sensors

MDPI

Article

Enhancing Digital Twin Fidelity Through Low-Discrepancy Sequence and Hilbert Curve-Driven Point Cloud Down-Sampling

Yuening Ma, Liang Guo * and Min Li

School of Mathematics and Statistics, Shandong University, Wen Hua Xi Road 180, Weihai 264209, China; yuening_ma@mail.sdu.edu.cn (Y.M.); minli@sdu.edu.cn (M.L.)
* Correspondence: liangguo@sdu.edu.cn

Abstract: This paper addresses the critical challenge of point cloud down-sampling for digital twin creation, where reducing data volume while preserving geometric fidelity remains an ongoing research problem. We propose a novel down-sampling approach that combines Low-Discrepancy Sequences (LDS) with Hilbert curve ordering to create a method that preserves both global distribution characteristics and local geometric features. Unlike traditional methods that impose uniform density or rely on computationally intensive feature detection, our LDS-Hilbert approach leverages the complementary mathematical properties of Low-Discrepancy Sequences and space-filling curves to achieve balanced sampling that respects the original density distribution while ensuring comprehensive coverage. Through four comprehensive experiments covering parametric surface fitting, mesh reconstruction from basic closed geometries, complex CAD models, and real-world laser scans, we demonstrate that LDS-Hilbert consistently outperforms established methods, including Simple Random Sampling (SRS), Farthest Point Sampling (FPS), and Voxel Grid Filtering (Voxel). Results show parameter recovery improvements often exceeding 50% for parametric models compared to the FPS and Voxel methods, nearly 50% better shape preservation as measured by the Point-to-Mesh Distance (than FPS) and up to 160% as measured by the Viewpoint Feature Histogram Distance (than SRS) on complex real-world scans. The method achieves these improvements without requiring feature-specific calculations, extensive pre-processing, or task-specific training data, making it a practical advance for enhancing digital twin fidelity across diverse application domains.

Keywords: low-discrepancy sequence; Hilbert curve; point cloud down-sampling

check for updates

Academic Editor: Kenneth Loh

Received: 23 April 2025
Revised: 6 June 2025
Accepted: 10 June 2025
Published: 11 June 2025

Citation: Ma, Y.; Guo, L.; Li, M. Enhancing Digital Twin Fidelity Through Low-Discrepancy Sequence and Hilbert Curve-Driven Point Cloud Down-Sampling. *Sensors* **2025**, *25*, 3656. https://doi.org/10.3390/s25123656

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Light Detection and Ranging (LiDAR) technology is essential for creating digital twins—high-fidelity virtual representations of physical environments. However, the massive point clouds generated by modern LiDAR systems, often containing millions or billions of points, create a fundamental computational bottleneck that impedes the efficient creation and utilization of these digital models [1,2]. Down-sampling techniques, which reduce the density of point clouds while preserving key features, are therefore critical for balancing computational efficiency with model fidelity [3]. The choice of down-sampling method significantly impacts the quality and usability of the resulting digital twin, with inadequate methods potentially compromising essential geometric features, spatial relationships, and surface characteristics that are vital for downstream applications [4,5].

To mitigate these challenges, down-sampling techniques are crucial. By reducing the density of point clouds while preserving key features, down-sampling lowers the computational load and enables faster processing [3]. This, in turn, reduces power consumption, making digital twin generation and maintenance more sustainable [6]. The choice of down-sampling method significantly impacts the fidelity of the resulting point clouds. Common approaches include Simple Random Sampling, which is simple but vulnerable to randomness; Voxel Grid Down-sampling, which can regularize density but struggles to preserve sharp edges; feature-preserving techniques, which are effective but computationally expensive; Farthest Point Sampling, which enforces uniform spacing but tends to lose details in high-density regions; and learning-based methods, which can achieve state-of-the-art results but require large training datasets.

Despite the variety of existing techniques, a fundamental challenge remains: achieving an optimal balance between computational efficiency, preservation of salient geometric features, and faithful representation of the point cloud's overall spatial distribution [4,5].

This paper addresses this gap by proposing a novel down-sampling approach, LDS-Hilbert, driven by Low-Discrepancy Sequences (LDS) combined with Hilbert curve ordering. While LDS havs been explored for point cloud index selection [7], and space-filling curves (SFCs) have been utilized for points encoding and sorting [8,9], our method introduces a distinct synergistic framework. It uniquely first establishes strong, global 3D spatial coherence by directly mapping the entire point cloud with a Hilbert curve, and then employs an LDS-derived reordering mechanism for a globally uniform yet structure-aware selection from this coherent 1D sequence. Leveraging the inherent uniformity properties of LDS and the superior locality-preserving characteristics of Hilbert curves, our method aims to generate a down-sampled point cloud that better preserves both the global structure and local details of the original data compared to existing methods. By ensuring a more uniformly distributed yet structure-aware sampling, the LDS-Hilbert approach can potentially enhance the fidelity of digital twins constructed from the down-sampled data.

Our extensive experiments demonstrate that the proposed LDS-Hilbert method consistently outperforms traditional down-sampling approaches (SRS, FPS, Voxel) across diverse geometric data and evaluation metrics. Notably, it achieves significant error reductions in parametric surface fitting and superior shape and feature preservation in mesh reconstruction tasks, particularly for complex CAD models and real-world laser scans, without requiring feature-specific calculations or training data.

The key contributions of this work are as follows:

1. **A novel synergistic down-sampling framework:** We propose LDS-Hilbert, a method that uniquely combines Hilbert curve ordering for spatial coherence with Low-Discrepancy Sequences (LDS)-based reordering for globally uniform and representative point selection.
2. **Enhanced fidelity for digital twins:** We demonstrate through comprehensive experiments that LDS-Hilbert significantly improves the fidelity of digital twins constructed from down-sampled point clouds, excelling in both parametric parameter recovery and mesh reconstruction accuracy across various object types.
3. **Robustness and practicality:** The proposed method is deterministic, requires minimal parameter tuning beyond the target sample size, and shows strong performance on noisy real-world scan data, making it a practical solution for diverse applications.
4. **Insight into ordering strategies:** Through ablation studies (LDS-Hilbert vs. LDS-PCA), we provide insights into the importance of locality-preserving spatial ordering (via the Hilbert curve) as a precursor to LDS-based sampling, especially for complex geometries and larger datasets.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 details the proposed LDS-Hilbert methodology. Section 4 presents the experimental setup and results. Finally, Section 5 concludes the paper and discusses future work.

2. Literature Review

2.1. Point Cloud Down-Sampling Techniques

Point cloud down-sampling is a fundamental pre-processing step in 3D data processing, essential for managing the computational burden associated with large, dense datasets captured by technologies like LiDAR. The primary objective is to reduce the number of points from an initial set while preserving the essential geometric information and structural fidelity required for subsequent tasks, such as digital twin creation, reverse engineering, object recognition, scene reconstruction, or robotic navigation [10,11]. This section surveys established and contemporary down-sampling techniques, categorizing them based on their core strategies, while acknowledging the inherent overlaps and prevalence of hybrid approaches.

2.1.1. Random Sampling Methods

Conceptually the simplest approach, Random Sampling involves selecting a subset of points purely at random from the original point cloud [12]. Its main advantages are implementation ease and high computational efficiency. However, its stochastic nature ignores the spatial distribution and structural importance of points. This can lead to uneven point density, the potential loss of critical geometric features (especially in complex regions), and insufficient coverage, rendering it unsuitable for high-precision applications demanding high fidelity [13].

2.1.2. Spatially Structured Methods

To address the lack of spatial awareness in random methods, spatially structured techniques impose a regular or adaptive partitioning of the 3D space.

- **Voxel Grid Down-sampling**: This widely adopted method discretizes the space encompassing the point cloud into a uniform grid of cubic cells (voxels) [14]. Within each voxel containing points, a single representative point is retained, often the centroid or the point nearest the centroid. This guarantees a minimum spatial separation, effectively regularizes point density, and preserves the coarse structure. However, its primary limitation lies in potential oversimplification, particularly the loss of fine details when multiple significant points fall within a single voxel [4]. The method's performance is highly sensitive to the chosen voxel size: too large, and details are lost; too small, and the reduction ratio diminishes. Recent work explores dynamic or adaptive voxel sizes to mitigate these issues [4].
- **Octree-based Down-sampling**: Employing a hierarchical data structure, octree methods recursively subdivide the 3D space into eight octants, adapting the partitioning depth based on the local point density [15,16]. Points within leaf nodes are typically represented by a single point (e.g., centroid). This allows for more adaptive sampling than uniform voxel grids, better handling variations in point density. However, implementation is more complex, and performance depends on criteria like the maximum tree depth or points per leaf node, which still require careful tuning to balance detail preservation and computational efficiency [17].

2.1.3. Clustering-Based Methods

Leveraging the spatial distribution of points, clustering-based approaches group nearby points into clusters using algorithms like k-means or DBSCAN [18]. A representative

point (e.g., cluster centroid) is then selected from each cluster. This strategy adapts well to non-uniform densities, preserving dense regions while potentially discarding isolated points. However, it is sensitive to clustering parameters (k-value, distance thresholds) and algorithm choice, which can affect performance on varying point cloud structures. Furthermore, the computational cost, especially for iterative algorithms like k-means, can be significant for very large datasets [10].

2.1.4. Feature-Preserving Methods

These techniques prioritize the retention of points deemed geometrically significant, crucial for preserving the shape's salient characteristics.

