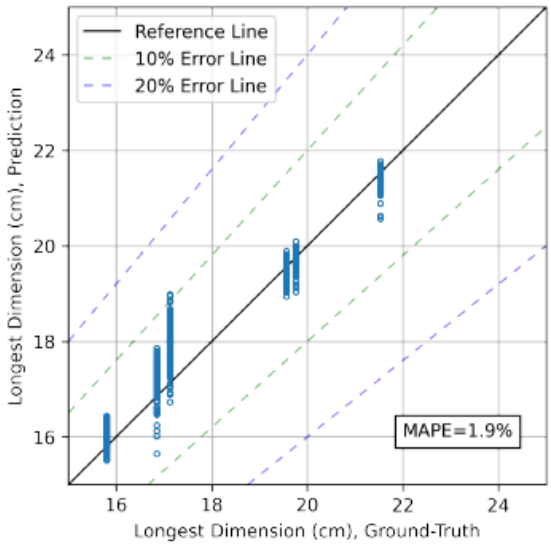

(d) Longest Dimension

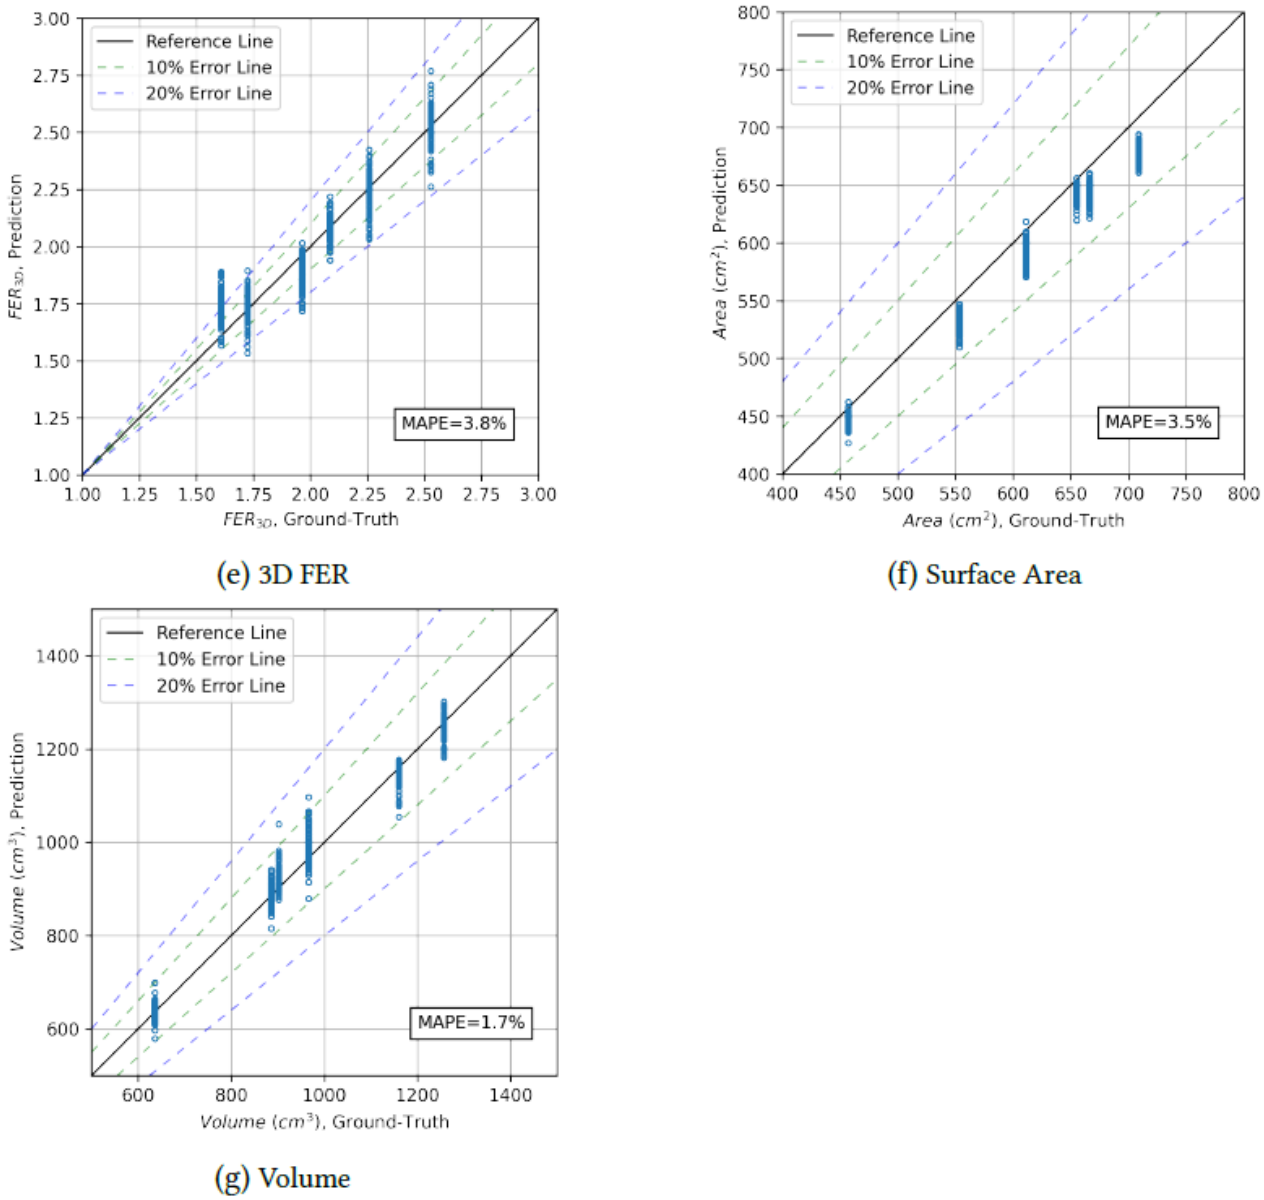

Figure 28. Graph. Comparisons of macroscale metrics between the completed shapes and ground-truth shapes in the test set.

IMPROVEMENT OF MORPHOLOGICAL ANALYSIS RESULTS USING SHAPE PERCENTAGE THRESHOLDING

By further investigating the aggregate shape characteristics in a stockpile, it was observed that the particle aggregate shapes usually exhibit very different visibility levels from a stockpile observation. For example, particles in a flat-layered stockpile setting tend to have better visibility due to larger void space and less occlusion between adjacent aggregates; meanwhile particles are likely to have lower visibility in a densely stacked stockpile setting.

In this regard, a quantitative method of characterizing the field “visibility” of aggregate shapes was developed. The intuition behind this visibility concept comes from the observation that partial and complete aggregate shapes demonstrate great differences in their spatial occupation patterns. A complete shape is a watertight surface such that a ray originating from the centroid will hit the enclosed surface in any arbitrary direction. However, for partial shapes, the rays originating from the centroid will either hit (for existing regions) or miss (for missing regions).

The field visibility of aggregate shapes is calculated using a modified ray-casting scheme similar to the one previously described. By calculating the ratios between the number of hit rays and the total number of cast rays, a visibility indicator named shape percentage (SP) was developed to quantify the partial shape observation, as described below:

- Step 1: Initializing a directional sphere at the centroid of the aggregate. Note that the centroid is approximated as the centroid of the partial shape, which may not be exactly at the centroid of the true shape; but is the best prediction based on partial observations. Then, a directional sphere with N = 1,000 equally distributed surface vertices is created at the centroid.

- Step 2: Ray casting for shape surface intersection. For each vertex on the sphere surface, the directional vector from the centroid to the vertex coordinates forms a ray direction. The ray-surface intersection is then conducted to indicate if the current direction contains a valid shape surface. If the ray hits the surface, the number of ray hits increments; otherwise, this direction represents a missing region.

- Step 3: Calculating the SP value. After completing ray casting for all N = 1, 000 directional vectors, the SP value is calculated as the ratio between the number of ray hits and the total number of rays.

The demonstration of the SP concept is presented in Figure 29. The blue region on the directional sphere represents the directions that have ray hits with the partial surface, while the orange region illustrates the missing space from the partial observation.

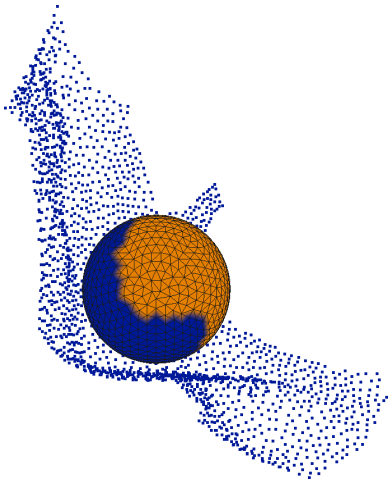

Figure 29. Illustration. Shape percentage concept of a partial shape.

