Evaluating the robustness of 3-D space-carving seed reconstruction under image noise

by

Pavan Sai Keshav Chandana

B.Tech., National Institute of Technology, Warangal, 2022

A REPORT

submitted in partial fulfillment of the
requirements for the degree

MASTER OF SCIENCE

Department of Computer Science
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Approved by:

Major Professor
Dr. Mitchell L. Neilsen

Copyright

Abstract

Three-dimensional (3-D) seed reconstruction is a critical component of automated plant phenotyping, enabling accurate measurement of seed volume and surface morphology from multi-view image data. Among various reconstruction approaches, space carving provides a simple yet powerful voxel-based method that integrates multiple binary silhouettes to estimate the true object geometry. However, the accuracy of such reconstructions depends strongly on the quality and consistency of the image masks used as input. In practical imaging pipelines, image errors arise unavoidably due to lighting variation, sensor artifacts, and imperfect thresholding or learning-based models.

This study investigates the robustness of 3-D space-carving reconstruction under controlled iamge perturbations. A single wheat seed was imaged from 36 viewpoints at 10° intervals, and reconstruction was performed using the first 10 silhouettes—sufficient for convergence of the carved volume—via the multi-threaded carving algorithm of Nielsen[1]. Binary masks were systematically modified using three perturbation types: morphological erosion, morphological dilation, and random salt-and-pepper noise. A 3×3 kernel was applied, and the percentage of pixels altered (0.2–1.0%) was varied to ensure consistent noise scaling across all perturbation modes.

The results show that space carving is highly stable under both erosion and dilation, with minor deviations even at the highest perturbation levels, confirming robustness to moderate under- and over-segmentation. In contrast, salt-and-pepper noise caused large deviations in surface area (up to 150%) and moderate reductions in volume (about 13%) due to the formation of small holes and fragmented silhouettes. These findings reveal that while the method is resilient to smooth boundary variation, it is highly sensitive to loss of silhouette connectivity.

Table of Contents

List of Figures

List of Tables

Acknowledgments

First, I would like to express my gratitude to my committee members Dr. Mitchell Neilsen, Dr. Torben Amtoft, and Dr. Lior Shamir for generously offering their time and support to serve on my committee throughout this project. I am particularly grateful to my Major Professor, Dr. Mitchell Neilsen, for his unreserved confidence in my abilities and consistent support since my second semester at Kansas State University. I also wish to thank Dr. Doina Caragea and Mr. John Keller for their expert guidance and support during my tenure as a Graduate Teaching Assistant.

I am forever indebted to my family members, in particular my mother, Udayini Chandana, and my father, Satish Chandana, for their steadfast support and encouragement that made this journey possible. Their guidance has been invaluable in my life and career thus far.

I would also like to express my appreciation to my friends Sai Surya Ravipati, Nirmal Gelal, Samhitha Burugupalli, Aryan Dalal and Bindu Naraharasetti for their constant support and companionship that has made the past two years extremely memorable and enjoyable.

Chapter 1

Introduction

1.1 Problem Definition

In automated seed phenotyping, three-dimensional measurements such as volume and surface geometry are important for analyzing seed quality, varietal differences, and developmental characteristics. The multi-view space carving method provides an effective way to reconstruct the 3-D shape of a seed from multiple 2-D images captured around a turntable. However, the accuracy of the reconstruction depends directly on the quality of the masks extracted from each image.

In practice, segmentation is often performed automatically using image processing techniques or machine learning models. These methods are prone to different kinds of errors, such as under-segmentation, over-segmentation, and random noise, caused by lighting variations, threshold sensitivity, and image artifacts. These errors can propagate through the reconstruction process and lead to inaccurate estimates of seed volume and surface area, which can reduce the reliability of downstream phenotypic analysis.

Therefore, it is necessary to understand how sensitive the 3-D space carving algorithm is to mask inaccuracies, and to what extent reconstruction measurements remain stable under different types of noise.

1.2 Objective

The primary objective of this project is to evaluate the reliability of a 3-D seed reconstruction method under variations in mask quality. The study focuses on the multi-view space carving framework, which reconstructs a seed's volume from a set of silhouette images captured at evenly spaced rotational angles. Since the accuracy of this reconstruction depends heavily on the precision of the masks derived from these images, the project aims to analyze how different forms of image noise affect the resulting 3-D model.

To achieve this, a baseline set of segmented images is first used to perform reconstruction and compute reference measurements of volume and surface area. The masks are then systematically modified using morphological erosion and dilation to simulate under- and over-segmentation, as well as salt-and-pepper noise to introduce random pixel-level distortion. Each variant of the mask set is processed through the same space-carving pipeline, and the resulting reconstructed models are evaluated and compared against the baseline.

This study provides insight into how robust the space carving algorithm is when mask quality varies and identifies the magnitude and type of errors that significantly impact downstream phenotypic measurements. The findings contribute to improving the reliability of automated seed phenotyping systems in practical environments where image imperfections are unavoidable.

Chapter 2

Background and Related Work

2.1 Seed Phenotyping and the Importance of 3-D Measurements

Seed morphology, including traits such as volume, shape, and surface structure, plays a critical role in the evaluation of plant performance, genetic variation, and market quality. Traditional seed phenotyping methods often rely on manual measurement tools or 2-D imaging systems. Although widely used, these approaches suffer from limitations: manual measurements are labor intensive and subjective, while 2-D imaging only captures projected silhouettes and lacks depth information, resulting in incomplete or distorted shape characterization.

To address these limitations, three-dimensional (3D) phenotyping methods have gained increasing attention. 3D reconstruction enables precise quantification of seed structure, improving the precision of morphological comparison and supporting large-scale breeding, quality assessment, and trait analysis. Several computational approaches, including volumetric modeling, shape-from-shading, ellipsoid fitting, and voxel carving, have been explored to produce reliable 3-D representations of biological samples.

2.2 Early Approaches to Volume Estimation

Early work in agricultural product phenotyping primarily used mathematical approximation methods to estimate volume. For instance, Koc[2] evaluated the volume of a watermelon using ellipsoid slicing and image-derived contour extraction. Geometric parameters such as length and diameter were substituted into ellipsoid-based formulas, while the contour-based image processing approach estimated volume using numerical integration. Comparisons showed that the image-derived method more closely matched physical water displacement measurements, demonstrating the value of digital image analysis in shape estimation.

Similarly, Sankaran[3] estimated dimensions of a chickpea using watershed-based segmentation to isolate seeds in clustered images. Their measurements showed a strong correlation to ground-truth physical measurements, further validating digital imaging approaches for agricultural phenotyping.

These studies demonstrate the feasibility of image-based geometric measurement, but they remain fundamentally two-dimensional and cannot capture the full spatial characteristics of seeds with irregular or non-ellipsoidal shapes.

2.3 Shape-from-Silhouette and Multi-View Reconstruction

The shape-from-silhouette (SfS) or space carving approach has been one of the most influential techniques for reconstructing volumetric structure from images. In SfS, multiple 2-D silhouettes of an object captured from different viewpoints are used to iteratively carve away regions of a 3-D voxel grid that cannot belong to the object surface. The intersected volume of all back-projected silhouettes yields an approximation of the object known as the visual hull.

The foundations of silhouette carving were formalized by Potmesil[4], who proposed constructing octree-based volumetric representations from sequences of object silhouettes. This work demonstrated that silhouette intersections could approximate arbitrary 3-D shapes using only 2-D boundary information, avoiding the need for explicit surface depth cues.