- **Curvature/Normal-based Sampling**: Points located in regions of high surface curvature or significant surface normal variation are assigned higher importance and are preferentially retained [19]. Such points often correspond to edges, corners, or areas of intricate detail. While effective at preserving sharp features critical for accurate modeling [20], these methods are computationally more intensive due to the need for local neighborhood analysis to estimate geometric properties. They can also be sensitive to noise, which might be misinterpreted as high curvature, and require careful parameter tuning (e.g., neighborhood size, curvature thresholds). Recent efforts focus on balancing feature preservation with achieving a more uniform point distribution [21].
- **Structure-aware/Salient-point Sampling:** This category encompasses methods that identify important points based on various geometric descriptors or structural roles. Recent approaches may use advanced techniques like Gaussian processes on Riemannian manifolds [22] or implicitly identify salient points within learning frameworks focused on reconstruction quality [23].

2.1.5. Learning-Based Methods

Representing the state-of-the-art frontier, deep learning methods train neural networks to learn optimal sampling strategies. These can be broadly categorized as follows:

- Generative vs. Score-Based: Some networks directly generate the coordinates of the down-sampled points, while others assign an importance score to each original point, followed by a selection process (e.g., top-k scoring points) [24].
- Task-Agnostic vs. Task-Specific: Task-agnostic methods aim for general geometric fidelity, often optimized via reconstruction losses [24] or leveraging distilled 2D information [5]. Task-specific methods optimize sampling for downstream applications like object detection [25], classification [26], registration [11], or surface reconstruction [23].
- Advanced Architectures: Techniques like attention mechanisms and transformers are increasingly employed to capture complex inter-point relationships for more informed sampling [24,25]. Methods also address challenges like imbalanced data distributions [27].

While powerful and adaptable, learning-based methods typically require substantial labeled training data, significant computational resources for training, and their generalization to novel object types or sensor characteristics must be carefully validated [4].

2.1.6. Hybrid Methods

Recognizing that no single strategy is universally optimal, hybrid methods combine elements from different categories to leverage their respective strengths. Examples include using clustering with feature extraction [10] or combining learned sampling strategies with traditional pooling mechanisms [27].

2.1.7. Summary and Motivation

While numerous down-sampling techniques exist, each presents trade-offs between computational efficiency, data reduction ratio, and the preservation of geometric fidelity. Random methods are fast but structurally naive. Voxel and octree methods enforce spatial constraints but can lose details or require careful parameterization. Clustering adapts to density but is parameter-sensitive. Feature-based methods preserve details but are computationally heavier and noise-sensitive. Learning-based methods offer high potential but demand significant data and computational resources, with potential generalization issues.

For digital twin applications demanding high fidelity, these limitations become critical. An ideal down-sampling method for these domains should effectively reduce the data volume while robustly preserving both fine-grained surface details and the overall global structure, maintain the original point cloud's spatial density variations where meaningful, and remain computationally tractable without extensive training phases. The persistent challenges and trade-offs associated with existing methods motivate the exploration of alternative paradigms.

This paper investigates (1) Low-Discrepancy Sequences, a mathematically grounded approach focused on spatial uniformity, and (2) the Hilbert curve, a mathematically grounded method to map a multi-dimensional space (like 2D or 3D) onto a 1D path while preserving spatial locality, to address this gap, proposing a novel Low-Discrepancy-Sequences-based method designed to enhance digital twin fidelity by generating a more representative and structurally sound down-sampled point cloud.

2.2. Low-Discrepancy Sequences and Quasi-Monte Carlo Sampling

Low-Discrepancy Sequences (LDS), also known as quasi-random sequences, offer a deterministic alternative to pseudo-random number sequences (PRNGs), designed to minimize clustering and improve spatial uniformity [28]. The uniformity of a point set is formally quantified by its discrepancy [29,30]. LDS are mathematically constructed to achieve low discrepancy, ensuring points fill space more evenly than typical PRNGs [31–34].

The primary application of LDS is in Quasi-Monte Carlo (QMC) methods [28,35]. Beyond numerical integration, their inherent uniformity makes LDS valuable for various sampling tasks. For instance, in selecting a smaller, representative subset from a larger dataset, LDS can guide selection for better spatial coverage than random selection [36]. Xu et al. [7] explored using LDS-generated indices for point cloud down-sampling after an initial PCA-based ordering. While this demonstrates the potential of LDS for structured subset selection, the efficacy of such an approach heavily relies on the initial ordering's ability to capture the most relevant spatial characteristics; PCA, by focusing on global variance, may not always optimally preserve the local geometric details and intricate manifold structures crucial for high-fidelity reconstruction, especially in complex point clouds. LDS are also widely used in computer graphics, signal processing, and experimental design [36–40]. In essence, for tasks requiring evenly spread representative points, LDS offer a principled approach [41,42].

2.3. Hilbert Curve for Spatial Coherence

Beyond primary down-sampling strategies, the method used to structure or order a point cloud prior to sampling significantly influences the outcome. Space-Filling Curves (SFCs) map multi-dimensional data to a 1D sequence while attempting to preserve spatial locality [43,44]. The Hilbert curve is renowned for its superior locality-preserving properties compared to alternatives like the Z-order curve (used in [9]), minimizing large jumps in the 1D sequence for points close in the original d-space [45–47].

This property has been exploited for efficient spatial indexing [47] and adapting sequential deep learning models for point cloud analysis [48]. For instance, Chen et al. [8] introduced HilbertNet, which applies 2D Hilbert curves slice-wise to Z-axis partitioned voxels for efficient 2D convolutional processing. While innovative for dimensionality reduction, this voxel-based, slice-by-slice 2D Hilbert mapping can lead to loss of fine-grained point cloud details and may not optimally preserve overall 3D spatial locality. Separately, Li et al. [9] proposed PointSCNet, which employs Z-order curve coding on an initially sampled (e.g., via FPS) and grouped point set, followed by an equal-interval sampling along the Z-ordered sequence to select points for feature learning in classification and segmentation tasks. While Z-order provides some geometric ordering, its locality preservation is generally inferior to Hilbert curves. Moreover, the equal-interval sampling can be susceptible to aliasing if its stride resonates with periodic patterns in the data. Furthermore, PointSCNet primarily targets feature learning from sparse point sets. Its efficacy in preserving the comprehensive geometric detail required for high-fidelity digital twins from large-scale, dense point clouds remains unexplored.

The challenge in point cloud down-sampling lies in effectively leveraging these mathematical tools for optimal geometric fidelity. While LDS ensures uniform selection based on indices, its success hinges on the quality of the initial data ordering. Similarly, while Hilbert curves provide excellent spatial ordering of the entire point cloud, a subsequent naive or simplistic sampling from this order may not be optimal for capturing both global distribution and local detail. This paper investigates how a direct 3D Hilbert curve ordering of the raw point cloud, preserving its detailed geometric structure without initial voxelization or aggressive pre-sampling, can serve as a robust foundation. We then explore how this globally coherent, locality-preserving 1D sequence can be synergistically combined with LDS principles to achieve a down-sampled set that is both comprehensively representative and rich in local features.

3. Proposed Method

Throughout this paper, we use the following notation conventions: Calligraphic letters (e.g., $\mathcal{P}$) denote point sets; bold lowercase letters (e.g., $\mathbf{p}$) represent individual points or vectors; bold uppercase letters (e.g., $\mathbf{M}$) indicate meshes or matrices; and non-bold lowercase letters (e.g., d) represent scalar values. Subscripts indicate indices (e.g., $\mathbf{p}_i$ is the i-th point) or specific subsets (e.g., $\mathcal{P}_{\text{sorted}}$ is the sorted point cloud), while superscripts denote alternative versions (e.g., $\mathcal{P}'$ is the down-sampled point cloud).

We formulate the research problem as minimizing the discrepancy between the reconstructed digital twin and the true underlying surface:

$$\min_{f} D(\mathbf{M}', \mathbf{M}_{\text{true}}) = \min_{f} D(g(f(\mathcal{P})), \mathbf{M}_{\text{true}})$$

where f is the down-sampling function, g is the reconstruction function, $\mathcal{P}$ is the original point cloud, $\mathbf{M}'$ is the reconstructed digital twin, and $\mathbf{M}_{\text{true}}$ is the true underlying surface.

3.1. Algorithm Description

Our algorithm operates in three phases:

1. Ordering with Hilbert Curve Mapping to create spatial coherence.
2. LDS-driven reordering to ensure uniform coverage.
3. Contiguous subsequence selection to preserve local feature relationships.

Algorithm 1 summarizes the process of the proposed point cloud down-sampling method. An illustration of the proposed method can be found in Appendix A.

Algorithm 1 LDS-Hilbert: A point cloud down-sampling method.