Based on the SP indicator, the analysis results can be interpreted from a more effective perspective for all segmented shapes in the stockpile. An example analysis on RR3R-S1 stockpile is shown in Figure 30. The size of the datapoints is proportional to the SP value of each partial shape after segmentation. More intuitively, the SP values (in percentage) are directly labeled on each data point.

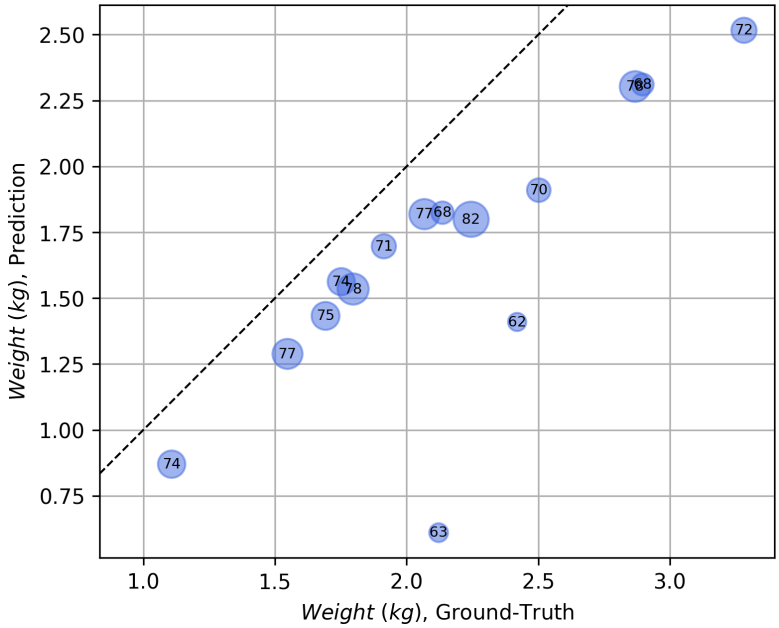

Figure 30. Graph. Shape percentage (in percentage) analysis of RR3R-S1 stockpile results.

As observed in Figure 30, most of the outliers with high deviation from the ground truth have relatively low SP values (e.g., below 70%). As previously discussed, the shape completion process is probabilistic and learning based. Therefore, the higher the shape visibility (i.e., the greater the portion of a shape that can be observed), the better the reliability and robustness of the shape completion results. When interpreting the results for practical use, it is important to screen and select the effective data, i.e., reliable results with high confidence level in the analysis.

Therefore, an SP thresholding process was developed to improve the morphological analyses procedure. Based on the segmented particle shapes from the laboratory and field stockpile analysis, the most common range of SP values is 60% to 85%. Accordingly, a SP threshold series of {65%, 70%, 75%, and 80%} was used to investigate the effect of SP thresholding. The morphological analysis results for field stockpiles at various SP levels are presented in Figure 31. For field stockpile data with only weight metric, the comparison between Figure 31(a) and Figure 31(d) shows the MAPE error decreases from 22.2% to around 15% (at SP levels of 75% and 80%).

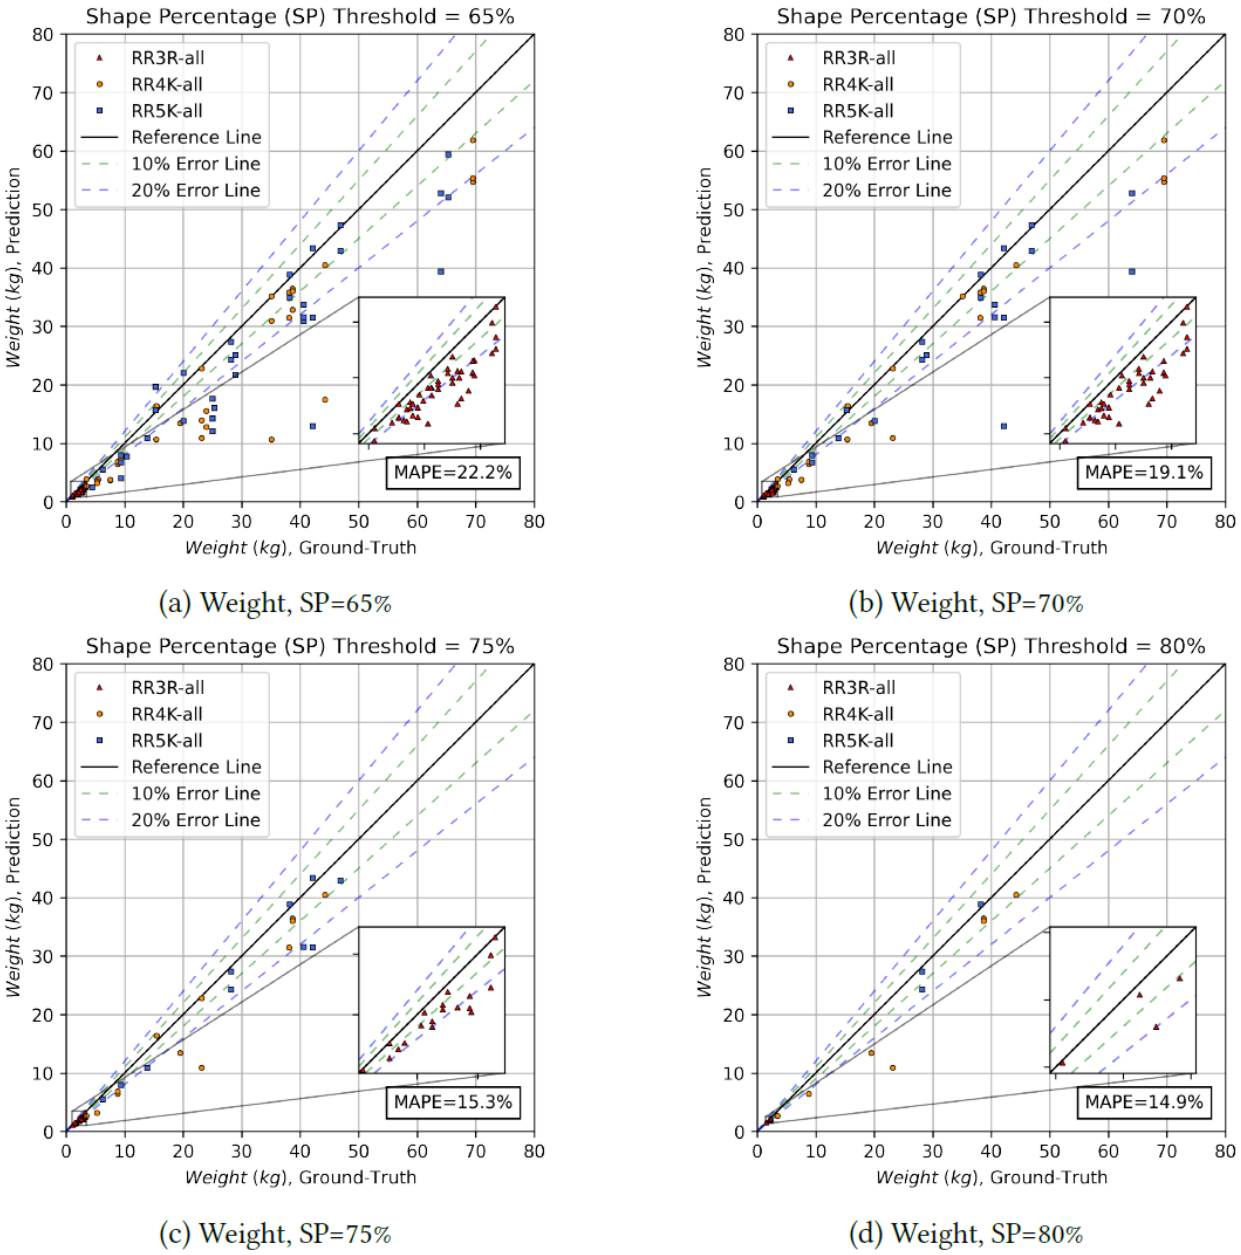

Figure 31. Graph. Effect of shape percentage thresholding on high-dimensional metric (weight) for field stockpile data.

The MAPE statistics indicate that the SP thresholding process effectively improves the results of the reconstruction-segmentation-completion framework in estimating morphological properties. In addition to the MAPE error evaluation (which only gives an average error estimate of all the data points), error bound analysis was also conducted to better reveal the improvements. As shown in Figure 31, ±10% and ±20% error lines were added to indicate the range of deviation with respect to ground truth. The SP thresholding process plays a crucial role in screening most of the less reliable predictions (i.e., those with low SP and limited partial shape observation) and improving the overall confidence level of the morphological analysis. This observation coincides perfectly with the intuition

of the shape percentage concept. Namely, when an aggregate shape is observed with limited visibility (e.g., SP below 60%), the shape completion result can only represent one of the best-effort guesses based on the partial observation. Conversely, as the partial shape approaches toward a relatively complete observation (e.g., SP over 75%), the shape completion results may better capture the true aggregate shape with increasing confidence. Based on Figure 31, a SP threshold of 75% is recommended for volume and weight estimation to ensure that a sufficient number of representative data points are kept with improved reliability.