Roussel[5] applied this method to reconstruct seed shapes as small as 200 µm. Although successful, their system processed only one seed at a time and required carefully controlled imaging conditions and camera positioning, resulting in low throughput.

Later, Cao[6] introduced a more accessible single-seed reconstruction setup using 3-D printed supports and a standard digital camera. Seeds were placed on a rotating platform against a uniform background while images were captured at fixed angular increments. The reconstruction relied on silhouette-based space carving, and validation experiments using spherical ceramic reference objects demonstrated volume estimation within 3% error, confirming both practicality and accuracy.

Zhao[7] extended this system by improving the imaging platform and segmentation procedure, enabling more consistent silhouette extraction under varied lighting conditions. This resulted in an affordable, repeatable, and accurate single-seed reconstruction pipeline, which formed the basis of the system used in the present work.

2.4 Space Carving and Its Computational Refinements

The space carving algorithm reconstructs a 3-D object by iteratively removing voxels that are inconsistent with silhouette constraints across all viewpoints[4]. When applied to seed imaging, each silhouette defines a region in 3-D space where the seed must exist. By carving away voxels that fall outside any silhouette, the true 3-D shape emerges.

However, traditional implementations can be computationally expensive, particularly when processing large voxel grids. Nielsen[1] introduced a multi-threaded C-based carving implementation designed to accelerate the reconstruction process while preserving accuracy. Their implementation partitions the voxel grid and processes it in parallel, allowing faster reconstruction without requiring specialized hardware such as GPUs. This made 3-D seed reconstruction both efficient and accessible for standard laboratory environments.

2.5 Segmentation in Seed Phenotyping

Accurate silhouette extraction is critical for the success of space carving. Segmentation can be performed using:

1. Traditional digital image processing (DIP) — e.g., HSV color thresholding, morphological filtering
2. Deep learning-based segmentation — e.g., convolutional neural networks trained for seed contour detection

Zhao[?] investigated traditional image processing and convolutional neural network (CNN) methods for seed edge detection. Their experiments demonstrated that CNN-based segmentation models, particularly U-Net and Mask R-CNN variants, were significantly more robust under lighting variations and complex backgrounds, though at the cost of greater computational overhead and the need for labeled training datasets. Despite these advances, HSV segmentation continues to be used in practical laboratory environments due to its ease of deployment and low computational requirements, especially in low-volume phenotyping workflows.

However, both approaches are subject to segmentation errors caused by:

- Shadows and lighting variation,
- Partial occlusion or seed reflectance,
- Sensitivity to threshold selection,

- Image noise introduced by camera sensors.

These segmentation inaccuracies directly propagate into the space-carving stage because the voxel occupancy decision at each viewing angle depends solely on the correctness of the silhouette mask.

2.6 Motivation for the Present Study

Although prior work has successfully demonstrated 3-D seed reconstruction systems and explored alternative segmentation strategies, there has been limited investigation into how segmentation errors affect the accuracy and stability of the reconstruction output. In real-world phenotyping pipelines, segmentation is rarely perfect, especially when automated at scale. Thus, it is important to understand:

- How much segmentation error the system can tolerate,
- Which types of segmentation errors are most impactful,
- How reconstructed volume and surface area respond to specific noise conditions.

This research addresses this gap by conducting a systematic robustness study. Controlled perturbations (morphological erosion, dilation, and pixel noise) are applied to the segmentation masks, and their effects on reconstructed volume and surface geometry are quantitatively evaluated.

Chapter 3

Methodology

This chapter describes the dataset, imaging setup, preprocessing pipeline, mask creation procedure, 3-D reconstruction method, noise perturbation strategy, and evaluation metrics used in this study. The workflow is designed to replicate a standard single seed phenotyping system while allowing controlled analysis of how image quality influences the reconstructed 3-D morphology.

3.1 System Overview

The core objective of this work is to analyze the robustness of 3-D seed reconstruction when masks are degraded systematically. To enable this, the study utilizes a controlled single seed imaging setup in which a wheat seed is placed on a motorized turntable and captured from multiple angular viewpoints. The acquired images are segmented to extract silhouettes, which are then used as input to a voxel-based space-carving algorithm. The volume and surface area of the reconstructed seed model are computed and used as quantitative measures of reconstruction accuracy under different types of image noise.

The experimental workflow consists of five primary stages:

1. Image Acquisition

2. Image Preprocessing and Coregistration

3. Creation of Silhouette Masks

4. 3-D Reconstruction Using Multi-Threaded Space Carving

5. Noise-Based Perturbation and Performance Evaluation

Figure 3.1: System Overview

A summary of the full workflow is shown in Fig.3.1

3.2 Image Acquisition Setup

3.2.1 Seed Type

A single wheat seed was selected as the subject for reconstruction. Wheat seeds are small, non-spherical, and relatively anisotropic in shape, making them well suited for evaluating sensitivity to image noise.

3.2.2 Imaging Platform

The imaging setup consists of:

- A standard digital camera mounted in a fixed side-view position
- A motorized rotation stage (turntable)
- A uniformly colored background to improve silhouette contrast

- Controlled laboratory lighting to minimize shadow artifacts

Figure 3.2: Imaging platform setup as seen in Neilsen et al.[1]

The seed is placed upright on the turntable and rotated at fixed angular increments. The camera remains stationary, ensuring that all variation in the images arises solely from the change in rotation angle.

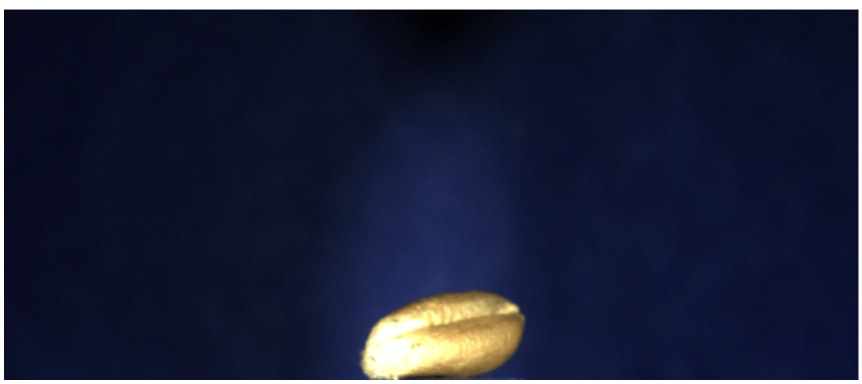

Figure 3.3: Image of a wheat seed taken in the imaging setup

3.2.3 Number of Views

A total of 36 images were captured, corresponding to a 10° rotational step between consecutive images, covering a full 360° rotation:

$$\theta = \frac{360^\circ}{36} = 10^\circ$$

The number of images is in line with prior studies[1].

3.3 Image Preprocessing and Coregistration

Raw images are captured at a resolution of 640×277 pixels. To ensure the silhouette extraction and space-carving process operate correctly, the images are aligned and cropped around the seed using the coregister and crop procedure.

3.3.1 Cropping

A rectangular region containing the seed is manually identified in the first frame. This crop rectangle is applied across all images to remove background elements and reduce processing time. The resulting cropped images are standardized to 200×200 pixels.