Require: Original point cloud $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^{N} \subset \mathbb{R}^d$ ($d \geq 2$), Target size M ($M < N$)
Ensure: Down-sampled point cloud $\mathcal{P}' = \{\mathbf{p}'_j\}_{j=1}^{M}$

Step 1: Initial Ordering via Hilbert Curve Mapping
$N \leftarrow |\mathcal{P}|$
$\mathcal{P}_{\text{norm}} \leftarrow \emptyset$
for $j \leftarrow 1$ to d **do**
    $\min_j \leftarrow \min_{\mathbf{x} \in \mathcal{P}} \mathbf{x}_j$
    $\text{range}_j \leftarrow (\max_{\mathbf{x} \in \mathcal{P}} \mathbf{x}_j) - \min_j$
end for
for $\mathbf{p}_i \in \mathcal{P}$ **do**
    $\mathbf{p}_{\text{norm},i} \leftarrow$ vector where $\mathbf{p}_{\text{norm},i,j} = (\mathbf{p}_{i,j} - \min_j)/\text{range}_j$ for $j = 1 \ldots d$
    $\mathcal{P}_{\text{norm}} \leftarrow \mathcal{P}_{\text{norm}} \cup \{\mathbf{p}_{\text{norm},i}\}$
end for
$s \leftarrow \lfloor \log_{2^d}(N) + 0.5 \rfloor$
$\mathcal{H}_s \leftarrow \text{ConstructHilbertCurve}(s, d)$
$\mathcal{D}_{\text{Hilbert}} \leftarrow$ array of size N
$\mathcal{T} \leftarrow \text{BuildKDTree}(\mathcal{H}_s)$
for $i \leftarrow 1$ to N **do**
    $\mathbf{h}_i \leftarrow \text{NearestNeighbor}(\mathbf{p}_{\text{norm},i}, \mathcal{T})$
    $d_i \leftarrow \text{DistanceAlongCurve}(\mathcal{H}_s, \mathbf{h}_i)$
    $\mathcal{D}_{\text{Hilbert}}[i] \leftarrow d_i$
end for
$\mathcal{P}_{\text{sorted}} \leftarrow \text{Sort}(\mathcal{P}, \mathcal{D}_{\text{Hilbert}})$

Step 2: Inverse Sorting Vector Computing for LDS
$\mathbf{a} \leftarrow$ array of size N
for $i \leftarrow 1$ to N **do**
    $\mathbf{a}[i] \leftarrow (i \cdot e) \mod 1$
end for
$\mathbf{a}' \leftarrow \text{Sort}(\mathbf{a})$
$\mathbf{r} \leftarrow$ array of size N
for $i \leftarrow 1$ to N **do**
    Find k such that $\mathbf{a}'[k] = \mathbf{a}[i]$
    $\mathbf{r}[i] \leftarrow k$
end for

Step 3: Reordering Point Cloud with LDS Inverse Sorting Vector
$\mathcal{P}_{\text{reordered}} \leftarrow ()$
for $k \leftarrow 1$ to N **do**
    $\mathcal{P}_{\text{reordered}} \leftarrow \mathcal{P}_{\text{reordered}} \cup (\mathcal{P}_{\text{sorted}}[\mathbf{r}[k]])$
end for

Step 4: Select Subsequence
$\mathcal{P}' \leftarrow \mathcal{P}_{\text{reordered}}[1 \ldots M]$
return $\mathcal{P}'$

During the initial sorting phase (Step 1), the vertices of the s-order Hilbert curve $\mathcal{H}_s$ can be pre-computed and stored for commonly used orders s and dimensions d to accelerate the process.

When projecting points onto the Hilbert curve, a KD-Tree is constructed from the vertices of $\mathcal{H}_s$. This data structure is then used to efficiently find the nearest point $\mathbf{h}_i$ on the Hilbert curve for each normalized point $\mathbf{p}_{\text{norm},i}$. The Hilbert index d_i is the distance along the curve from its starting point to $\mathbf{h}_i$.

In Step 2, for generating LDS, Euler's number ($e \approx 2.71828$) is used as the transcendental number due to its simplicity and good practical performance in generating sequences with favorable distribution properties.

The inverse sorting vector $\mathbf{r}$ computed in Step 2 essentially maps the original LDS indices to their sorted ranks. When this vector is used to reorder the Hilbert-sorted point cloud in Step 3, it effectively shuffles the spatially coherent sequence $\mathcal{P}_{\text{sorted}}$ in a quasi-random manner, breaking up local clusters while attempting to maintain a good global distribution for any sub-sequence taken from the beginning.

3.2. Theoretical Foundation and Analysis

The superior performance of our LDS-Hilbert method derives from the complementary mathematical properties of Low-Discrepancy Sequences and Hilbert curve ordering, which together create a synergistic approach to point cloud down-sampling.

3.2.1. Hilbert Curve Properties

Step 1 leverages the Hilbert curve's locality-preserving property, which has been rigorously established by Moon et al. [46] and applied successfully to point cloud processing by Chen et al. [49] and Pandhare [48]. The Hilbert curve mapping $\mathcal{H} : [0,1] \to [0,1]^d$ preserves locality in a precise mathematical sense: for any $\delta > 0$, there exists $\epsilon > 0$ such that if $|s - t| < \epsilon$, then $|\mathcal{H}(s) - \mathcal{H}(t)| < \delta$.

As shown by Gotsman [50], the Hilbert curve minimizes the average locality distortion among all space-filling curves, making it optimal for preserving neighboring relationships. This mapping ensures that points spatially close in the original 3D space remain close in the 1D sequence, preserving neighborhood relationships through the transformation.

Our implementation uses the bit-manipulation approach introduced by Skilling [51] rather than recursive geometric constructions, making it computationally efficient with $O(s \cdot d)$ complexity per point. This efficiency is critical for processing large point clouds with millions of points.

3.2.2. Low-Discrepancy Sequence Properties

Low-Discrepancy Sequences inherently excel at achieving uniform coverage of sampling domains through their fundamental property: minimizing the deviation between the empirical distribution of points and the uniform distribution. This deviation, formally quantified as discrepancy, is mathematically expressed as:

$$D_N(P) = \sup_{B \in \mathcal{B}} \left| \frac{|P \cap B|}{N} - \lambda(B) \right|$$

where P is a point set with N points, $\mathcal{B}$ is a family of measurable subsets, and $\lambda(B)$ is the Lebesgue measure of set B. LDS construction ensures this discrepancy decreases at a rate of $O((\log N)^d / N)$ compared to the $O(1/\sqrt{N})$ rate of pseudo-random sequences, providing theoretically superior domain coverage with fewer samples [28,34].

The use of Euler's number e as the base for our LDS follows the well-established additive recurrence relation method for generating Low-Discrepancy Sequences [28,52,53], which guarantees optimal uniformity of the sampling distribution.

3.2.3. Synergistic Combination

The synergy in LDS-Hilbert arises from addressing the distinct challenges of achieving both local feature preservation and global sampling uniformity. While Hilbert ordering (Section 3.2.1) provides a spatially coherent 1D path, naive subset selection from this path can be biased and does not inherently guarantee optimal discrepancy for the selected

M points. Conversely, directly applying LDS principles (Section 3.2.2) to an unordered point cloud might disregard the intrinsic manifold structure and local density variations of the data.

Our method leverages the Hilbert curve to first establish this crucial 1D spatial coherence. Then, the LDS-based reordering ensures that points are selected from all segments along this coherent path in a quasi-random, uniformly distributed manner. This two-stage process thereby samples the local spatial relationships (preserved by Hilbert) with global representativeness (ensured by LDS). This integrated strategy fundamentally distinguishes our approach from prior methods that might employ SFCs primarily for voxel flattening and 2D processing [8] or use LDS with simpler, less locality-aware initial orderings, such as the PCA-based ordering explored by Xu et al. [7]. The unique contribution of the Hilbert curve component within our synergistic framework is empirically evaluated by comparing its performance against LDS-PCA [7] in our experiments.

3.2.4. Starting Index Selection and Contiguous Sampling

Our decision to select the first M points (setting $s = 1$) from the LDS-reordered sequence is theoretically justified by the mathematical properties of Low-Discrepancy Sequences. After the reordering in Step 3, the point cloud $\mathcal{P}_{\text{reordered}}$ already embodies the optimal distribution properties of the LDS, effectively distributing sample points uniformly across the entire Hilbert-ordered space. Due to this property, any contiguous block of M points from this sequence will maintain representative coverage of the original point cloud's domain.

Our decision to select the first contiguous block of M points is not only straightforward but also theoretically grounded. Since the Hilbert curve maps spatially proximate points to nearby positions in the 1D sequence, selecting the first contiguous block preserves feature clusters and geometric relationships that might otherwise be fragmented if points were selected at fixed intervals or based on rank thresholds. Therefore, in our implementation, we use a deterministic approach by consistently setting the starting index to $s = 1$ across all experiments.

3.2.5. Computation and Storage Complexity

Computationally, the primary costs of the LDS-Hilbert algorithm stem from the sorting steps and the Hilbert curve mapping. Sorting the original points by their Hilbert indices and sorting the LDS sequence to compute the inverse sorting vector both have a typical time complexity of $O(N \log N)$, where N is the number of input points.

Regarding storage for the Hilbert curve mapping, the strategy involves pre-computing vertices for commonly used orders s. A d-dimensional Hilbert curve of order s has $N_H = (2^d)^s$ vertices. Storing these N_H vertices, each with d coordinates, requires $O(N_H \cdot d)$ space, and loading them from disk is an efficient I/O operation. The order s is typically chosen such that the $N_H \approx O(N)$ to provide sufficient granularity for mapping. Consequently, the storage for pre-computed vertices scales as $O(N \cdot d)$.

The computational aspect of Hilbert curve mapping, using these N_H pre-computed vertices, involves building a spatial indexing structure (e.g., a KD-Tree) on them, which takes $O(N_H \log N_H)$ time. Since $N_H \approx O(N)$, this step is $O(N \log N)$. Subsequently, for each of the N input points, querying this structure to find the nearest curve vertex and projecting the point to the relevant curve segments takes approximately $O(\log N_H + d)$ per point, leading to a total of $O(N(\log N + d))$ for all input points.

Thus, the overall theoretical time complexity, combining sorting and the Hilbert mapping with pre-computed vertices, is dominated by these steps, generally scaling as

$O(N \log N + N \cdot d)$. This makes the method efficient relative to approaches with quadratic or higher complexity, especially for large datasets.