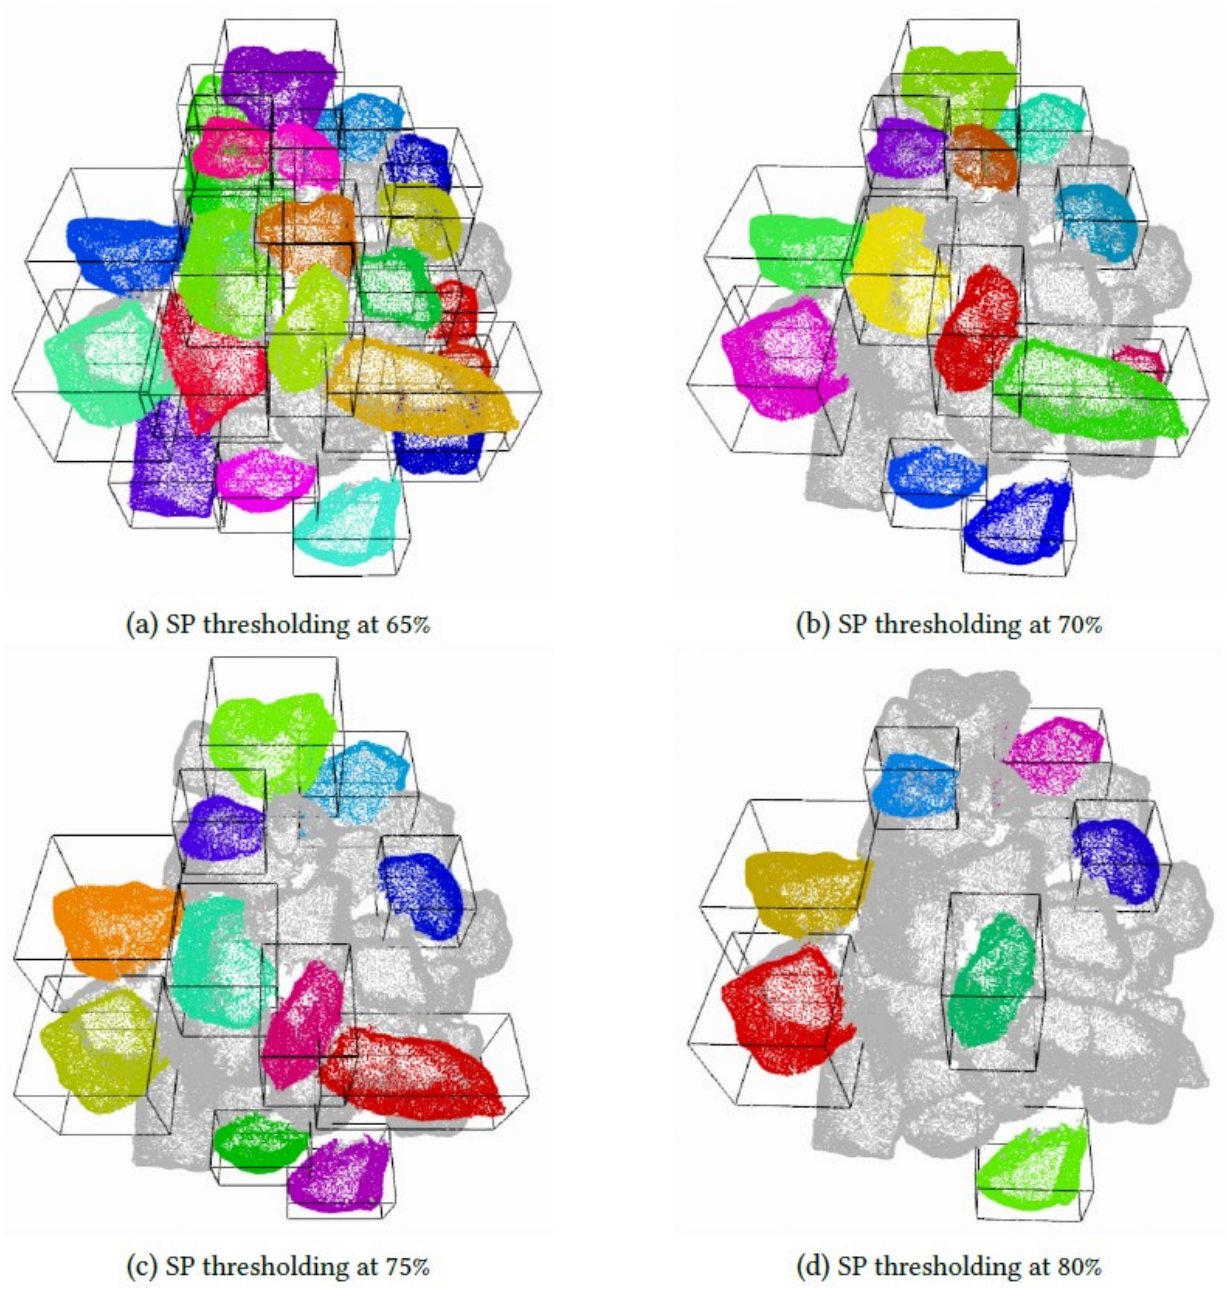

Figure 32. Illustration. Effect of SP thresholding on segmented aggregate shapes in a stockpile (RR4-S6).

Furthermore, Figure 32 illustrates the effect of the SP thresholding process on the stockpile analysis results. As the SP threshold increases, the number of effective aggregates decreases, with the remaining shapes at locations of less occupancy and larger open space. Aggregates at these protruding positions typically have better visibility with a large portion of the shape accessible from multi-view observation. Partial shapes segmented from a stockpile, in either densely stacked or flat-layered forms, can exhibit different SP values, i.e., visibility levels. Generally, the flat-layered stockpile form gives higher SP values or better shape visibility than the densely stacked form. This is because the aggregates in a flat-layered form usually have fewer occlusions from the stacking of particles. Therefore, when a certain SP threshold is used to screen the segmentation and completion results, flat-layered stockpiles are expected to have more effective aggregates (i.e., aggregates with SP greater than the threshold) than densely stacked stockpiles, given the same number of total aggregates in the stockpile.

SUMMARY

This chapter introduced the selected state-of-the-art strategy that is applicable to learning irregular aggregate shapes. To generate partial-complete shape pairs for deep learning, the varying-visibility and varying-view ray-casting schemes were developed, and an aggregate shape completion dataset was prepared. The selected deep learning framework was implemented and trained on the partial-complete shape pairs to learn the shape completion of aggregates. The shape completion network demonstrated good performance on both novel views of the known shapes as well as completely unseen shapes. Moreover, a shape percentage prediction threshold was developed toward practical interpretation of the morphological analysis results after the instance segmentation and shape completion steps. The final analysis results were quantitatively validated against the ground truth. Overall, the framework demonstrated robustness and reliability in characterizing morphological properties of aggregates in stockpiles.

CHAPTER 6: SOFTWARE DEVELOPMENT FOR 3D RIPRAP STOCKPILE CHARACTERIZATION

INTRODUCTION

This chapter presents the software development procedure and details for the 3D riprap stockpile image analysis program developed for IDOT and practitioners in Illinois. The primary outcome of this research study is a computer vision software application to conduct size and shape analyses of riprap and large-sized aggregates in the field. The deliverables integrated in the software include: (a) a user-friendly 3D reconstruction module for efficient field data preprocessing to generate calibrated 3D point clouds for aggregate stockpiles; (b) an automated segmentation module to extract individual aggregates from the stockpile; (c) a shape completion module to generate the complete 3D shape of the segmented particles; (d) a morphological analysis module to characterize particle size and shape; (e) a user-friendly graphical user interface (GUI) design for data input and output and visualization; and (f) a user manual including best practice and guidance for capturing images and demonstrating software usage.

REVIEW OF SOFTWARE ENGINEERING CONCEPTS

Software development is a complicated process that requires careful planning and execution to transform the goals into a software system. Developers must react timely and aggressively to meet ever-changing requirements (Rehman and Paul 2003). Maintaining software quality hinders fast-paced software development, as many testing cycles are necessary to ensure quality products (Pusuluri 2006).

Software Life Cycle

Software development, in contrast to common thought, is not just “writing code.” The life cycle of one software development project comprises many stages, which can be summarized as follows: 1) requirement gathering, 2) writing functional specifications, 3) creating architecture and design documents, 4) implementation and coding, 5) testing and assurance, 6) software release, and 7) documentation. All these stages/activities are necessary for a project to go through to produce a high-quality software system (Rehman and Paul 2003).

Following the predictable path of the software development life cycle, developers first understand and evaluate the problem or goal. Next, they identify the requirements of a solution to the problem. These requirements are studied in-depth to develop the architecture of the solution that will satisfy the requirements. The solution architecture provides the organizational details for designing the individual solutions (Eddie 2007). Once the architecture and detailed solution designs have been finalized, construction/implementation is the next focus. All work should be carefully tested before integration into the software system. Any further modification to the software system will be treated as maintenance, which will start a new iteration of the software life cycle (Eddie 2007).