3.3.2 Coregistration

Although the seed remains centered on the rotating platform, small shifts in seed position may occur due to camera alignment, vibration, or rotation axis variation. The coregistration step computes offsets that adjust each frame to keep the seed centered across all views. These offsets are later used when projecting the volume during reconstruction.

3.4 Mask Creation Using HSV Thresholding

Silhouette masks are generated using a color-space thresholding approach in HSV (Hue, Saturation, Value) space. Wheat seeds exhibit distinct reflectance and color characteristics relative to the background; therefore, HSV thresholding is chosen for its robustness to moderate brightness and contrast changes.

The mask creation procedure is:

1. Convert the cropped RGB image to HSV space.

2. Apply lower and upper HSV threshold bounds to isolate the seed.

3. Convert the threshold result to a binary mask.

4. Apply morphological closing to remove small holes inside the silhouette.

5. Save binary masks in white-foreground / black-background format, where:

These masks form the baseline input for reconstruction.

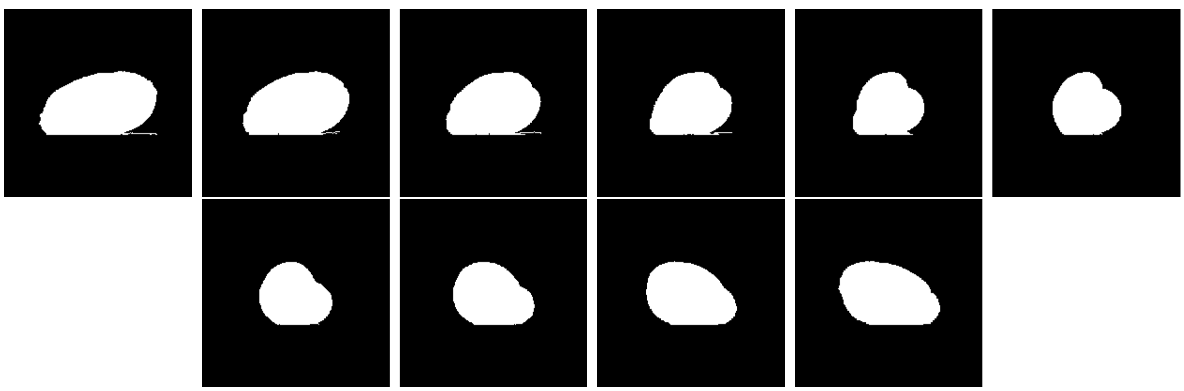

Figure 3.4: Example masks used in the 3D reconstruction process. Masks are shown at evenly spaced rotation angles to illustrate the change in silhouette as the seed rotates on the turntable.

3.5 3-D Reconstruction via Multi-Threaded Space Carving

The reconstructed 3D seed volume is represented as a voxel grid of size $125 \times 100 \times 125$. These dimensions were chosen to balance spatial resolution with computation time.

3.5.1 Initial Volume Construction

A total of 36 images were acquired while the seed was rotated in 10° increments on a calibrated turntable. However, following the convergence analysis reported in Figure 5 of Nielsen[1], only the first 10 silhouettes were used during reconstruction. The prior study demonstrated that once 8–10 viewpoints have been incorporated, the carved volume changes negligibly with additional views due to the redundancy of overlapping silhouette constraints. This allows a substantial reduction in computation time with no meaningful loss of geometric accuracy.

Figure 3.5: Space carving with masks 0, 1, 2, 3, 4, 5, 6, and 7/8. Figure adapted from Neilsen et al.[1]

Accordingly, in this work, the voxel carving process begins with a volumetric grid initialized to the bounding cuboid dimensions (12.5 mm $\times$ 10.0 mm $\times$ 12.5 mm). Each of the first 10 binary masks is then used to remove voxels inconsistent with the corresponding silhouette projection. After approximately the eighth mask, the volume converges to a stable shape, with further carving iterations producing changes of less than 0.5% in the measured volume or surface area. Thus, the reconstruction step efficiently exploits viewpoint redundancy while maintaining accuracy.

3.5.2 Projection Matrices

For each image, a projection matrix maps voxel coordinates to image coordinates, accounting for camera parameters, crop offsets, and rotation angle.

3.5.3 Carving Operation

For each view:

- Voxels that project onto background pixels in the silhouette mask are removed.
- Voxels that remain consistent across all silhouettes form the final seed model.

The voxel removal step is performed using a compiled C implementation, which employs multi-threading to accelerate carving across slices of the volume.

3.6 Noise Perturbation Strategy

To evaluate the robustness of the reconstruction process to image inaccuracies, controlled distortions were applied to the baseline binary mask sequence. Three perturbation families were introduced, corresponding to under-segmentation, over-segmentation, and random pixel-level noise.

Noise Type	**Purpose**	**Implementation Details**
Erosion (0.2–1.0%)	Simulates under-segmentation (foreground shrinkage)	Randomly erodes 0.2–1.0% of foreground pixels using a 3×3 kernel; 10 random seeds
Dilation (0.2–1.0%)	Simulates over-segmentation (foreground expansion)	Randomly dilates 0.2–1.0% of foreground pixels using a 3×3 kernel; 10 random seeds
Salt-and-Pepper (0.2–1.0%)	Simulates unstructured pixel-level noise	Randomly flips 0.2–1.0% of total pixels; 30 random seeds

Table 3.1: Summary of the image noise perturbation strategy.

Morphological erosion and dilation operate using a structuring element, or kernel, which defines the neighborhood used to modify the binary mask. This kernel is a small matrix that slides over the mask and its shape and size determine how the operation affects the mask. In erosion, for each pixel in the mask, if all pixels under the kernel are ”on” (white in this case), the central pixel remains ”on”. Otherwise, it becomes ”off” (black). This shrinks foreground objects and can remove small noise. Meanwhile in dilation, for each pixel in the mask, if at least one pixel under the kernel is ”on”, the central pixel becomes ”on”. This expands foreground objects and can fill small holes or connect broken parts.

For erosion and dilation, the i value (0.2–1.0%) controls the proportion of pixels randomly eroded or dilated, and j (1–10) seeds the random number generator to introduce slight variations across trials. For salt-and-pepper noise, 30 independently seeded trials were generated for each percentage level.

Each perturbed mask set was reconstructed using the same shape-from-silhouette pipeline as the baseline, allowing direct comparison of volume and surface area measurements against the original unperturbed reconstruction.

3.7 Volume and Surface Area Computation

3.7.1 Volume

Volume is determined by counting occupied voxels and multiplying by the physical size of each voxel. Because voxel spacing is isotropic, volume is computed directly as:

$$V = N_{\text{voxels}} \cdot s^3$$

where s is the voxel side length (mm).

3.7.2 Surface Area Estimation via Marching Cubes

To obtain an accurate estimate of the seed surface geometry, the reconstructed volumetric data was converted into a triangular mesh representation using the Marching Cubes algorithm. Let $V(x, y, z)$ denote the binary volumetric occupancy grid obtained after space carving, where $V(x, y, z) = 1$ indicates seed volume and $V(x, y, z) = 0$ indicates empty space. The marching cubes operation extracts a polygonal iso-surface at the object boundary by interpolating voxel intensity transitions.

$$(\mathbf{v}, \mathbf{f}) = \mathrm{MarchingCubes}(V, \mathrm{level} = 0.5, \mathrm{spacing} = (s_x, s_y, s_z))$$

Here, $\mathbf{v}$ is the set of mesh vertices expressed in millimeter units, $\mathbf{f}$ is the set of triangular face indices, and (s_x, s_y, s_z) represent the voxel dimensions in mm:

$$s_x = \frac{W}{N_x}, \quad s_y = \frac{H}{N_y}, \quad s_z = \frac{D}{N_z}$$

where W, H and D are the physical dimensions of the reconstruction volume, and N_x, N_y, N_z are the number of voxels per axis.