3.2.6. Summary of Advantages

This theoretical foundation explains our method's empirical superiority:

1. Global distribution preservation: By sampling uniformly across the Hilbert-ordered sequence rather than in 3D space directly, our method maintains the original density variations that often encode important structural information.
2. Local feature preservation: The locality preservation property of the Hilbert curve ensures that important local features are not arbitrarily split across distant indices, making them less susceptible to being eliminated during the sampling process.
3. Deterministic coverage guarantees: Unlike stochastic methods, our approach provides mathematical guarantees on sampling coverage through the discrepancy bounds of LDS, ensuring no significant regions of the point cloud are overlooked.

We conduct four experiments to test this theoretical analysis across diverse geometry types and metrics.

4. Materials and Experiments

4.1. Digital Twin Categorization and Fidelity Measurement

Digital twins represent virtual replicas of physical objects, and their construction methodology significantly impacts their fidelity and utility. As illustrated in Figure 1, we categorize digital twins ($\mathbf{M}'$) created from down-sampled point clouds ($\mathcal{P}'$) into three distinct types based on the original object's geometry ($\mathbf{M}_{\text{true}}$) and the reconstruction approach ($g : \mathcal{P}' \rightarrow \mathbf{M}'$).

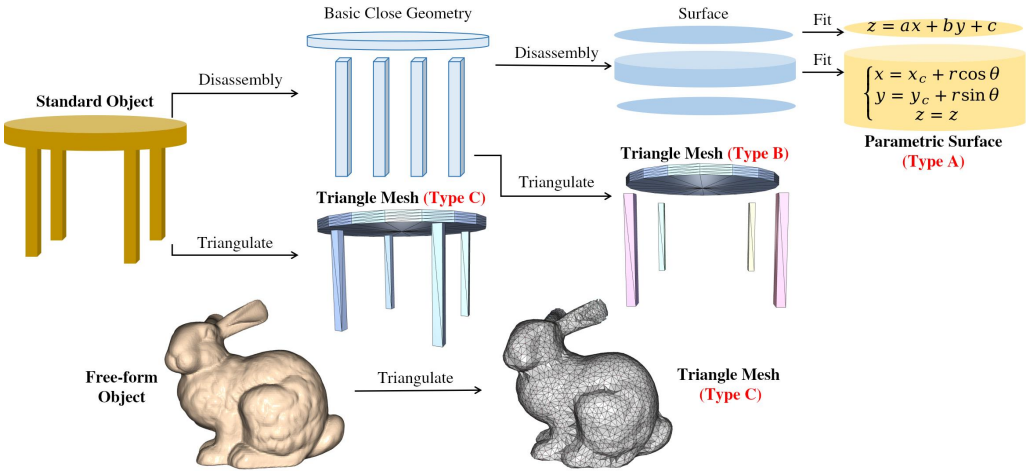


Figure 1. Digital twin categorization based on object complexity and reconstruction approach.

For objects with predominantly primitive or standard surface geometry, two reconstruction approaches are viable. Type A digital twins employ parametric representation, where the model is defined by fitted geometric primitives (planes, spheres, cylinders, etc.) or standard parametric surfaces. This approach involves segmenting the down-sampled point cloud and applying least-squares methods to estimate parameters ($\mathbf{w}_{\text{est}}$) that define these geometric entities. As shown in Figure 1, surfaces like planes and cylinders can be effectively represented by their mathematical equations, yielding compact yet precise digital representations.

Alternatively, Type B digital twins utilize mesh-based representation for parameterizable objects. Although the underlying geometry could be described parametrically, this approach directly applies surface reconstruction algorithms (such as Ball Pivoting, Poisson

Surface Reconstruction, or Alpha Shapes) to the down-sampled point cloud. The resulting triangle mesh provides visual fidelity while maintaining a direct connection to the original geometric properties. Figure 1 demonstrates how a standard object like a table can be disassembled into basic closed geometries (cylindrical top and rectangular legs), which can then be individually triangulated to form a complete mesh model.

For objects with free-form or organic geometry (such as the bunny model shown in Figure 1), Type C digital twins are inherently represented as triangle meshes since their complexity makes parametric fitting impractical. This approach applies surface reconstruction algorithms directly to the down-sampled point cloud to generate a mesh that captures the intricate shapes and features of complex objects. The figure clearly illustrates how triangulation transforms the smooth, organic bunny shape into a detailed mesh representation capable of preserving its distinctive features.

The fidelity assessment for each digital twin type requires specialized metrics tailored to their representation. For Type A twins, we evaluate parameter accuracy, while Types B and C necessitate geometric comparison measures between reconstructed surfaces and ground truth data. This systematic categorization forms the foundation for our experimental evaluation of different down-sampling methods and their impact on digital twin fidelity.

The assessment of fidelity, denoted as $D(\mathbf{M}', \mathbf{M}_{\text{true}})$, employs metrics specifically chosen based on the type of digital twin being evaluated.

For Type A digital twins, where the reconstructed model $\mathbf{M}'$ is represented by a fitted surface defined by a parameter vector $\mathbf{w}$, fidelity is measured by the discrepancy between the estimated parameters $\mathbf{w}_{\text{est}}$ and the ground truth parameters $\mathbf{w}_{\text{true}}$. A robust method for quantifying this difference is the L1 norm, chosen for its resilience to outliers. This metric is calculated as $D_{\text{L1}} = \|\mathbf{w}_{\text{est}} - \mathbf{w}_{\text{true}}\|_1 = \sum_{i=1}^{k} |w_{\text{est},i} - w_{\text{true},i}|$, where k represents the number of parameters in the vector $\mathbf{w}$.

When the digital twin $\mathbf{M}'$ is represented as a mesh (applicable to both Type B and Type C), the fidelity assessment relies on metrics that quantify the geometric discrepancy between the reconstructed representation and the ground truth. This involves comparing either the down-sampled point cloud $\mathcal{P}'$ or the reconstructed mesh $\mathbf{M}'$ against the ground truth mesh $\mathbf{M}_{\text{true}}$ or the ground truth point cloud $\mathcal{P}_{\text{true}}$. Several geometric metrics are utilized for this purpose.

One such metric is the Point-to-Mesh Distance (D_{p2m}). This asymmetric measure calculates the average distance from each point $\mathbf{p}'_i$ in the down-sampled point cloud $\mathcal{P}' = \{\mathbf{p}'_i\}_{i=1}^{M}$ to its nearest corresponding point $\mathbf{p}_{\text{true}}$ on the continuous surface $\mathcal{S}(\mathbf{M}_{\text{true}})$ of the ground truth mesh. It effectively quantifies how closely the sampled points conform to the actual surface, providing insight into reconstruction accuracy from the sample's perspective. The formulation is $D_{\text{p2m}}(\mathcal{P}', \mathbf{M}_{\text{true}}) = \frac{1}{M} \sum_{i=1}^{M} \min_{\mathbf{p}_{\text{true}} \in \mathcal{S}(\mathbf{M}_{\text{true}})} \|\mathbf{p}'_i - \mathbf{p}_{\text{true}}\|_2$, and it is often computed using tools like Metro [54].

Another key metric is the Chamfer Distance (D_{CD}), a symmetric measure evaluating the similarity between two point sets, $\mathcal{P}'$ and $\mathcal{P}_{\text{true}}$. It provides a global assessment of alignment by summing the average distance from each point in one set to its nearest neighbor in the other set, and vice versa. Using non-squared distances, the formulation is $D_{\text{CD}}(\mathcal{P}', \mathcal{P}_{\text{true}}) = \frac{1}{M} \sum_{\mathbf{p}' \in \mathcal{P}'} \min_{\mathbf{p}_{\text{true}} \in \mathcal{P}_{\text{true}}} \|\mathbf{p}' - \mathbf{p}_{\text{true}}\|_2 + \frac{1}{N} \sum_{\mathbf{p}_{\text{true}} \in \mathcal{P}_{\text{true}}} \min_{\mathbf{p}' \in \mathcal{P}'} \|\mathbf{p}_{\text{true}} - \mathbf{p}'\|_2$. Although the original concept is older [55], it has been widely adopted, particularly in deep learning contexts [56], often using the sum of the average squared distances in implementations.

To address the sensitivity of the standard Hausdorff Distance [57] to outliers, the Percentile Hausdorff Distance ($D_{\text{H,p}}$) is employed. This more robust metric calculates a specific percentile (e.g., 95th) of the nearest neighbor distances rather than the max-

imum. We consider its two directional components separately, as discussed in metric evaluation surveys [58]. The sample-to-ground truth component ($D_{\text{H,95, s2g}}$) measures the 95th percentile of distances from points in $\mathcal{P}'$ to their nearest neighbors in $\mathcal{P}_{\text{true}}$, indicating how well $\mathcal{P}'$ is covered by the true surface while tolerating some outliers in $\mathcal{P}'$. Conversely, the ground truth-to-sample component ($D_{\text{H,95, g2s}}$) measures the 95th percentile of distances from points in $\mathcal{P}_{\text{true}}$ to their nearest neighbors in $\mathcal{P}'$, indicating how well the sample $\mathcal{P}'$ covers the true surface, tolerating some unrepresented points in $\mathcal{P}_{\text{true}}$. Letting $d_{i\to\text{true}} = \min_{\mathbf{p}_{\text{true}}\in\mathcal{P}_{\text{true}}} \|\mathbf{p}'_i - \mathbf{p}_{\text{true}}\|_2$ and $d_{j\to\text{sample}} = \min_{\mathbf{p}'\in\mathcal{P}'} \|\mathbf{p}_{\text{true},j} - \mathbf{p}'\|_2$, these are formulated as $D_{\text{H,95, s2g}}(\mathcal{P}', \mathcal{P}_{\text{true}}) = \text{Percentile}_{95\%}(\{d_{i\to\text{true}}\}_{i=1}^{M})$ and $D_{\text{H,95, g2s}}(\mathcal{P}', \mathcal{P}_{\text{true}}) = \text{Percentile}_{95\%}(\{d_{j\to\text{sample}}\}_{j=1}^{N})$.