Lifecycle Model of Software Development

A software process is a detailed series of steps during the software development project. To facilitate reasoning about software processes following the software life cycle, it is helpful to create specific abstract groupings of software processes, which is commonly known as the life cycle model (Feiler et al. 1993).

Because the software development for this project does not require extensive coding, the most-used model, code and fix, is selected as the desirable model. The code-and-fix model involves little to no overhead and is applicable for small projects and short-lived prototypes. The code-and-fix model starts with sufficient requirements and system specification analysis to begin coding. Once the software system works, developers rework and add more code to the system until it meets all requirements of the project.

DESIGN AND CONTROL LOGIC OF GRAPHICAL USER INTERFACE

The stockpile aggregate images captured by the end users need to be processed to yield a 3D point cloud representation of the stockpile that accurately replicates the physical stockpile scene. First, the collected images of a target riprap stockpile will be fed into the preprocessing module to conduct the 3D reconstruction step. During this step, the 3D point cloud of the stockpile is generated and cleaned after being calibrated to physical scale. The point cloud is then processed through the user-independent 3D segmentation and completion modules to analyze the complete 3D shape of each segmented particle. The output shapes are then used for determining the morphological indices of individual aggregate particles, such as principal dimensions, 3D FER, and volume and weight. The visual description of different modules is depicted in Figure 33, and the current version of the GUI is illustrated in Figure 34. The control logic of the GUI is divided into three modules: 3D reconstruction, 3D segmentation, and 3D completion.

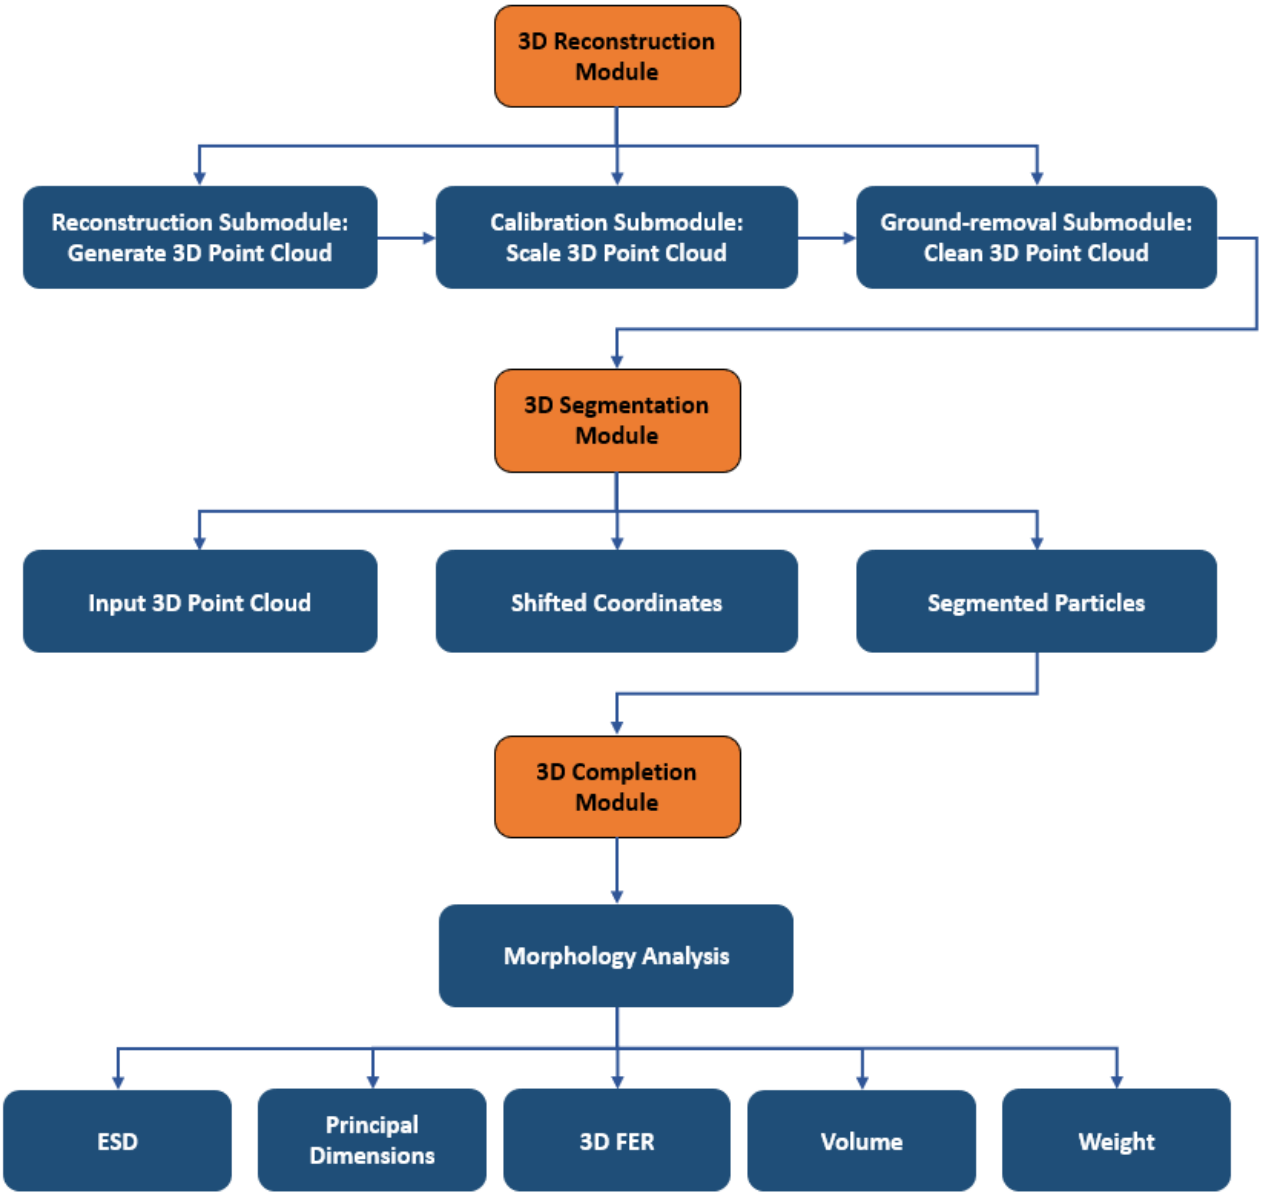

Figure 33. Illustration. Schematic drawing of the graphical user interface control logic.

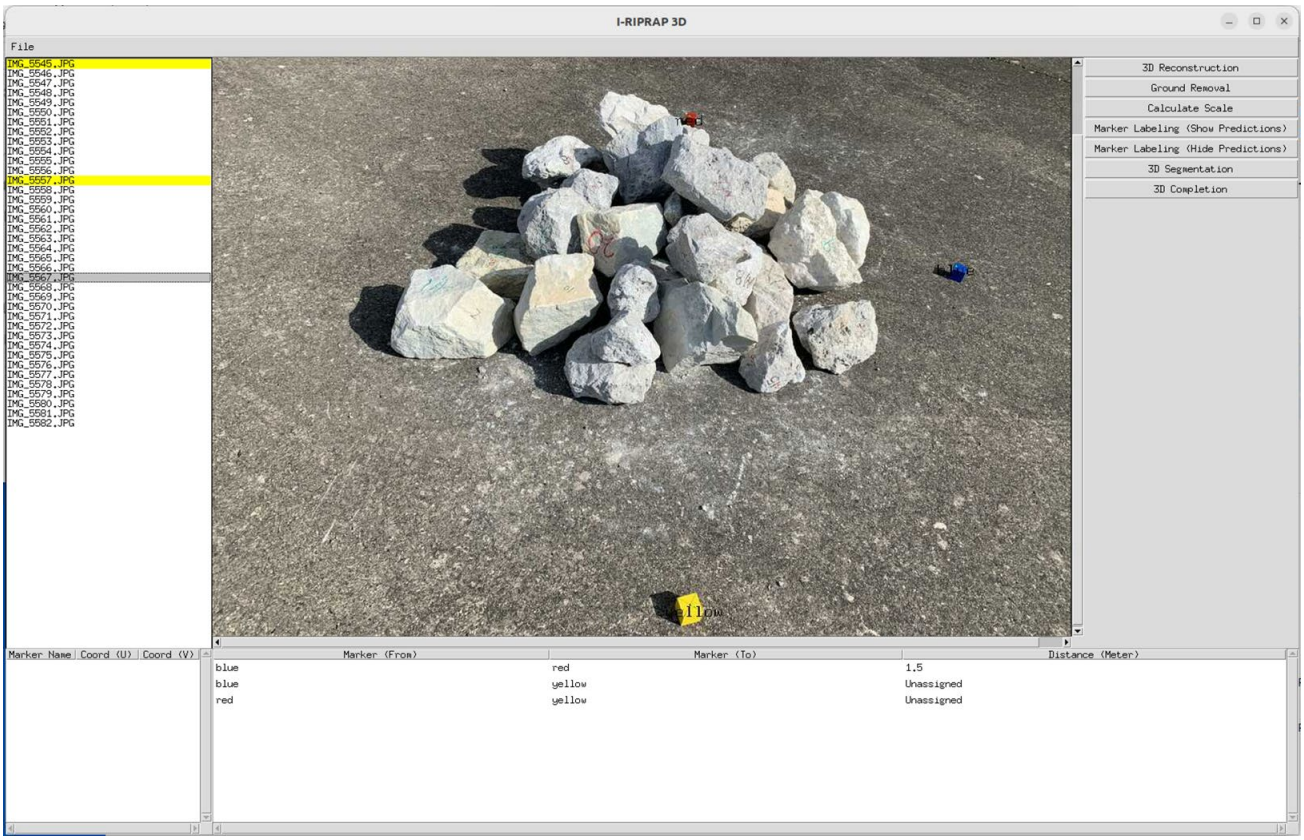

Figure 34. Illustration. Graphical user interface of the software.