For each triangular face, the local surface area is computed as:

$$A_i = \frac{1}{2} \left\| (\mathbf{v}_{i,2} - \mathbf{v}_{i,1}) \times (\mathbf{v}_{i,3} - \mathbf{v}_{i,1}) \right\|$$

The total seed surface area A is then obtained by summing over all N_f triangular faces:

$$A = \sum_{i=1}^{N_f} A_i$$

This mesh-based approach provides a more accurate surface representation than voxel-boundary counting and reduces discretization-induced overestimation of surface area.

3.8 Evaluation Protocol

For each mask condition, the reconstructed volume and surface area are compared to baseline values. Percent error is computed as:

$$\text{Error} = \frac{|M - M_{\text{baseline}}|}{M_{\text{baseline}}} \times 100\%$$

where M is volume or surface area under noise.

Chapter 4

Results and Analysis

4.1 Overview

This chapter presents the experimental results obtained from the 3D reconstruction and robustness evaluation pipeline described in the methodology. The primary objective of the analysis is to assess the sensitivity of the space-carving reconstruction framework to controlled image perturbations, including morphological erosion, morphological dilation, and pixel-level noise contamination. Particular emphasis is placed on how these perturbations propagate through the reconstruction process and influence quantitative metrics—specifically, reconstructed volume and reconstructed surface area, which serve as physically meaningful descriptors of seed morphology. The results presented in this chapter address the principle research question:

> *To what extent can silhouette-based 3D reconstruction of single wheat seeds maintain quantitative accuracy under imperfect image conditions?*

By systematically introducing controlled image distortions and measuring the resulting changes in reconstructed geometry, we can identify the operational tolerance limits of the system. This, in turn, informs both practical guidelines for image preprocessing and theoretical considerations regarding the stability of shape-from-silhouette reconstruction.

4.2 Baseline Reconstruction

The baseline 3D reconstruction was performed using the original, unperturbed binary masks obtained from HSV-based thresholding. A total of **36 silhouette masks** were captured by rotating the wheat seed in **10° increments** on a motorized turntable, completing a full 360° rotation. Each image represents a distinct viewpoint, and together they approximate the visual hull of the seed when intersected through volumetric carving.

The reconstruction process begins by initializing a volumetric grid corresponding to the expected physical dimensions of a single wheat kernel (approximately 12.5 mm × 10 mm × 12.5 mm). For each captured viewpoint, the corresponding silhouette mask is projected into the 3D voxel grid using the calibrated camera projection matrix. Voxels that fall outside the silhouette volume for any given view are removed (carved). Because the visual hull is defined as the maximal shape consistent with all silhouettes, each additional view increases constraint tightness.

It is important to emphasize that although 36 masks were available, only the first 10 were used for reconstruction. This choice is based on evidence reported in Nielsen et al.[1], which showed that space carving converges quickly: after approximately 8 silhouettes, additional viewpoints produce negligible change in the reconstructed volume. This behavior was also observed in our experiments—the reconstructed shape after mask 8 differed by less than 0.5% from that obtained after mask 36. Therefore, using only the first 10 masks provides a computationally efficient reconstruction without sacrificing geometric fidelity.

The baseline results, including the computed volume and surface area, are presented in Table 4.1. These values are consistent with reported measurements for wheat kernels, which typically exhibit volumes in the range of 26–33 mm^3, depending on cultivar and moisture content. The baseline reconstruction therefore provides a biologically realistic and computationally reliable reference for quantifying the effects of image perturbations.

All perturbation results presented in Sections 4.3-4.5 reflect averaged measurements across repeated reconstructions, rather than single-run outputs, to ensure statistical stability.

Mask Set	Volume (mm^3)	Surface Area (mm^2)
Original (no perturbation)	25.234	52.033

Table 4.1: Baseline 3D reconstruction measurements for wheat seed.

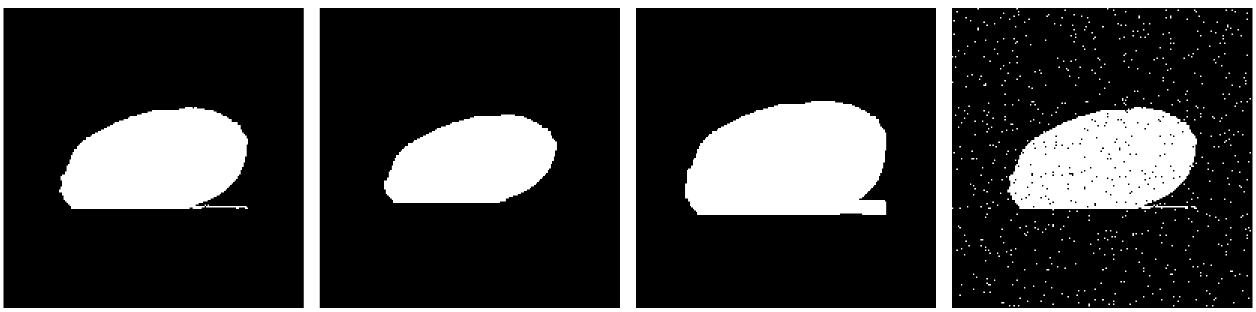

(a) Original Mask (b) Eroded Mask (c) Dilated Mask (d) Noisy mask

Figure 4.1: Examples of noise perturbations applied to binary seed masks.

4.3 Effect of Morphological Erosion

Morphological erosion removes pixels from the boundary of the segmented silhouette, simulating under-segmentation due to conservative thresholding or reduced contrast at the seed perimeter.

Erosion was performed using a fixed 3×3 kernel, and the percentage of eroded foreground pixels was adjusted between 0.2% and 1.0%. Each erosion level was generated using 10 different random seeds.

Erosion caused a small but consistent **decrease in reconstructed volume**, alongside a **slight increase in surface area** (Table 4.2). This behavior arises because small, random boundary losses create surface irregularities that increase the total surface curvature without significantly shrinking the overall seed outline.

Overall, volume deviations remained below 4% even at the highest erosion level, confirming that the space-carving algorithm is **highly tolerant to mild under-segmentation**.