Finally, the Viewpoint Feature Histogram (VFH) Distance (D_{VFH}) assesses similarity based on global shape and viewpoint characteristics. VFH itself is a global descriptor for 3D point clouds, encoding geometric information derived from surface normals and viewpoint directions [59,60]. The D_{VFH} is not a direct spatial distance between points but rather a distance computed between the VFH descriptor histograms ($\mathbf{h}_{\text{VFH}}$) of the reconstructed model/point cloud ($\mathcal{P}'$ or $\mathbf{M}'$) and the ground truth model/point cloud ($\mathcal{P}_{\text{true}}$ or $\mathbf{M}_{\text{true}}$). A smaller distance, typically calculated using the L2 norm $d_{\text{hist}}(\mathbf{h}_1, \mathbf{h}_2) = \|\mathbf{h}_1 - \mathbf{h}_2\|_2$, indicates greater similarity in the overall shape and viewpoint properties captured by the histograms: $D_{\text{VFH}}(\mathcal{P}', \mathcal{P}_{\text{true}}) = d_{\text{hist}}(\mathbf{h}_{\text{VFH}}(\mathcal{P}'), \mathbf{h}_{\text{VFH}}(\mathcal{P}_{\text{true}}))$.

4.2. Experimental Design

To rigorously evaluate the efficacy of the proposed Low-Discrepancy Sequences-driven down-sampling method using Hilbert curve ordering (LDS-Hilbert), we designed a series of four comprehensive experiments. These experiments assess the method's performance across a spectrum of object complexities and digital twin reconstruction tasks, comparing it against three widely recognized benchmark down-sampling techniques: Simple Random Sampling (SRS), Farthest Point Sampling (FPS), and Voxel Grid Filtering (Voxel). Furthermore, to isolate the contribution of the Hilbert curve ordering component within our proposed framework, we include an ablation study in each experiment by evaluating a variant of our method that replaces the Hilbert curve sorting step with sorting based on the first principal component obtained via Principal Component Analysis (PCA), referred to as LDS-PCA.

All experiments were implemented in Python 3.8 on a laptop with an Intel Core i7-11800H CPU (Intel, Santa Clara, CA, USA). Core point cloud and mesh operations utilized Open3D [61] and NumPy [62]. VFH computation leveraged the Python interface to the Point Cloud Library (PCL) [14]. Farthest Point Sampling (FPS) was implemented via the fpsample library [63], and Hilbert curve generation used the numpy-Hilbert-curve package [64] which relies on Skilling's work [51]. Additional libraries, such as Scikit-learn [65] and SciPy [66], were employed for specific tasks like PCA and parametric fitting.

To provide a direct comparison of LDS-Hilbert against the other methods (SRS, FPS, Voxel, LDS-PCA), we introduce a **Performance Ratio** for each fidelity metric and each object. Let $D_{k,j}(\text{Method})$ denote the calculated value of the k-th fidelity metric (i.e., D_{L1}, D_{p2m}, D_{CD}, $D_{\text{H,95, s2g}}$, $D_{\text{H,95, g2s}}$, D_{VFH}) for the j-th object using a specific down-sampling method. For a given comparison method Comp (where Comp $\in$ SRS, FPS, Voxel, LDS-PCA), the Performance Ratio relative to LDS-Hilbert is calculated as:

$$R_{k,j}(\text{Comp vs. LDS-Hilbert}) = \frac{D_{k,j}(\text{Comp})}{D_{k,j}(\text{LDS-Hilbert})}$$

Since lower values indicate higher fidelity for all metrics used in these experiments, a ratio $R_{k,j} > 1$ signifies that the LDS-Hilbert method achieved a better (lower) score and thus outperformed the comparison method 'Comp' for that specific metric k and object j.

To summarize the overall performance within each experiment, we calculate the percentage of objects for which LDS-Hilbert outperforms each comparison method on a given metric. Let N_{obj} be the total number of objects in an experiment. For a specific metric k and comparison method Comp, we count the number of objects N_{win}, where $R_{k,j}(\text{Comp vs. LDS-Hilbert}) > 1$. The **Win Percentage** is then:

$$\text{Win Percentage}_k(\text{LDS-H vs. Comp}) = \frac{N_{\text{win}}}{N_{\text{obj}}} \times 100\%$$

A Win Percentage greater than 50% indicates that LDS-Hilbert outperformed the comparison method 'Comp' on that specific metric for the majority of objects within that experiment. This systematic evaluation across diverse datasets, reconstruction tasks, and metrics aims to provide robust evidence regarding the efficacy of the proposed LDS-Hilbert down-sampling method in enhancing digital twin fidelity.

4.3. Experiment 1: Parametric Surface Fidelity

This experiment evaluated the capability of various down-sampling methods (LDS-Hilbert, LDS-PCA, SRS, FPS, Voxel) to retain geometric information essential for accurately recovering the parameters of simple, known shapes (Type A Digital Twin). The dataset consisted of synthetic point clouds generated by sampling points uniformly from seven distinct analytical surfaces: plane, sphere, ellipsoid, torus, cylinder, cone, and paraboloid, each defined by specific ground truth parameters (e.g., sphere radius = 1, ellipsoid semi-axes = 3,2,1; see Appendix B for complete details). Points were sampled within defined parameter ranges for each surface type. To simulate measurement error, controlled Gaussian noise ($\sigma = 0.05$ relative to object scale) was added independently to the coordinates of each point.

Point clouds were generated at three sizes (N = 9180, N = 37,240 and N = 95,760), resulting in 21 unique synthetic objects. Each down-sampling method reduced the point clouds to 10% and 20% of their original size (ratio = 0.10 and ratio = 0.20). The resulting down-sampled point cloud $\mathcal{P}'$ was then used to fit the corresponding known parametric surface equation using a least-squares optimization method (see Appendix B for details). Performance was measured using the L1 norm of the parameter difference (D_{L1}) between the estimated parameters $\mathbf{w}_{\text{est}}$ derived from $\mathcal{P}'$ and the known ground truth parameters $\mathbf{w}_{\text{true}}$.

Figure 2 illustrates the seven parametric surfaces used in our evaluation. These surfaces represent a diverse set of geometric primitives commonly encountered in CAD models and industrial environments, ranging from simple planar surfaces to more complex curved geometries like tori and paraboloids. This diversity enables comprehensive evaluation of each down-sampling method's ability to preserve parametric characteristics across different surface types.

4.3.1. Quantitative Results

Tables 1 and 2, respectively, present the Performance Ratio and Win Percentage metrics for the L1 parameter distance. The Performance Ratio quantifies the relative error of each method compared to LDS-Hilbert, with values greater than 1.0 indicating LDS-Hilbert's superior performance. The Win Percentage shows the proportion of test cases where LDS-Hilbert outperformed each comparison method. To analyze the effects of point cloud size (N) and the down-sampling ratio on the results, the data are grouped by these factors, with each value representing the average across all surface types.

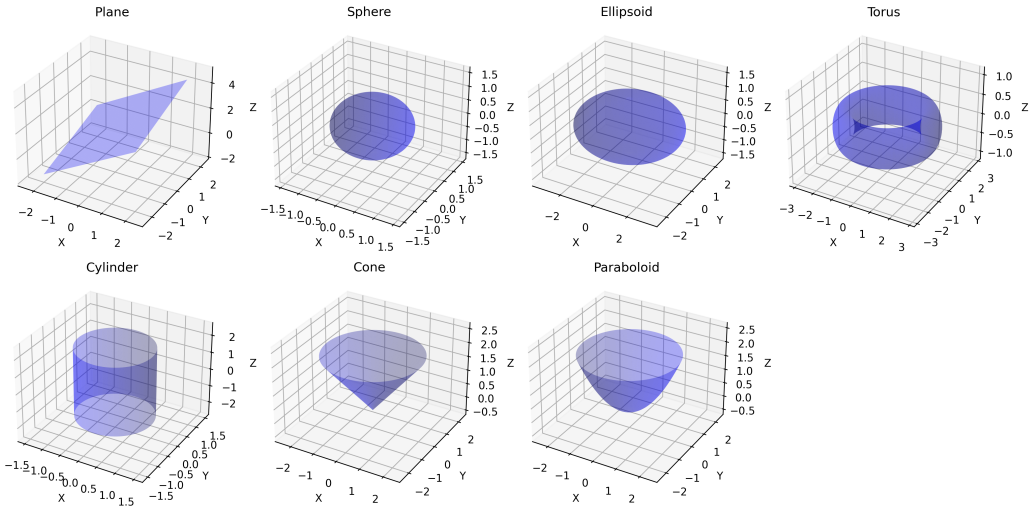


Figure 2. The seven parametric surfaces used in Experiment 1.

Table 1. Performance Ratios for LDS-Hilbert over other methods on parametric surface fidelity (D_{L1}).

N	Ratio	SRS	FPS	Voxel	LDS-PCA
9180	10%	1.40	2.59	2.68	1.47
	20%	0.97	1.66	1.48	0.84
37,240	10%	0.90	2.04	1.81	0.90
	20%	1.17	2.15	1.98	0.96
95,760	10%	1.29	3.07	2.97	1.15
	20%	1.20	2.54	2.50	1.18
Average:		1.16	2.34	2.24	1.08

Table 2. Win Percentage of LDS-Hilbert over other methods on parametric surface fidelity (D_{L1}).