3D Reconstruction Module

As a critical preprocessing step, the 3D reconstruction module was implemented to convert the input 2D image sequence (or video frames) to a 3D point cloud.

Initial-Reconstruction Submodule

To initiate 3D reconstruction, the user needs to click the “3D Reconstruction” button located at the control pane list (positioned at the upper right corner of the main window). To provide better user experience and separate different functions, a pop-up window will be prompted. After setting the required configuration, a preliminary version of the 3D point cloud will be computed as the output of this submodule.

Calibration Submodule

A major piece of missing information from the 2D image sequence (or video frame) is the scale factor of the point cloud. To obtain the physical scale, the user needs to manually annotate/label the location of the calibration objects on a subset of images and provide the corresponding measured distances between the calibration objects. During the labeling stage, after adding two images with the calibration objects, the software will compute the projected 3D location of them and back-project to other unlabeled images, which is termed “prediction.” The user can choose to show/hide the predictions to assess the accuracy of the calibration step. By labeling more images with the

calibration objects, the accuracy of the prediction increases. To achieve a relatively high accuracy, the authors recommend labeling three to five images from different viewing angles, at minimum. After checking the accuracy of the labeling, the software will compute the physical scaling and calibrate the raw point cloud.

Ground-removal Submodule

To clean up the non-stockpile points in the 3D point cloud, which are mainly the ground plane and the calibration objects, the user needs to manually remove those points. This step is necessary and important for a precise analysis due to the sensitivity of the module toward noise and outliers. In addition, due to the complicated and unpredictable field conditions, the surroundings of the target riprap stockpile may be noisy and random. This randomness necessitates that this step is designed to be manual instead of automated, so that the end user can ensure the right ground and background are being removed.

After completing the steps above, the output of the 3D reconstruction module is an accurate, clean, and calibrated 3D point cloud reconstructed from the input image sequence (or video frames).

3D Segmentation Module

After preprocessing the input images and specifying the output folder, the user can run the segmentation module with one click, and all results will be displayed in an interactive window presenting visual results at each step during 3D instance segmentation, as shown in Figure 35, where the calibrated 3D point cloud is the left image, the shifted coordinates of the point cloud are shown in the center image, and the segmented particles with bounding boxes are shown with different colors on the right image. A discussion of the algorithmic details was presented in Chapter 4.

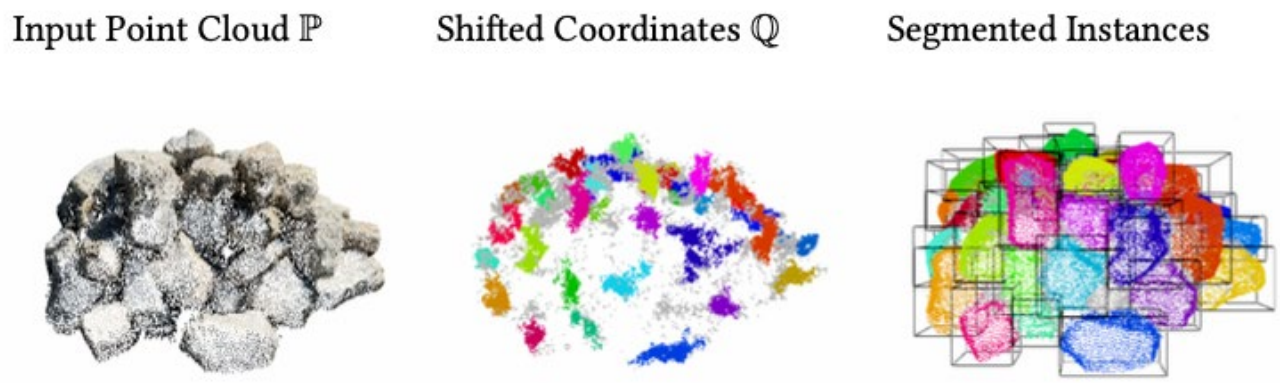

Figure 35. Illustration. Example of 3D segmentation result window.

3D Completion Module

After running the segmentation module, the user can run the 3D completion module in one click. Each individual segmented particle will be analyzed to predict the unseen portion and obtain the full shape. A detailed discussion on the algorithmic details was presented in Chapter 5. The SP

thresholding process will be applied at this step to filter out the highly occluded shapes, and the completed 3D model for each particle will be saved as the output data.

After the completion step, the software will automatically compute the morphological indices of the stockpile, such as ESD, principal dimensions, 3D FER, and volume and weight. A summary spreadsheet of the morphological analysis results will be saved in the output folder.

SUMMARY

This chapter summarized the software development effort toward a practical, user-independent, and user-friendly 3D riprap analysis software by integrating the proposed reconstruction-segmentation-completion approach. A software application was developed as the aggregate stockpile evaluation tool for size and shape analyses. The research team reviewed and followed professional software engineering guidelines during the developmental stages. 3D riprap stockpile analysis algorithms were programmed and integrated into a standalone software application with a simple user interface and interactive windows. The design of the GUI and control logic was intended to provide a fluent workflow for use by practitioners, including convenient file input/output (I/O), synchronous results display and switching, and user interaction. A user manual is also provided for software usage and technical support.

CHAPTER 7: CONCLUSIONS AND RECOMMENDATIONS

INTRODUCTION

The main benefit of this project includes the ability to conduct riprap and large-sized aggregate size and shape property testing in the field, at a quarry or construction site, with improved accuracy in determining the size and shape properties from imaging without having to weigh rocks. This is a pressing need because engineers and inspectors at quarries and government agencies have been spending considerable time and effort in riprap sizing at job sites.

Implementation of the advanced machine vision system developed in this project helps to accurately determine if a riprap material meets gradation and 3D-volume-based weight specifications. The proposed approach is sustainable for field applications by utilizing state-of-the-art machine vision and stockpile segmentation and shape completion techniques. At the end, better property characterization and optimized material selection can be achieved to improve designs through effective quality control, reduced costs, and energy consumption. Major cost savings in terms of personnel time, transportation, laboratory equipment, and facility use can be realized.

SUMMARY OF FINDINGS

Based on the 3D aggregate stockpile research study, the following findings can be summarized:

- A marker-based 3D reconstruction approach was developed as a cost-effective and flexible procedure to allow full 3D reconstruction of individual aggregates and aggregate stockpiles. A 3D aggregate particle library of 46 RR3 and 36 RR4 aggregate samples collected from field studies was established using this approach. The approach was further extended to field applications of reconstructing aggregate stockpiles, with an additional scale reference procedure to provide real physical scale of the reconstructed scene.

- A state-of-the-art 3D instance segmentation network was implemented, trained, and tested. The 3D instance segmentation network achieved an average completeness of 78% and an average intersection over union (IoU) precision of 82%, with a standard deviation of 6.3% and 4.8%, respectively. The segmentation network effectively learns the per-point offset vector to shift the original point cloud into an optimized clustered coordinate space, from which the instance proposals are generated.

- A state-of-the-art 3D shape completion network was implemented, trained, and tested. A dataset of partial-complete shape pairs was generated from the particle library based on varying-visibility and varying-view ray-casting schemes, at 7 visibility levels and 16 model orientations for each of the 82 models in the particle library. The completion network effectively learns the global and local shape context of the partial input point cloud and predicts the missing regions with fine-grained details.

- Based on the development of neural networks, an integrated 3D reconstruction-segmentation-completion approach was developed. The robustness and reliability of the approach were validated against re-engineered stockpiles and field stockpiles. Further, a shape percentage (SP) thresholding study was conducted to quantitatively characterize the partial observation process and improve the quality of morphological analysis.

- A practical, user-independent, and user-friendly 3D riprap analysis software was developed by integrating the established reconstruction-segmentation-completion approach. The software program is envisioned as a convenient tool that provides a fluent workflow for 3D aggregate stockpile analysis in industry and research applications.