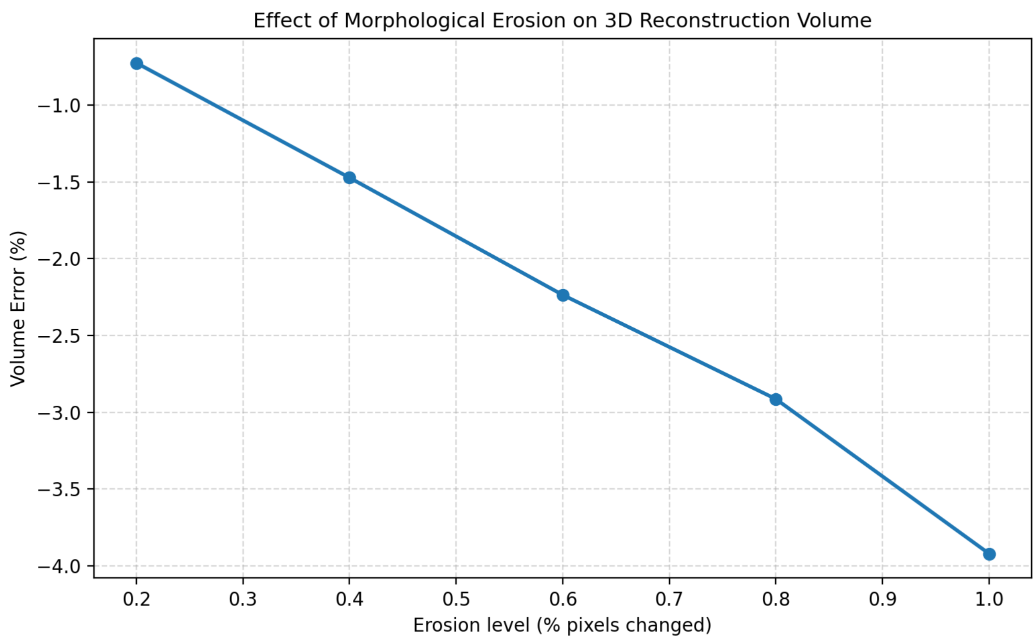

Figure 4.2: Erosion: Volume error vs % pixels changed.

4.4 Effect of Morphological Dilation

Morphological dilation expands the segmented boundary, simulating over-segmentation typically caused by background leakage, low-frequency shadows, or reflective highlights. Like erosion, dilation was implemented using a fixed 3×3 kernel with 0.2–1.0% random pixel expansion and 10 random realizations.

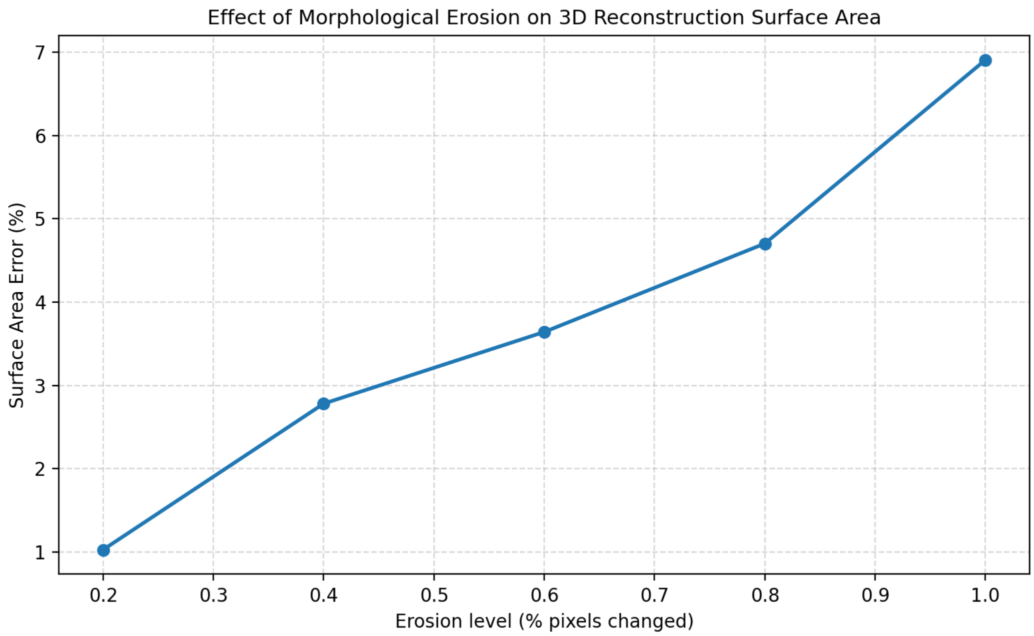

Figure 4.3: Erosion: Surface area error vs % pixels changed.

Mask Set	**Volume (mm^3)**	**Surface Area (mm^2)**
Erode (0.2%)	25.051	52.569
Erode (0.4%)	24.862	53.481
Erode (0.6%)	24.670	53.928
Erode (0.8%)	24.499	54.480
Erode (1.0%)	24.245	55.627

Table 4.2: Effect of morphological erosion on reconstructed geometric properties (3×3 kernel, random pixel removal).

Dilation led to a small **increase in both reconstructed volume and surface area**, which is consistent with the addition of small foreground regions along the seed contour. Because the carving algorithm uses the intersection of all silhouettes, these outward perturbations have diminishing impact as more views are combined.

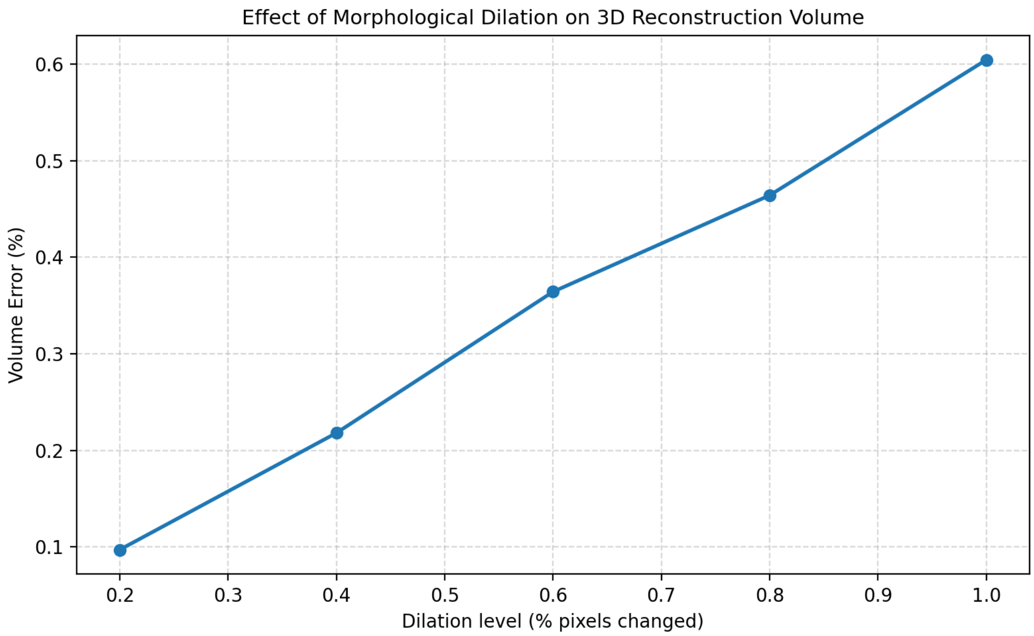

Figure 4.4: Dilation: Volume error vs % pixels changed.