N	Ratio	SRS	FPS	Voxel	LDS-PCA
9180	10%	57.14%	100.00%	100.00%	42.86%
	20%	28.57%	71.43%	71.43%	14.29%
37,240	10%	42.86%	71.43%	85.71%	42.86%
	20%	57.14%	100.00%	85.71%	57.14%
95,760	10%	100.00%	100.00%	100.00%	71.43%
	20%	100.00%	100.00%	100.00%	100.00%
Average:		64.29%	90.48%	90.48%	54.76%

4.3.2. Key Findings

Three critical patterns emerge from this analysis:

1. **Scale-dependent advantage of LDS-Hilbert:** The performance benefit of LDS-Hilbert becomes more pronounced as point cloud size (N) increases. At $N = 9180$, its advantage over methods like SRS or LDS-PCA is marginal (e.g., ratios of 0.97 vs. SRS and 0.84 vs. LDS-PCA at 20% sampling). However, at N = 95,760, LDS-Hilbert significantly outperforms them (e.g., ratios up to 1.29 vs. SRS and 1.18 vs. LDS-PCA, with 100% win rates against both in some cases). This scale dependency likely arises because larger N allows the Hilbert curve (with a correspondingly higher order s) to provide a more granular, locality-preserving 1D mapping, and the LDS reordering can more effectively distribute selection ranks across this longer, more detailed sequence.

2. **Consistent superiority over structured methods:** Regardless of point cloud size, LDS-Hilbert significantly outperforms both the FPS and Voxel methods across nearly all test conditions. Performance ratios consistently remain above 1.24 and reach as high as 3.07 for FPS at N = 95,760 with 10% down-sampling. Table 2 demonstrates that for the largest point clouds, LDS-Hilbert achieves a 100% win rate against these methods, indicating universal superiority for parameter recovery tasks at scale.
3. **Evolving role of Hilbert vs. PCA ordering:** The comparison with LDS-PCA (using Principal Component Analysis for initial ordering) shows that at smaller scales ($N = 9180$), LDS-PCA can be competitive to LDS-Hilbert. This suggests PCA's global variance capture might suffice for simpler distributions. However, at N = 95,760, LDS-Hilbert consistently surpasses LDS-PCA. This indicates that as point density increases, the Hilbert curve's superior ability to preserve local neighborhood relationships across the entire manifold becomes more critical for accurate fitting than the global structure captured by PCA.

Notably, while the relative performance ranking of the methods remains largely consistent across different down-sampling ratios (10% vs. 20%), LDS-Hilbert demonstrates particular robustness and often a greater magnitude of performance advantage at the more aggressive 10% ratio. This highlights its effectiveness precisely when substantial data reduction is required, suggesting its advantages become more pronounced under stricter down-sampling constraints.

See Section 4.7 for a paired t-test of the significance level used to measure the superiority of LDS-Hilbert.

4.3.3. Implications

These results provide strong evidence that the LDS-Hilbert method significantly improves parametric surface reconstruction fidelity compared to traditional methods. Its advantages become most apparent when handling large point clouds requiring substantial down-sampling—precisely the scenario most relevant to practical digital twin applications where computational efficiency must be balanced with model accuracy.

4.4. Experiment 2: Mesh Reconstruction from Closed Prismatic Geometry

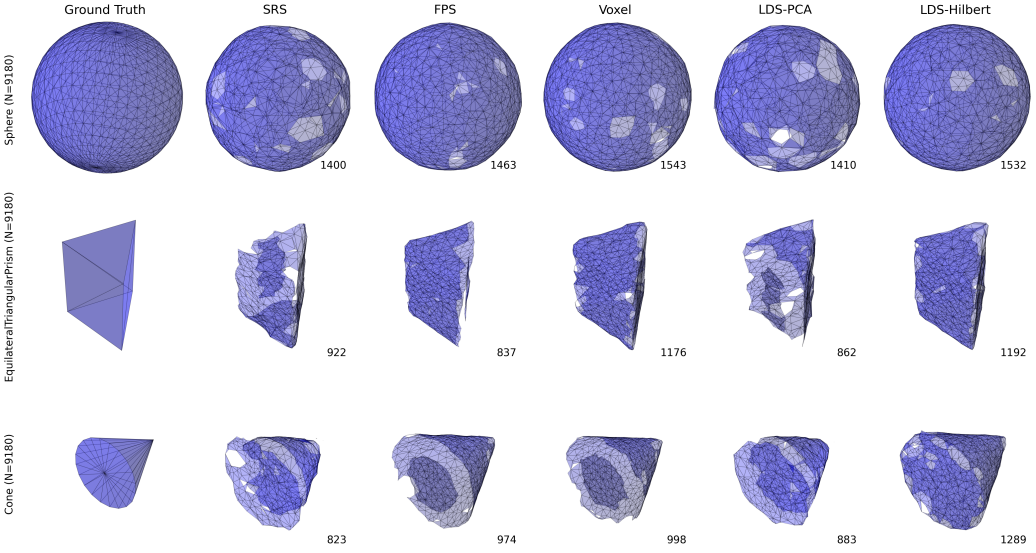
This experiment investigated the performance of down-sampling methods in reconstructing mesh-based digital twins (Type B) from point clouds representing regular, closed geometric shapes. These shapes featured both planar and curved faces, as well as sharp edges. The dataset was generated from 15 distinct mesh-based geometric models, including shapes like spheres, cylinders, various prisms (triangular, cube, rectangular, pentagonal, hexagonal), structural beams (L, T, U, H types), cones, tori, and basic polyhedra (tetrahedron, octahedron), with precisely defined dimensions based on a scale factor (see Appendix C for more details). Before sampling, each mesh model underwent a random rotation. Points were then sampled relatively uniformly from the surface of each rotated mesh using the Poisson disk sampling method, generating clouds at three sizes (N = 9180, N = 37,240 and N = 95,760) for a total of 45 objects. Gaussian noise with a standard deviation of 1 cm (assuming meter scale for the objects) was added to the coordinates of each sampled point to simulate LiDAR scanning errors. Each down-sampling method reduced the input cloud $\mathcal{P}$ to 10% (ratio = 0.10), as Experiment 1 demonstrated that the ratio had minimal impact on the relative performance of methods.

The resulting down-sampled cloud $\mathcal{P}'$ was then used as input for surface reconstruction via the Ball Pivoting algorithm (g) to create a mesh digital twin $\mathbf{M}'$ (see Appendix D for Ball Pivoting details).

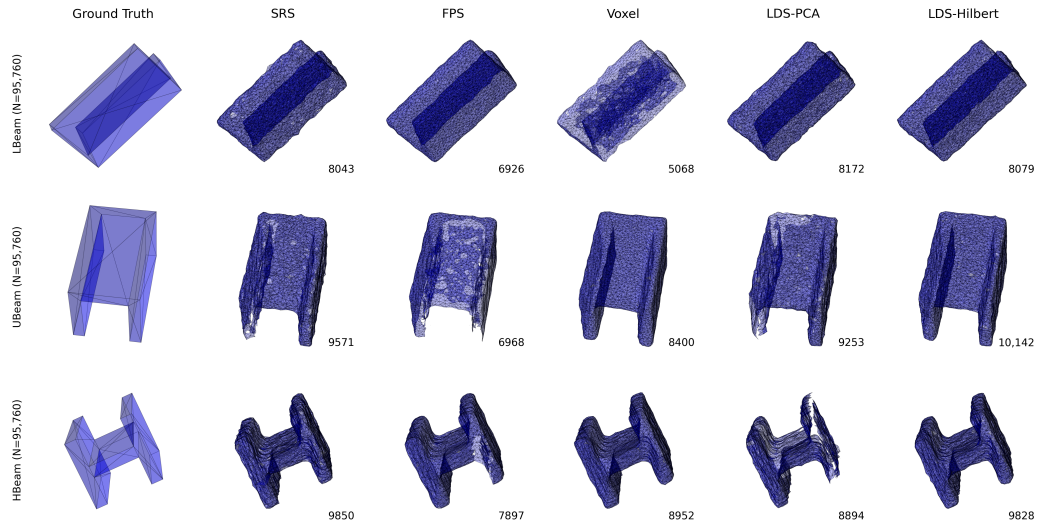
Fidelity was evaluated by comparing $\mathcal{P}'$ against the ground truth mesh $\mathbf{M}_{\text{true}}$ and the ground truth point cloud $\mathcal{P}_{\text{true}}$ (vertices of the original mesh before noise) using a suite of geometric metrics: Point-to-Mesh Distance (D_{p2m}), Chamfer Distance (D_{CD}), the 95th Percentile Hausdorff Distance components ($D_{\text{H95,s2g}}$ and $D_{\text{H95,g2s}}$), and Viewpoint Feature Histogram Distance (D_{VFH}).

4.4.1. Visual Assessment

Figure 3a,b provide visual comparisons of the reconstruction results across different down-sampling methods. Figure 3a shows the results for simpler geometries (sphere, triangular prism, and cone), while Figure 3b presents more complex structural shapes (L-beam, U-beam, and H-beam).



(**a**)



(**b**)

Figure 3. Visual comparison of surface reconstruction results for: (**a**) Sphere, triangular prism, and cone; (**b**) Structural beams (L, U, and H profiles).