RECOMMENDATIONS FOR FUTURE RESEARCH

Due to the data-driven nature of the deep learning networks in this framework, the performance is expected to gain progressive and scalable improvement with increased dataset size. The potential directions of improvements include:

- Extending the 3D particle library by collecting more aggregates from different origins and size groups (e.g., ballast, gravel, etc.). The size and quality of the 3D particle library will directly reflect on the quality of the stockpile dataset and the partial-complete shape dataset.

- Exploring and improving the 3D instance segmentation and 3D shape completion networks. By utilizing cutting-edge advancements in the computer vision research domain, the robustness and accuracy of the networks in the context of aggregate studies are expected to improve with innovative network designs.

Furthermore, during field visits and collaboration with aggregate producers, the researchers discovered great research potential for integrating the developed approach with intelligent sensing technologies. The current approach adopts the traditional SfM techniques for obtaining the 3D point clouds of aggregate stockpiles. Note that the three major components of the approach are standalone. Namely, as long as the input to the 3D segmentation and shape completion networks follows the point cloud format, the approach does not need to be bonded to certain 3D reconstruction methods or devices. With the rapid development in 3D visualization and augmented reality, it is expected that more advanced technologies for 3D sensing will be readily available in the future. For example, potential methods for the 3D reconstruction step can be further developed as follows: (i) LiDAR devices that directly capture the point cloud and (ii) dense simultaneous localization and mapping (SLAM) techniques that leverage RGB-D sensors and optical flow methods.

On the other hand, the data acquisition devices are not limited to handheld sensors. For example, to embed the developed approach into the aggregate production line, it is best to attach sensors to the conveyor system, which allows better statistical coverage of most aggregates before they are added to the stockpiles. Further, intelligent methods of acquiring stockpile aggregate images can be integrated with advanced aerial photography techniques. For example, unmanned aerial vehicles can greatly help with the image acquisition step for multi-spot or all-around inspection of a large stockpile, especially when intelligent route planning techniques are used. By dividing the large

stockpile into chunks, the entire stockpile can be analyzed by using the developed approach incrementally.

Also note the great potential of using unmanned aerial vehicles for calibration-free reconstruction. Commercial or industry level unmanned aerial vehicles usually have an open-source Software Developer ToolKit that records the internal inertial measurement unit data. With such flight route data integrated into the 3D reconstruction step, it is very likely to achieve a calibration-free reconstruction of aggregate stockpiles for segmentation as well as size and shape analyses of individual aggregate particles.

REFERENCES

Agisoft, LLC. 2021. *Agisoft Metashape User Manual: Professional Edition, Version 1.7*. Agisoft LLC.

Andrew, Alex M. 2001. "Multiple View Geometry in Computer Vision." *Kybernetes* 30 (9/10): 1333–1341. https://doi.org/10.1108/k.2001.30.9_10.1333.2.

Anochie-Boateng, Joseph K., Julius J. Komba, and Gculisile M. Mvelase. 2013. "Three-Dimensional Laser Scanning Technique to Quantify Aggregate and Ballast Shape Properties." *Construction and Building Materials* 43: 389–398. https://doi.org/10.1016/j.conbuildmat.2013.02.062.

Armeni, Iro, Ozan Sener, Amir R. Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 2016. "3D Semantic Parsing of Large-Scale Indoor Spaces." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*: 1534–1543.

ASTM D448. 2017. *Standard Classification for Size of Aggregate for Road and Bridge Construction*. West Conshohocken, PA: ASTM International.

ASTM D2940. 2015. *Standard Specification for Graded Aggregate Material for Bases or Subbases for Highway or Airports*. West Conshohocken, PA: ASTM International.

ASTM D5519. 2015. *Standard Test Methods for Particle Size Analysis of Natural and Man-Made Riprap Materials*. West Conshohocken, PA: ASTM International.

ASTM D6473. 2015. *Standard Test Method for Specific Gravity and Absorption of Rock for Erosion Control*. West Conshohocken, PA: ASTM International.

ASTM D6092. 2021. *Standard Practice for Specifying Standard Sizes of Stone for Erosion Control*. West Conshohocken, PA: ASTM International.

Barrett, P. J. 1980. "The Shape of Rock Particles, a Critical Review." *Sedimentology* 27 (3): 291–303. https://doi.org/10.1111/j.1365-3091.1980.tb01179.x.

Bartelt, D. 2018. *Wolman Count (Riprap Gradation Test)—Set-Up*. [Video]. YouTube. Accessed November 12, 2021. https://www.youtube.com/watch?v=aIiMb5UMHf4.

Bernhardsen, Tor. 2002. *Geographic Information Systems: An Introduction*. Hoboken, NJ: John Wiley & Sons.

Bessa, Iuri S., Verônica T. F. Castelo Branco, Jorge B. Soares, and José A. Nogueira Neto. 2015. "Aggregate Shape Properties and Their Influence on the Behavior of Hot-Mix Asphalt." *Journal of Materials in Civil Engineering* 27 (7). https://doi.org/10.1061/(ASCE)MT.1943-5533.0001181.

Blodgett, J. C., and C. E. McConaughy. 1986. *Water-Resources Investigations Report 86-4128: Rock Riprap Design for Protection of Stream Channels Near Highway Structures-Volume 2 Evaluation of Riprap Design Procedures*. Sacramento, CA: U.S. Geological Survey.

Chang, Angel X., Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. 2015. "ShapeNet: An Information-Rich 3D Model Repository." In: arXiv: 1512.03012 [cs]. https://doi.org/10.48550/arXiv.1512.03012.

Chen, Jingsong. 2011. “Discrete Element Method (DEM) Analyses for Hot-Mix Asphalt (HMA) Mixture Compaction.” PhD diss., University of Tennessee—Knoxville.

Chiew, Yee-Meng. 1995. “Mechanics of Riprap Failure at Bridge Piers.” *Journal of Hydraulic Engineering* 121 (9): 635–643. https://doi.org/10.1061/(ASCE)0733-9429(1995)121:9(635).

Choi, S., Q. Y. Zhou, and V. Koltun. 2015. “Robust Reconstruction of Indoor Scenes.” In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5556–5565.

Cignoni, P., M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia. 2008. “MeshLab: An Open-Source Mesh Processing Tool.” http://doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136

Eddie, B. 2007. *The Software Development Process*. http://sce2.umkc.edu/BIT/burrise/pl/software-process/. Accessed September 22, 2019.

Feiler, Peter H., and Watts S. Humphrey. 1993. “Software Process Development and Enactment: Concepts and Definitions*.” Proceedings of the Second International Conference on the Software Process-Continuous Software Process Improvement*, pp. 28–40.

Feret, L. R. 1930. La Grosseur Des Grains Des Matières Pulvérulentes [Grain Size Powdery Materials]. Eidgen. Materialprüfungsanstalt ad Eidgen. Technischen Hochschule.

Gates, L., E. Masad, R. Pyle, and D. Bushee. 2011. *Aggregate Imaging Measurement System 2 (AIMS2)*. Washington, DC: Federal Highway Administration.

Geiger, A., P. Lenz, C. Stiller, and R. Urtasun. 2013. “Vision Meets Robotics: The KITTI Dataset.” *International Journal of Robotics Research* 32 (11): 1231–1237.

Ghauch, Ziad. 2014. “Micromechanical Finite Element Modeling of Asphalt Concrete Materials Considering Moisture Presence.” Master’s thesis, University of Illinois Urbana-Champaign.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press.

Graham, Benjamin, Martin Engelcke, and Laurens van der Maaten 2018. “3D Semantic Segmentation with Submanifold Sparse Convolutional Networks.” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*: 9224–9232.

Griwodz, Carsten, Simone Gasparini, Lilian Calvet, Pierre Gurdjos, Fabien Castan, Benoit Maujean, Gregoire De Lillo, and Yann Lanthony. 2021. “AliceVision Meshroom: An Open-Source 3D Reconstruction Pipeline.” *Proceedings of the 12th ACM Multimedia Systems Conference*: 241–247. https://doi.org/10.1145/3458305.3478443

Guo, Y., H. Wang, Q. Hu, H. Liu, L. Liu, and M. Bennamoun. 2020. “Deep Learning for 3D Point Clouds: A Survey.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43(12): 4338–4364.

Han, Lei, Tian Zheng, Lan Xu, and Lu Fang. 2020. “OccuSeg: Occupancy-Aware 3D Instance Segmentation.” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*: 2940–2949. https://doi.org/10.48550/arXiv.2003.06537.

He, Y., H. Yu, X. Liu, Z. Yang, W. Sun, Y. Wang, ... and A. Mian. 2021. “Deep Learning Based 3D Segmentation: A Survey.” arXiv Preprint. https://doi.org/10.48550/arXiv.2103.05423.

Hiller, Priska Helene. 2017. “Riprap Design on the Downstream Slopes of Rockfill Dams.” PhD diss., Norwegian University of Science and Technology.

Hryciw, Roman D., Junxing Zheng, Hyon-Sohk Ohm, and Jia Li. 2014. “Innovations in Optical Geocharacterization.” *Geo-Congress 2014 Keynote Lectures: Geo-Characterization and Modeling for Sustainability*: 97–116. https://doi.org/10.1061/9780784413289.005

Huang, Hai. 2010. “Discrete Element Modeling of Railroad Ballast Using Imaging Based Aggregate Morphology Characterization.” PhD diss., University of Illinois Urbana-Champaign.