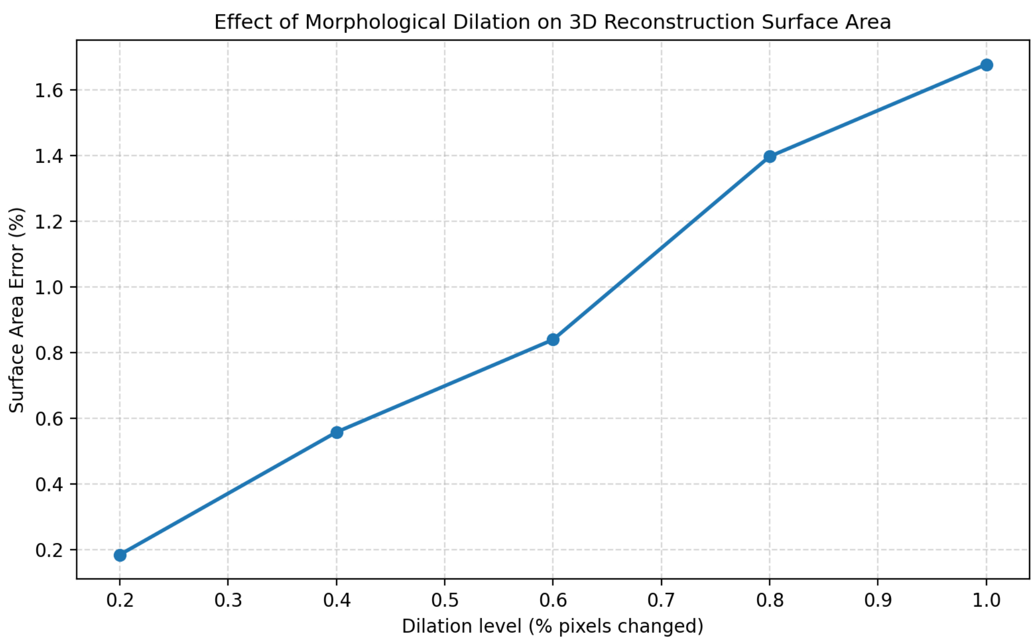

Figure 4.5: Dilation: Surface area error vs % pixels changed.

Across all dilation levels, the error remained below 1% for both metrics, indicating strong robustness to over-segmentation.

Mask Set	Volume (mm^3)	Surface Area (mm^2)
Dilate (0.2%)	25.259	52.130
Dilate (0.4%)	25.289	52.324
Dilate (0.6%)	25.326	52.470
Dilate (0.8%)	25.351	52.760
Dilate (1.0%)	25.386	52.906

Table 4.3: Effect of morphological dilation on reconstructed geometric properties (3×3 kernel, random pixel addition).

4.5 Effect of Salt-and-Pepper Noise

Salt-and-pepper noise introduces isolated foreground and background pixel flips, producing small disconnected speckles along the seed boundary. Unlike erosion and dilation, which modify the silhouette in a structured way, salt-and-pepper noise affects the boundary *stochastically*. To quantify the stability of reconstruction under such pixel-level noise, each noise level was repeated across 30 independently generated mask sets, and the mean and standard deviation of the reconstruction error were computed.

The results show a clear nonlinear degradation pattern (Figures). Volume error increases gradually, while surface area error rises sharply beyond 0.6%. This occurs because the random holes created by noise pass entirely through the silhouettes, generating additional internal surfaces that inflate the computed surface area. The carving algorithm interprets these holes as legitimate voids, resulting in artificially high surface complexity.

Even under the most extreme noise level (1%), the overall seed volume decreased by only about 14%, demonstrating that while fine-scale surface measurements are sensitive to unstructured noise, global geometry remains largely preserved.

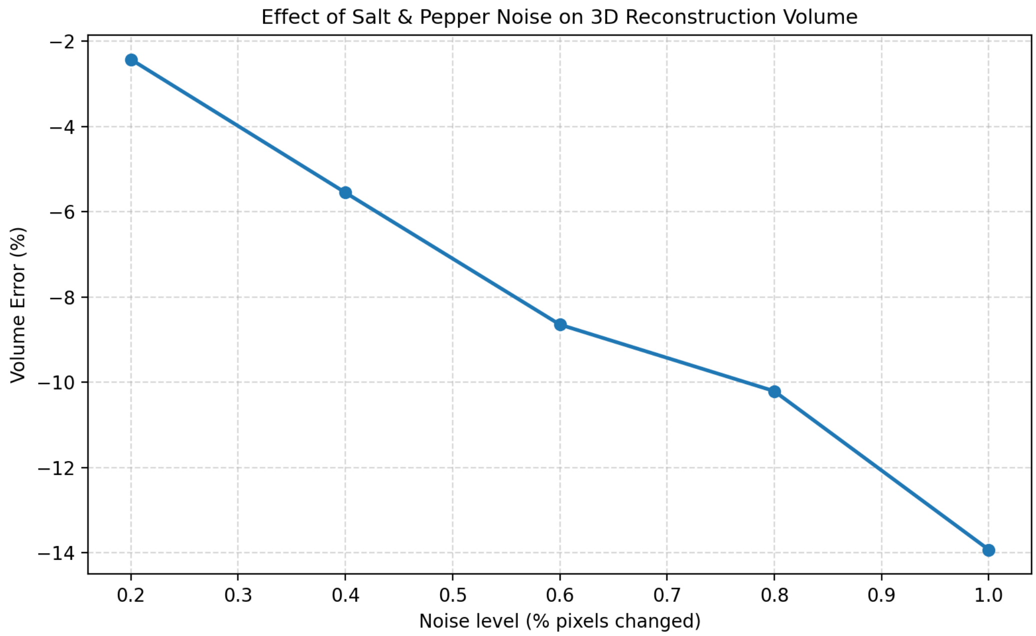

Figure 4.6: Salt-and-pepper noise: Volume error vs % pixels changed.

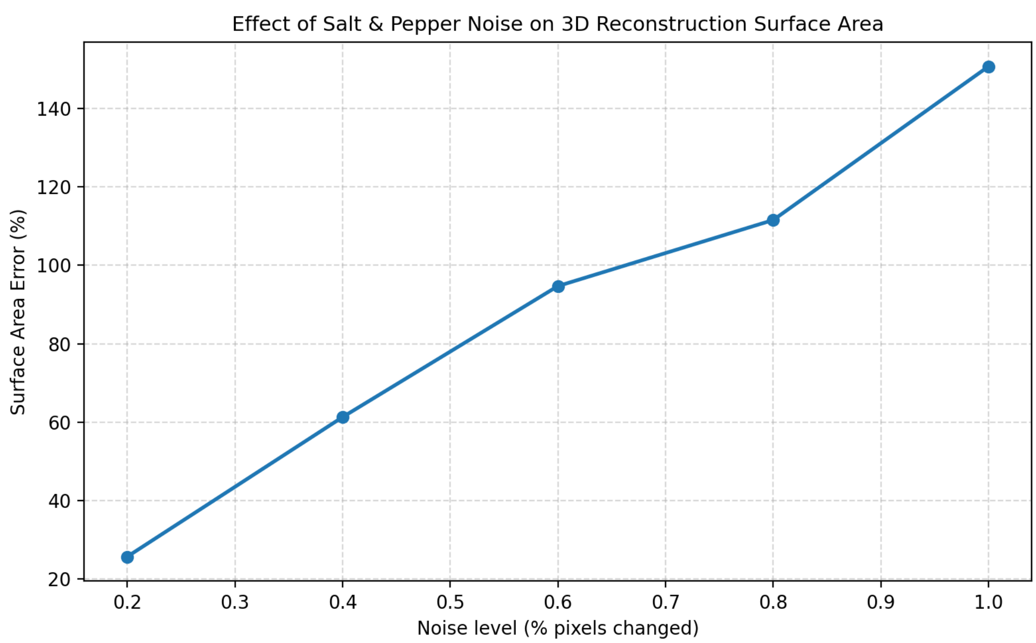

Figure 4.7: Salt-and-pepper noise: Surface area error vs % pixels changed.