Visual inspection of Figure 3a reveals that for the sphere ($N = 9180$), all methods produce reasonably complete reconstructions, but SRS, FPS, and to some extent LDS-PCA, create noticeable holes or irregularities on the surface. The LDS-Hilbert method

produces a more complete and uniform sphere reconstruction with a high face count (1532), comparable to the Voxel method (1543).

For shapes with sharp features, such as the triangular prism ($N = 9180$) and cone ($N = 9180$), the differences become more pronounced. SRS produces highly irregular reconstructions with significant gaps, while FPS creates smoother but sometimes incomplete reconstructions. The LDS-Hilbert method achieves the highest face counts (1192 for the triangular prism and 1289 for the cone), indicating more complete mesh reconstruction that better preserves the sharp geometric features.

Figure 3b demonstrates even more significant differences with complex structural shapes. The SRS method tends to produce irregular meshes with inconsistent densities. While FPS generates smoother surfaces, it often fails to maintain proper thickness in the thin sections of these complex shapes. The Voxel method struggles with the interior corners of the U-beam and H-beam. In contrast, the LDS-Hilbert method consistently achieves the highest face counts (8079 for L-beam, 10,142 for U-beam, and 9828 for H-beam) and visually preserves both the overall structure and fine details of these challenging geometries.

4.4.2. Quantitative Results

Tables 3 and 4 present the Performance Ratio and Win Percentage metrics for each evaluation criterion, grouped by point cloud size.

The Performance Ratio quantifies the relative error of each comparison method to LDS-Hilbert (values greater than 1.0 indicate LDS-Hilbert's superior performance), while the Win Percentage shows the proportion of test cases where LDS-Hilbert outperformed each method.

The quantitative metrics align with the visual observations from Figure 3a,b, demonstrating several clear patterns across different metrics and point cloud sizes.

In addition, we also provide computational time comparisons in Appendix E to validate the efficiency of our proposed method.

Table 3. Performance Ratio on closed prismatic geometry mesh fidelity.

N	Metrics	SRS	FPS	Voxel	LDS-PCA
9180	D_{p2m}	1	1.39	0.89	1
	D_{CD}	1.08	1.06	0.94	1.07
	$D_{\text{H,95, s2g}}$	0.99	1.24	1.04	0.99
	$D_{\text{H,95, g2s}}$	1.21	0.88	0.87	1.19
	D_{VFH}	1.06	0.99	1.09	1.08
37,240	D_{p2m}	1	1.67	1.17	1
	D_{CD}	1	1.2	0.99	0.99
	$D_{\text{H,95, s2g}}$	0.99	1.36	1.16	0.99
	$D_{\text{H,95, g2s}}$	1	0.84	0.78	0.99
	D_{VFH}	1.13	1.05	1.04	1.11
95,760	D_{p2m}	1	1.64	1.36	1
	D_{CD}	1.04	1.33	1.18	1.04
	$D_{\text{H,95, s2g}}$	1	1.39	1.26	1.01
	$D_{\text{H,95, g2s}}$	1.14	1.04	0.97	1.14
	D_{VFH}	1.96	1.96	1.85	1.86
Average:	D_{p2m}	1	1.57	1.14	1
	D_{CD}	1.04	1.2	1.04	1.04
	$D_{\text{H,95, s2g}}$	1	1.33	1.15	1
	$D_{\text{H,95, g2s}}$	1.12	0.92	0.88	1.11
	D_{VFH}	1.39	1.33	1.33	1.35

Table 4. Win Percentage on closed prismatic geometry mesh fidelity.

N	Metrics	SRS	FPS	Voxel	LDS-PCA
9180	D_{p2m}	46.67%	100.00%	6.67%	40.00%
	D_{CD}	100.00%	93.33%	0.00%	100.00%
	$D_{\text{H,95, s2g}}$	26.67%	100.00%	53.33%	46.67%
	$D_{\text{H,95, g2s}}$	100.00%	0.00%	0.00%	100.00%
	D_{VFH}	60.00%	46.67%	86.67%	73.33%
37,240	D_{p2m}	40.00%	100.00%	93.33%	46.67%
	D_{CD}	20.00%	100.00%	40.00%	13.33%
	$D_{\text{H,95, s2g}}$	33.33%	100.00%	100.00%	26.67%
	$D_{\text{H,95, g2s}}$	26.67%	0.00%	0.00%	26.67%
	D_{VFH}	66.67%	73.33%	33.33%	46.67%
95,760	D_{p2m}	66.67%	100.00%	100.00%	40.00%
	D_{CD}	100.00%	100.00%	100.00%	100.00%
	$D_{\text{H,95, s2g}}$	40.00%	100.00%	100.00%	60.00%
	$D_{\text{H,95, g2s}}$	100.00%	100.00%	13.33%	100.00%
	D_{VFH}	100.00%	100.00%	86.67%	100.00%
Average:	D_{p2m}	51.00%	100.00%	67.00%	42.00%
	D_{CD}	73.00%	98.00%	47.00%	71.00%
	$D_{\text{H,95, s2g}}$	33.00%	100.00%	84.00%	44.00%
	$D_{\text{H,95, g2s}}$	76.00%	33.00%	4.00%	76.00%
	D_{VFH}	76.00%	73.00%	69.00%	73.00%

4.4.3. Key Findings

1. **Progressive advantage with scale**: For nearly all metrics, LDS-Hilbert's advantage increases with point cloud size. This is particularly evident in the Point-to-Mesh Distance (D_{p2m}) metric, where the Performance Ratio against FPS increases from 1.39 at $N = 9180$ to 1.64 at N = 95,760, and against Voxel from 0.89 at $N = 9180$ to 1.36 at N = 95,760. Similarly, the Win Percentage against Voxel for this metric rises dramatically from 6.67% at $N = 9180$ to 100% at N = 95,760.
2. **Superior morphological fidelity**: The Viewpoint Feature Histogram (VFH) Distance shows the most dramatic improvement, with Performance Ratios at N = 95,760 reaching 1.96, 1.96, 1.85, and 1.86 against SRS, FPS, Voxel, and LDS-PCA, respectively. This substantial difference (nearly twice as good) indicates that LDS-Hilbert preserves the overall geometric characteristics and surface morphology significantly better than all comparison methods, with 100% win rates against SRS, FPS, and LDS-PCA at the largest scale.
3. **Trade-off in coverage metrics**: For the Hausdorff Distance from ground truth to sample ($D_{\text{H,95,g2s}}$), the FPS and Voxel methods achieved lower error rates than LDS-Hilbert at smaller scales, with Performance Ratios of 0.88 and 0.87, respectively, at $N = 9180$. This reflects these methods' design focus on ensuring uniform coverage. However, this advantage diminishes at larger scales, with ratios increasing to 1.04 for FPS at N = 95,760, suggesting that LDS-Hilbert's balanced approach becomes increasingly effective as complexity grows.
4. **Comprehensive distance metrics**: The Chamfer Distance (D_{CD}), which considers both how well sampled points represent the surface and how well the surface is covered by sampled points, shows LDS-Hilbert outperforming the comparison methods across most conditions. At N = 95,760, LDS-Hilbert achieves a 100% win rate against all comparison methods, demonstrating its balanced approach to addressing both objectives.

5. **Value of Hilbert ordering**: The comparison with LDS-PCA isolates the contribution of the Hilbert curve ordering. While LDS-PCA performs relatively well compared to other methods, LDS-Hilbert consistently outperforms it at the largest scale (N = 95,760) across all metrics, with Performance Ratios of 1.00–1.86 and Win Percentages of 40–100%. This confirms that the Hilbert curve's spatial coherence properties specifically enhance reconstruction quality.

The statistical significance of these results is analyzed in Section 4.7.

4.4.4. Implications and Synthesis

Combining the visual assessment with the quantitative results yields several important insights about LDS-Hilbert's performance in mesh reconstruction tasks:

1. LDS-Hilbert provides consistent, low-error results in surface reconstruction, particularly for large point clouds requiring significant down-sampling. The consistently higher face counts across diverse geometries indicate more complete reconstructions that better preserve the original shape characteristics.
2. While SRS produces down-sampled point clouds with moderate numerical differences from ground truth (Performance Ratios of 1.00–1.21 for most metrics), its high randomness and uneven distribution result in poor visual quality, evidenced by numerous holes and irregular surfaces in Figure 3a,b.
3. The FPS and Voxel methods create visually smoother reconstructions for simpler convex geometries but struggle with shapes having sharp features (triangular prisms, cones) and complex structures with concave regions (U-beams, H-beams). These limitations are reflected in their higher error rates at larger scales.
4. The most significant advantage of LDS-Hilbert appears in preserving the overall shape characteristics and surface morphology, as demonstrated by the substantial improvements in the VFH Distance at N = 95,760. This suggests that the method is particularly valuable for applications where faithful reproduction of an object's appearance and geometric features is critical.

These findings demonstrate that LDS-Hilbert offers substantial advantages for mesh reconstruction from closed prismatic geometries, with benefits that become increasingly pronounced as point cloud size and shape complexity increase—precisely the conditions encountered in practical digital twin applications.

4.5. Experiment 3: Mesh Reconstruction from Diverse CAD Models (ModelNet40)

This experiment focused on evaluating the generalization ability of the down-sampling methods across a broad spectrum of complex object shapes (Type C), often consisting of multiple components. The experiment utilized 40 CAD models selected from the ModelNet40 dataset, one from each class. A ground truth watertight mesh and corresponding point cloud were established for each model, and Gaussian noise was added to create the input point clouds. Following the standard procedure, each method down-sampled the input clouds to a 20% ratio. The resulting point clouds were then reconstructed into Type C Mesh digital twins using the Ball Pivoting algorithm. Performance was quantitatively evaluated using the same suite of geometric metrics as in Experiment 2, comparing the down-sampled cloud against the ground truth mesh and the ground truth point cloud.