Huang, H., M. Moaveni, S. Schmidt, E. Tutumluer, and John M. Hart. 2018. “Evaluation of Railway Ballast Permeability Using Machine Vision–Based Degradation Analysis.” *Transportation Research Record* 2672 (10): 62–73. https://doi.org/10.1177%2F0361198118790849

Huang, H., J. Luo, M. Moaveni, E. Tutumluer, J. M. Hart, S. Beshears, and A. J. Stolba. 2019. “Field Imaging and Volumetric Reconstruction of Riprap Rock and Large-Sized Aggregates: Algorithms and Application.” *Transportation Research Record* 2673 (9): 575–589. https://doi.org/10.1177/0361198119848704.

Huang, H., J. Luo, E. Tutumluer, J. M. Hart, and I. Qamhia. 2020a. *Size and Shape Determination of Riprap and Large-sized Aggregates Using Field Imaging*. Rantoul, IL: Illinois Center for Transportation. https://doi.org/10.36501/0197-9191/20-003

Huang, H., J. Luo, E. Tutumluer, J. M. Hart, and A. J. Stolba. 2020b. “Automated Segmentation and Morphological Analyses of Stockpile Aggregate Images Using Deep Convolutional Neural Networks.” *Transportation Research Record* 2674 (10): 285–298. https://doi.org/10.1177/0361198120943887.

Huang, H., J. Luo, I. Qamhia, E. Tutumluer, J. M. Hart, and A. J. Stolba. 2021. “I-RIPRAP Computer Vision Software for Automated Size and Shape Characterization of Riprap in Stockpile Images.” *Transportation Research Record* 2675 (9): 238–250. https://doi.org/10.1177/03611981211001375.

Hubel, David H., and Torsten N. Wiesel. 1962. “Receptive Fields, Binocular Interaction and Functional Architecture in the Cat’s Visual Cortex.” *Journal of Physiology* 160 (1): 106–154. https://doi.org/10.1113%2Fjphysiol.1962.sp006837.

Illinois Department of Transportation. 2019. *Approved/Qualified Producer List of Aggregate Sources*. Springfield, IL: Illinois Department of Transportation, Central Bureau of Materials.

Illinois Department of Transportation. 2021. *Manual of Test Procedures for Materials*. Springfield, IL: Illinois Department of Transportation, Central Bureau of Materials.

Illinois Department of Transportation. 2022a. *Standard Specifications for Road and Bridge Construction*. Springfield, IL: Illinois Department of Transportation, Central Bureau of Materials.

Illinois Department of Transportation. 2022b. *Policy Memorandum 14-08.3*. Springfield, IL: Illinois Department of Transportation, Central Bureau of Materials.

Jiang, Li, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. 2020. “PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation.” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*: 4867–4876. https://doi.org/10.48550/arXiv.2004.01658.

Jin, Can, Xu Yang, Zhanping You, and Kai Liu. 2018. “Aggregate Shape Characterization Using Virtual Measurement of Three-Dimensional Solid Models Constructed from X-Ray CT Images of Aggregates.” *Journal of Materials in Civil Engineering* 30 (3): 04018026. https://doi.org/10.1061/(ASCE)MT.1943-5533.0002210

Kazhdan, Michael, and Hugues Hoppe. 2013. “Screened Poisson Surface Reconstruction.” *ACM Transactions on Graphics (ToG)* 32 (3): 1–13.

Kazmee, Hasan, and Erol Tutumluer. 2015. *Evaluation of Aggregate Subgrade Materials Used as Pavement Subgrade/Granular Subbase*. Rantoul, IL: Illinois Center for Transportation. https://apps.ict.illinois.edu/projects/getfile.asp?id=3517

Komba, Julius J., Joseph K. Anochie-Boateng, and Wyand van der Merwe Steyn. 2013. “Analytical and Laser Scanning Techniques to Determine Shape Properties of Aggregates.” *Transportation Research Record* 2335 (1): 60–71. https://doi.org/10.3141%2F2335-07

Lagasse, Peter Frederick, L. W. Zevenbergen, James Douglas Schall, and P. E. Clopper. 2001. *Bridge Scour and Stream Instability Countermeasures. Experience, Selection, and Design Guidance*. Washington, DC: Federal Highway Administration.

Lagasse, P. F., P. E. Clopper, L. W. Zevenbergen, and J. F. Ruff. 2006. *Riprap Design Criteria, Recommended Specifications, and Quality Control*. Washington, DC: National Cooperative Highway Research Program.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. “Deep Learning.” *Nature* 521 (7553): 436–444.

Lippert, D. L. (2012). *Inspection of Stone for Erosion Protection, Sediment Control, and Rockfill*. Springfield, IL: Illinois Department of Transportation.

Liu, Chen, and Yasutaka Furukawa. 2019. “MASC: Multi-Scale Affinity with Sparse Convolution for 3D Instance Segmentation.” In: arXiv:1902.04478 [cs]. https://doi.org/10.48550/arXiv.1902.04478.

Liu, Shih-Hung, Shang-Yi Yu, Shao-Chi Wu, Hwann-Tzong Chen, and Tyng-Luh Liu. 2020. “Learning Gaussian Instance Segmentation in Point Clouds.” In: arXiv:2007.09860 [cs]. https://doi.org/10.48550/arXiv.2007.09860.

Liu, Yufeng, Harikrishnan Nair, Stephen Lane, Linbing Wang, and Wenjuan Sun. 2019. *Influence of Aggregate Morphology and Grading on the Performance of 9.5-Mm Stone Matrix Asphalt Mixtures*. Charlottesville, VA: Virginia Transportation Research Council.

Longuet-Higgins, H. Christopher 1981. “A Computer Algorithm for Reconstructing a Scene from Two Projections.” *Nature* 293 (5828): 133–135.

Lutton, Richard J., Billy J. Houston, and James B. Warriner. 1981. *Evaluation of Quality and Performance of Stone as Riprap or Armor*. Vicksburg, MS: Army Engineer Waterways Experiment Station Geotechnical Lab.

Maerz, Norbert H., Tom C. Palangio, and John A. Franklin. 1996. “WipFrag Image Based Granulometry System.” In *Measurement of Blast Fragmentation*, edited by John A. Franklin and Takis Katsabanis, 91–99.

Maerz, Norbert H., and Tom C. Palangio. 1999. “WipFrag System II–Online Fragmentation Analysis.”

FRAGBLAST 6, Sixth International Symposium for Rock Fragmentation by Blasting, 8–12.

Marschner, Steve, and Peter Shirley. 2018. *Fundamentals of Computer Graphics*. Boca Raton, FL: CRC Press.

Masad, Eyad. 2003. *The Development of a Computer Controlled Image Analysis System for Measuring Aggregate Shape Properties*. Washington, DC: Transportation Research Board.

Masad, Eyad, Taleb Al-Rousan, Manjula Bathina, Jeremy McGahan, and Cliff Spiegelman. 2007. "Analysis of Aggregate Shape Characteristics and Its Relationship to Hot Mix Asphalt Performance." *Road Materials and Pavement Design* 8 (2): 317–350. https://doi.org/10.1080/14680629.2007.9690077

Minnesota Department of Transportation. 2018. *Standard Specifications for Construction*. Saint Paul, MN: Minnesota Department of Transportation.

Moaveni, Maziar, Shengnan Wang, John M. Hart, Erol Tutumluer, and Narendra Ahuja. 2013. "Evaluation of Aggregate Size and Shape by Means of Segmentation Techniques and Aggregate Image Processing Algorithms." *Transportation Research Record* 2335 (1): 50–59. https://doi.org/10.3141%2F2335-06

Nevada Department of Transportation. 2014. *Standard Specifications for Road and Bridge Construction*. Carson City, NV: Nevada Department of Transportation.

Obaidat, M. T., Ghuzlan, K. A., and M. M. Alawneh. 2017. "Analysis of Volumetric Properties of Bituminous Mixtures Using Cellular Phones and Image Processing Techniques." *Canadian Journal of Civil Engineering* 44 (9): 715–726.

O'Rourke, Joseph. 1985. "Finding Minimal Enclosing Boxes." *International Journal of Computer & Information Sciences* 14 (3): 183–199.

Ozturk, Hande Isik, and Isfandiyar Rashidzade. 2020. "A Photogrammetry Based Method for Determination of 3D Morphological Indices of Coarse Aggregates." *Construction and Building Materials* 262: 120794. https://doi.org/10.1016/j.conbuildmat.2020.120794.

Paixão, André, Ricardo Resende, and Eduardo Fortunato. 2018. "Photogrammetry for Digital Reconstruction of Railway Ballast Particles–a Cost-Efficient Method." *Construction and Building Materials* 191: 963–976. https://doi.org/10.1016/j.conbuildmat.2018.10.048.