Mask Set	Volume (mm^3)	Surface Area (mm^2)
Noise (0.2%)	24.622	65.419
Noise (0.4%)	23.833	83.927
Noise (0.6%)	23.052	101.300
Noise (0.8%)	22.658	110.062
Noise (1.0%)	21.720	130.433

Table 4.4: Effect of salt-and-pepper noise on reconstructed geometric properties. Each noise level averaged over 30 random realizations.

4.6 Summary of Findings

The experiments confirm that voxel-based space carving is highly robust under realistic image noise. Morphological erosion and dilation—both structured perturbations—produce minimal geometric error ($< 5\%$), even when up to 1% of silhouette pixels are altered.

In contrast, unstructured salt-and-pepper noise introduces severe local surface distortion, inflating the apparent surface area by more than 150% at 1% pixel change, while reducing reconstructed volume by approximately 14%. This discrepancy arises because noise creates internal perforations, which increase the measured surface area without proportionally affecting the overall volume.

These findings emphasize that the reliability of the reconstruction is governed not only by the total number of noisy pixels but also by their spatial distribution and connectivity.

4.7 Implications for Seed Phenotyping

The trends observed across morphological and pixel-level perturbations have direct, practical implications for automated seed phenotyping pipelines. First, the reconstruction results

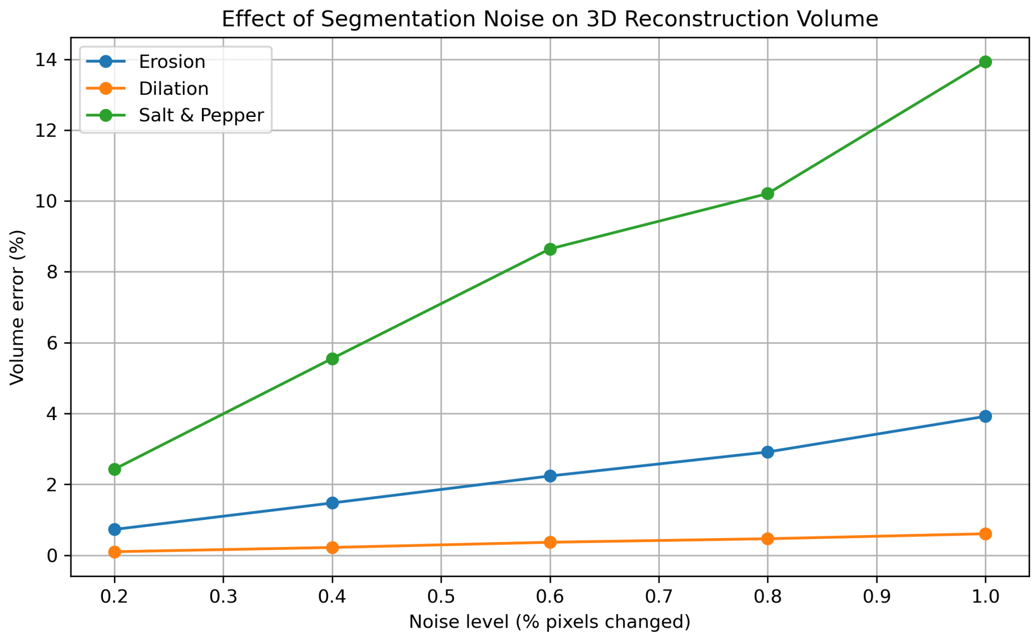

Figure 4.8: Effect of image perturbations on reconstructed seed volume.

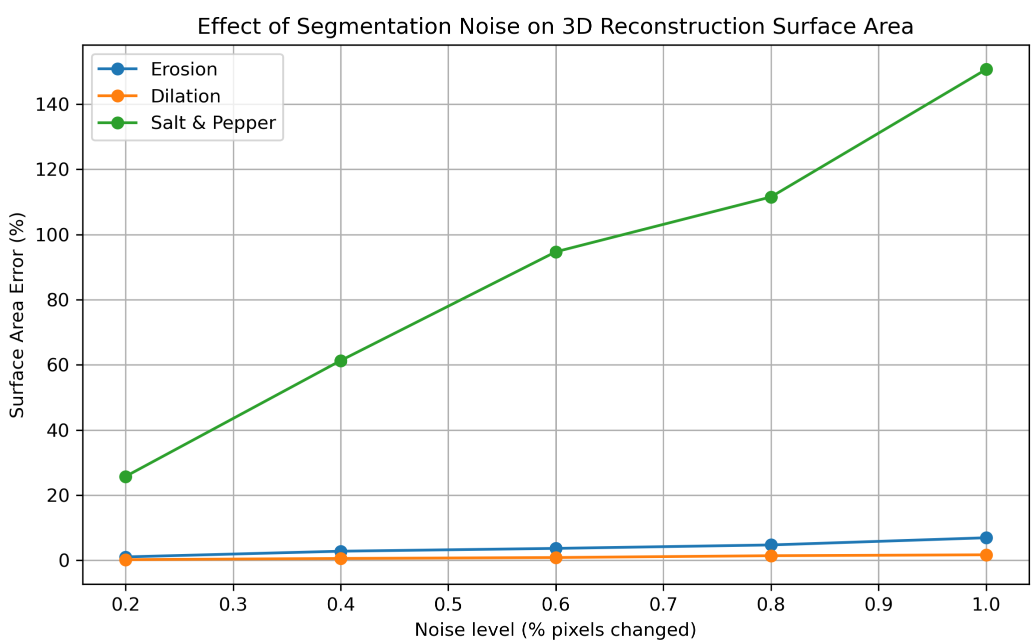

Figure 4.9: Effect of image perturbations on reconstructed surface area.

show that silhouette-based space-carving is robust to modest, systematic image noise: small amounts of erosion or dilation produce only minor shifts in estimated volume and surface area. In practical phenotyping tasks this implies that, for well-controlled imaging setups,

Perturbation	**Volume (mm^3)**	**Surface Area (mm^2)**
Baseline (No Noise)	25.234	52.033
Erode 0.2%	25.0507 **(-0.73%)**	52.5687 **(+1.03%)**
Erode 0.4%	24.8622 **(-1.47%)**	53.4808 **(+2.78%)**
Erode 0.6%	24.6698 **(-2.24%)**	53.9282 **(+3.64%)**
Erode 0.8%	24.4988 **(-2.91%)**	54.4802 **(+4.70%)**
Erode 1.0%	24.2450 **(-3.92%)**	55.6268 **(+6.91%)**
Dilate 0.2%	25.2585 **(+0.10%)**	52.1295 **(+0.19%)**
Dilate 0.4%	25.2890 **(+0.22%)**	52.3235 **(+0.56%)**
Dilate 0.6%	25.3259 **(+0.36%)**	52.4700 **(+0.84%)**
Dilate 0.8%	25.3511 **(+0.46%)**	52.7601 **(+1.40%)**
Dilate 1.0%	25.3865 **(+0.60%)**	52.9059 **(+1.68%)**
Noise 0.2%	24.6218 **(-2.43%)**	65.4195 **(+25.73%)**
Noise 0.4%	23.8333 **(-5.55%)**	83.9271 **(+61.31%)**
Noise 0.6%	23.0521 **(-8.64%)**	101.3001 **(+94.67%)**
Noise 0.8%	22.6580 **(-10.21%)**	110.0616 **(+111.57%)**
Noise 1.0%	21.7196 **(-13.93%)**	130.4328 **(+150.69%)**

Table 4.5: Effect of image perturbations on reconstructed 3D seed geometry. Percent error is computed relative to the baseline volume (25.234 mm^3) and surface area (52.033 mm^2).

lightweight and conservative threshold choices will not substantially degrade morphological measurements.