Pan, Tongyan, Erol Tutumluer, and Joseph Anochie-Boateng. 2006. "Aggregate Morphology Affecting Resilient Behavior of Unbound Granular Materials." *Transportation Research Record* 1952 (1): 12–20.

Phiel, Shari. 2015. "Work on North Jetty Underway at Mouth of Columbia River." *Daily News*, June 6, 2015. https://tdn.com/news/local/work-on-north-jetty-underway-at-mouth-of-columbia-river/article_345224d3-dbbf-5d7d-9e38-9212a73c2d13.html

Polat, Rýza, Mehrzad Mohabbi Yadollahi, A. Emre Sagsoz, and Seracettin Arasan. 2013. "The Correlation between Aggregate Shape and Compressive Strength of Concrete: Digital Image Processing Approach." *International Journal of Structural and Civil Engineering Research* 2: 63–80.

Powers, David. 2011. "Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation." *Journal of Machine Learning Technologies* 2 (1): 37–63.

https://doi.org/10.48550/arXiv.2010.16061

Prince, Simon J. D. 2012. *Computer Vision: Models, Learning, and Inference*. Cambridge University Press.

Pusuluri, N. R. 2006. *Software Testing Concepts and Tools.* Dreamtech Press.

Qi, Charles R., Hao Su, Kaichun Mo, and Leonidas J. Guibas. 2017a. "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 652–660. https://doi.org/10.48550/arXiv.1612.00593

Qi, Charles R., Li Yi, Hao Su, and Leonidas J. Guibas. 2017b. "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space." In: arXiv:1706.02413 [cs]. https://doi.org/10.48550/arXiv.1706.02413

Qian, Yu. 2014. "Integrated Computational and Experimental Framework for the Assessment of Railroad Ballast Life-Cycle Behavior." PhD thesis, University of Illinois Urbana-Champaign.

Qin, Xuebin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar R. Zaiane, and Martin Jagersand. 2020. "U2-Net: Going Deeper with Nested U-Structure for Salient Object Detection." *Pattern Recognition* 106: 107404. https://doi.org/10.1016/j.patcog.2020.107404

Quiroga, P. N., and D. W. Fowler. 2004. ICAR Research Report No. 104-1F: The Effects of Aggregate Characteristics on the Performance of Portland Cement Concrete. International Center for Aggregates Research.

Rao, Chetana, Erol Tutumluer, and In Tai Kim. 2002. "Quantification of Coarse Aggregate Angularity Based on Image Analysis." *Transportation Research Record* 1787 (1): 117–124.

Rehman, R. U., and C. Paul. 2003. *The Linux Development Platform: Configuring, Using, and Maintaining a Complete Programming Environment*. Prentice Hall Professional.

Richardson, Everett V., and Stanley R. Davis. 2001. *Evaluating Scour at Bridges*. Washington, DC: Federal Highway Administration.

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. "U-Net: Convolutional Networks for Biomedical Image Segmentation." In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, edited by Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, 234–241. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-24574-4_28.

Al-Rousan, Taleb, Eyad Masad, Leslie Myers, and Cliff Speigelman. 2005. "New Methodology for Shape Classification of Aggregates." *Transportation Research Record* 1913 (1): 11–23.

Al-Rousan, Taleb, Eyad Masad, Erol Tutumluer, and Tongyan Pan. 2007. "Evaluation of Image Analysis Techniques for Quantifying Aggregate Shape Characteristics." *Construction and Building Materials* 21 (5): 978–990. https://doi.org/10.1016/j.conbuildmat.2006.03.005.

Senthilkumaran, N., and R. Rajesh. 2009. "Edge Detection Techniques for Image Segmentation—A Survey of Soft Computing Approaches." *International Journal of Recent Trends in Engineering* 1 (2): 250.

Sillick, S., and RAC AASHTO. 2017. *Member Survey Results*. Helena, MT: Montana Department of

Transportation. https://research.transportation.org/rac-survey-detail/?survey_id=371.

Tchapmi, L. P., V. Kosaraju, H. Rezatofighi, I. Reid, and S. Savarese. 2019. "Top-Net: Structural Point Cloud Decoder." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition,* 383–392.

Thilakarathna, Petikirige Sadeep Madhushan, Shanaka Kristombu Baduge, Priyan Mendis, Egodawaththa Ralalage Kanishka Chandrathilaka, Vanissorn Vimonsatit, and Hyuk Lee. 2021. "Aggregate Geometry Generation Method Using a Structured Light 3D Scanner, Spherical Harmonics–Based Geometry Reconstruction, and Placing Algorithms for Mesoscale Modeling of Concrete." *Journal of Materials in Civil Engineering* 33 (8): 04021198.

Tutumluer, E., N. Ahuja, J. M. Hart, M. Moaveni, H. Huang, Z. Zhao, and S. Shah. 2017. *Field Evaluation of Ballast Fouling Conditions Using Machine Vision. Report No. Safety-27*. Washington, DC: Transportation Research Board.

Tutumluer, Erol, and Tongyan Pan. 2008. "Aggregate Morphology Affecting Strength and Permanent Deformation Behavior of Unbound Aggregate Materials." *Journal of Materials in Civil Engineering* 20 (9): 617–627. https://doi.org/10.1061/(ASCE)0899-1561(2008)20:9(617)

Tutumluer, Erol, Chetana Rao, and Joseph A. Stefanski. 2000. *Video Image Analysis of Aggregates*. Urbana, IL: University of Illinois Urbana-Champaign/Illinois Department of Transportation.

U.S. Army Corps of Engineers. 1990. *Engineering and Design: Construction with Large Stone*. Engineer Manual No. 1110-2-2302. Washington, DC: U.S. Army Corps of Engineers.

Vincent, L., and P. Soille. 1991. "Watershed in Digital Spaces: An Efficient Algorithm Based on Immersion Simulation." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13 (6): 583–598.

Wadell, Hakon. 1932. "Volume, Shape, and Roundness of Rock Particles." *Journal of Geology* 40 (5): 443–451.

Wang, Linbing, Wenjuan Sun, Erol Tutumluer, and Cristian Druta. 2013. "Evaluation of Aggregate Imaging Techniques for Quantification of Morphological Characteristics." *Transportation Research Record* 2335 (1): 39–49. https://doi.org/10.3141%2F2335-05.

Wang, Weiyue, Ronald Yu, Qiangui Huang, and Ulrich Neumann. 2018. "SGPN: Similarity Group Proposal Network for 3D Point Cloud Instance Segmentation." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*: 2569–2578. https://doi.org/10.48550/arXiv.1711.08588

Wani, M. A., and B. G. Batchelor. 1994. "Edge-Region-Based Segmentation of Range Images." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 16: 314–319.

Wnek, Michael A., Erol Tutumluer, Maziar Moaveni, and Eric Gehringer. 2013. "Investigation of Aggregate Properties Influencing Railroad Ballast Performance." *Transportation Research Record* 2374 (1): 180–189. https://doi.org/10.3141%2F2374-21

Wolman, M. Gordon. 1954. "A Method of Sampling Coarse River-bed Material." *EOS, Transactions American Geophysical Union* 35 (6): 951–956. https://doi.org/10.1029/TR035i006p00951

Wu, Changchang. 2011. “VisualSFM: A Visual Structure from Motion System.” http://ccwu.me/vsfm/index.html

Xiang, Peng, Xin Wen, Yu-Shen Liu, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Zhizhong Han. 2021. “SnowflakeNet: Point Cloud Completion by Snowflake Point Deconvolution with Skip-Transformer.” *Proceedings of the IEEE/CVF International Conference on Computer Vision*: 5499–5509.

Yang, Bo, Jianan Wang, Ronald Clark, Qingyong Hu, Sen Wang, Andrew Markham, and Niki Trigoni. 2019. “Learning Object Bounding Boxes for 3D Instance Segmentation on Point Clouds.” *Advances in Neural Information Processing Systems*. https://doi.org/10.48550/arXiv.1906.01140.

Yi, Li, Wang Zhao, He Wang, Minhyuk Sung, and Leonidas J. Guibas. 2019. “GSPN: Generative Shape Proposal Network for 3D Instance Segmentation in Point Cloud.” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*: 3947–3956. https://doi.org/10.48550/arXiv.1812.03320

Zhao, Hengshuang, Li Jiang, Jiaya Jia, Philip H. S. Torr, and Vladlen Koltun. 2021. “Point Transformer.” *Proceedings of the IEEE/CVF International Conference on Computer Vision*: 16259–16268. https://doi.org/10.48550/arXiv.2012.09164

Zheng, Junxing, and Roman D. Hryciw. 2014. “Soil Particle Size Characterization by Stereophotography.” *Geo-Congress 2014: Geo-Characterization and Modeling for Sustainability*: 64–73. https://doi.org/10.1061/9780784413272.007

Zheng, J., and R. D. Hryciw. 2017. “Soil Particle Size and Shape Distributions by Stereophotography and Image Analysis.” *Geotechnical Testing Journal* 40 (2): 317–328.


Illinois Department
of Transportation


ILLINOIS