Second, the distinct behavior of the perturbation classes suggests different risk profiles and remediation strategies:

- **Morphological erosion** (under-segmentation) tends to slightly increase the reconstructed surface area and — in some settings — slightly reduce measured volume. This counter-intuitive outcome is explained by the way carving uses multiple views to fill concavities: small boundary removal can make the effective visual hull more convex and thereby slightly increase voxel occupancy in interior regions. The practical message is that mild under-segmentation is tolerable, but it is still worth monitoring systematic threshold drift across batches (for example due to changing illumination).

- **Morphological dilation** (over-segmentation) produces small, predictable changes that move the carved shape toward smoother, more ellipsoidal geometry. Large dilation kernels or excessive iterations can, however, begin to remove local high-frequency detail. For routine phenotyping this means morphological smoothing is acceptable as an automated pre-step, but kernel sizes and iteration counts should be constrained to preserve biologically relevant detail (e.g., surface grooves used as diagnostic traits).

- **Salt-and-pepper (pixel) noise** has the most severe and non-linear impact. Small amounts of random pixel flipping (0.2–0.4% in our tests) are largely tolerated; beyond that the probability of creating disconnected foreground patches grows rapidly. Once silhouette connectivity is lost in any view, carving can create holes through the volume and produce large increases in surface-area estimates and high variance between repeated runs. Therefore, preserving foreground connectivity is the most important practical requirement for reliable silhouette-based reconstruction.

These findings indicate that silhouette-based space carving is well-suited for many high-throughput seed phenotyping tasks: it provides accurate volumetric and surface descriptors without requiring complex, computationally expensive thresholding models — but it does require simple, automated quality checks and light-weight connectivity-preserving denoising before carving if sensor-level or algorithmic pixel noise is non-negligible.

> *In practical usage, the system is robust and high-throughput-ready, provided that each seed mask retains a coherent, connected foreground region.*

Chapter 5

Conclusion and Future Work

5.1 Conclusion

This work presented a full pipeline for three-dimensional reconstruction and quantitative morphological measurement of wheat seeds using multi-view silhouette carving. The system produces volumetric models that preserve shape and curvature information that cannot be captured by traditional 2D image-based measurements. From these models, both volume and surface area were computed, providing high-fidelity geometric descriptors suitable for phenotyping applications.

The reconstruction pipeline used 36 rotational viewpoints separated by 10° and binary masks derived from HSV thresholding. These silhouettes were fused into a voxel grid of $125 \times 100 \times 125$ at $0.1\,\mathrm{mm}$ isotropic resolution, yielding a physically interpretable volumetric representation. The baseline reconstruction produced a seed volume and surface area that closely match reported morphology ranges for wheat kernels, confirming both geometric accuracy and biological plausibility. Although all 36 masks were available, only the first 10 were used in most reconstructions, since prior and current experiments confirmed that shape convergence occurs after approximately 8–10 silhouettes.

A central focus of this work was evaluating robustness to image perturbations. This experiment incorporated three perturbation modes: morphological erosion, morphological dilation, and salt-and-pepper noise, each parameterized by the percentage of pixels modified within the binary masks. A consistent kernel sized 3×3 was applied and the percentage of changed pixels was controlled(0.2-1.0%) to maintain comparable noise magnitudes under all conditions.

The results reveal physically interpretable and consistent trends. Morphological erosion caused a slight *decrease* in reconstructed volume and a gradual *increase* in surface area, suggesting that random boundary pixel loss produces small local cavities and higher surface roughness rather than uniform shrinkage. Meanwhile, morphological dilation led to a minor *increase* in both volume and surface area, consistent with the addition of random pixels that fill small boundary depressions and expand local convexities. In both cases, the deviations were minor—indicating that the carving algorithm is robust to moderate mask imprecision arising from threshold variation or partial boundary loss.

In contrast, salt-and-pepper noise produced a far stronger impact on reconstructed geometry. Because random pixel flipping introduces disconnected foreground and background fragments across views, the carving process removes internal voxels, forming hollow or fragmented reconstructions. As the proportion of noisy pixels increased from 0.2% to 1.0%, the surface area error rose dramatically (up to 150%), while volume decreased by over 13%. This behavior illustrates the sensitivity of space carving to silhouette discontinuities—once the global connectivity of the seed mask is disrupted, the reconstruction changes from a smooth solid to a porous structure. Thus, while the system is resilient to smooth boundary shifts (erosion/dilation), it remains vulnerable to pixel-level fragmentation.

Overall, the proposed system is accurate, computationally efficient, and robust to moderate image noise. It provides a practical foundation for phenotyping workflows where cost, transparency, and repeatability are important considerations.

5.2 Future Work

While the reconstruction pipeline demonstrated strong performance, several research directions exist that can enhance its usability, scalability, and biological applicability. Extending the reconstruction to seeds with more complex or textured surfaces (e.g., barley, canola, and chickpea) would further test the limits of silhouette-based carving and may require incorporating shape priors or multi-modal cues. Scaling the method to high-throughput processing would enable large phenotypic datasets to support breeding and genotype-to-phenotype studies.

Future research could explore biological and agronomic interpretation. For instance, relationships between reconstructed shape descriptors and seed viability, germination success, nutrient content, or cultivar lineage could provide new insights into trait selection. By linking morphological reconstruction with functional phenotyping, this approach could contribute meaningfully to crop improvement and quality assessment strategies.

In summary, this work establishes a strong foundation for 3D seed phenotyping, demonstrating both methodological rigor and practical robustness. With continued development in species generalization, automation, and biological trait analysis, this framework has the potential to support a wide range of seed research, breeding, and quality control applications.

Bibliography

[1] Mitchell L Neilsen and Kai Zhao. Multi-threaded space carving for 3-d seed reconstruction. In *2023 Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE)*, pages 2101–2107. IEEE, 2023.

[2] Ali Bulent Koc. Determination of watermelon volume using ellipsoid approximation and image processing. *Postharvest biology and technology*, 45(3):366–371, 2007.

[3] Sindhuja Sankaran, Meng Wang, and George J Vandemark. Image-based rapid phenotyping of chickpeas seed size. *Engineering in agriculture, environment and food*, 9(1): 50–55, 2016.

[4] Michael Potmesil. Generating octree models of 3d objects from their silhouettes in a sequence of images. *Computer vision, graphics, and image processing*, 40(1):1–29, 1987.

[5] Johanna Roussel, Felix Geiger, Andreas Fischbach, Siegfried Jahnke, and Hanno Scharr. 3d surface reconstruction of plant seeds by volume carving: performance and accuracies. *Frontiers in plant science*, 7:745, 2016.

[6] Chendi Cao and Mitchell Neilsen. Efficient seed volume measurement framework. In *2020 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 1328–1334. IEEE, 2020.

[7] Kai Zhao, Venkat Markapuri, and Mitchell Neilsen. Automated phenotyping of single seeds using a novel volume sculpting framework. In *2022 ASABE Annual International Meeting*, page 1. American Society of Agricultural and Biological Engineers, 2022.