4.5.1. Visual Assessment

Figure 4 provides a visual comparison of reconstruction results across different down-sampling methods for six representative objects: bottle, chair, cone, cup, flower pot, and piano. The ground truth is shown in the leftmost column, with reconstructed meshes

from each method in subsequent columns. The numbers indicate the face count of each reconstructed mesh.

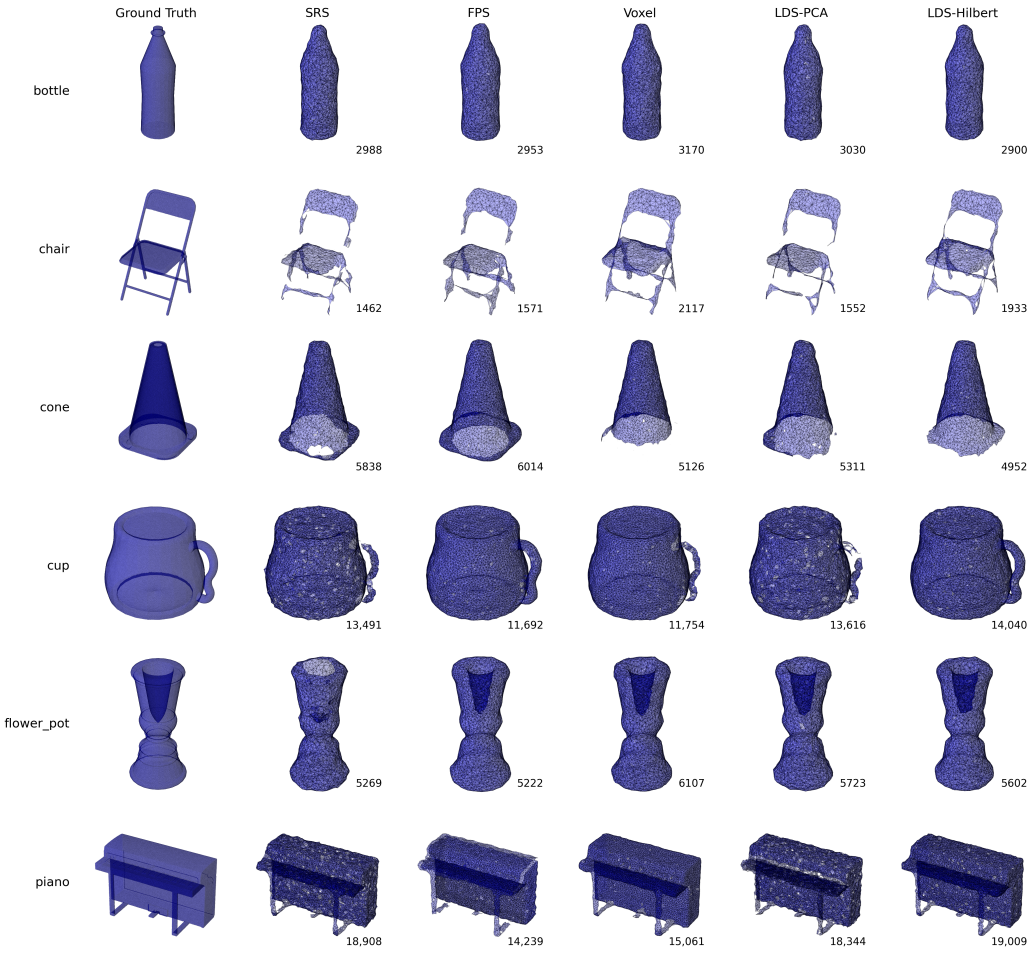


Figure 4. Visual comparison of surface reconstruction results for selected ModelNet40 objects.

Visual examination reveals clear differences in reconstruction quality across the methods. Similar to our observations in Experiment 2, the FPS and Voxel methods tend to produce smoother surfaces but frequently fail at modeling concave regions and sharp features. This is particularly evident in the chair model, where the FPS method (1571 faces) struggles to properly reconstruct the seat and back connection, while the LDS-Hilbert method (1933 faces) preserves more structural details and achieves a higher face count.

The SRS and LDS-PCA methods consistently generate reconstructions with numerous surface irregularities. For example, in the bottle model, both methods show noticeable bumps and inconsistencies along what should be smooth surfaces. While SRS occasionally achieves high face counts (as seen in the piano model with 18,908 faces), the visual quality of these reconstructions remains problematic with multiple gaps and irregular surfaces.

The LDS-Hilbert method demonstrates superior performance in preserving both smooth regions and detailed structures. This is clearly visible in the cup model, where LDS-Hilbert (14,040 faces) captures the hollow interior and handle with greater fidelity than other methods. For objects with more complex geometry, such as the piano, LDS-Hilbert (19,909 faces) achieves the highest face count and visually appears to best preserve the original shape, including the keyboard area and supporting structure.

4.5.2. Quantitative Results

Due to the extreme complexity of one model (Keyboard_0001), all methods failed to produce acceptable reconstructions, so this model was excluded from the quantitative analysis. For the remaining 39 models, Tables 5 and 6 present the Performance Ratio and Win Percentage metrics for each evaluation criterion.

Table 5. Performance Ratio on ModelNet40 mesh fidelity.

Metrics	SRS	FPS	Voxel	LDS-PCA
D_{p2m}	1	1.46	1.22	1
D_{CD}	1.03	1.16	1.07	1.03
$D_{\text{H,95, s2g}}$	1	1.27	1.19	1
$D_{\text{H,95, g2s}}$	1.08	0.93	0.89	1.08
D_{VFH}	1.18	1.21	1.44	1.17

Table 6. Win Percentage on ModelNet40 mesh fidelity.

Metrics	SRS	FPS	Voxel	LDS-PCA
D_{p2m}	61.54%	100.00%	79.49%	58.97%
D_{CD}	100.00%	97.44%	71.79%	97.44%
$D_{\text{H,95, s2g}}$	48.72%	100.00%	97.44%	53.85%
$D_{\text{H,95, g2s}}$	100.00%	2.56%	0.00%	100.00%
D_{VFH}	89.74%	66.67%	66.67%	87.18%

The Performance Ratio (Table 5) quantifies the relative error of each comparison method to LDS-Hilbert, with values greater than 1.0 indicating LDS-Hilbert's superior performance. The Win Percentage (Table 6) shows the proportion of test cases where LDS-Hilbert outperformed each comparison method. These metrics were calculated across all five geometric distance measures: Point-to-Mesh Distance (D_{p2m}), Chamfer Distance (D_{CD}), both components of the 95th percentile Hausdorff Distance ($D_{\text{H,95,s2g}}$ and $D_{\text{H,95,g2s}}$), and the VFH Distance (D_{VFH}).

4.5.3. Key Findings

1. **Comprehensive superior performance**: Table 5 shows that LDS-Hilbert outperforms all comparison methods across nearly all metrics. Compared to SRS, the LDS-Hilbert method achieves approximately 1.03 times better performance on the Point-to-Mesh and Chamfer Distance metrics, indicating better overall geometric fidelity. Against the FPS and Voxel methods, the advantage is even more pronounced, with Performance Ratios reaching 1.46 and 1.22, respectively, for the Point-to-Mesh Distance, and 1.16 and 1.07 for the Chamfer Distance.
2. **Trade-off in coverage vs. structure preservation**: While the FPS and Voxel methods show lower Hausdorff Distance from ground truth to sample ($D_{\text{H,95,g2s}}$) compared to LDS-Hilbert (Performance Ratios of 0.93 and 0.89, respectively), this advantage comes at the cost of poorer performance in other metrics. This pattern, consistent with our findings in Experiment 2, confirms that these methods prioritize uniform coverage at the expense of preserving the original geometric structure.
3. **Superior morphological fidelity**: For the VFH Distance, which captures overall shape similarities, LDS-Hilbert demonstrates substantial advantages over all comparison methods, with Performance Ratios of 1.18, 1.21, 1.44, and 1.17 against SRS, FPS, Voxel, and LDS-PCA, respectively. This indicates that LDS-Hilbert preserves the essential morphological characteristics of these complex models significantly better than other methods.

4. **Consistent Win Percentages**: Table 6 further supports LDS-Hilbert's superior performance, showing that it outperforms SRS in approximately 61.54% of cases for the Point-to-Mesh Distance and 100% for the Chamfer Distance. Against FPS and Voxel, LDS-Hilbert wins in 100% and 79.49% of cases, respectively, for the Point-to-Mesh Distance. For the comprehensive VFH metric, LDS-Hilbert achieves impressive win rates of 89.74%, 66.67%, 66.67%, and 87.18% against SRS, FPS, Voxel, and LDS-PCA, respectively.
5. **Value of Hilbert ordering**: The comparison with LDS-PCA isolates the contribution of the Hilbert curve ordering. While LDS-PCA performs relatively well compared to other methods, LDS-Hilbert consistently outperforms it across most metrics, with particularly notable advantages in the Point-to-Mesh Distance (58.97% win rate) and VFH Distance (87.18% win rate). This confirms that the Hilbert curve's spatial coherence properties specifically enhance the reconstruction quality for complex models.

For the corresponding significance test, please refer to Section 4.7.

4.5.4. Implications and Synthesis

Given the diversity of the ModelNet40 dataset, which includes objects ranging from simple geometric forms to complex multi-component assemblies, these results demonstrate that LDS-Hilbert's advantages extend beyond the basic parametric forms and closed prismatic geometries tested in previous experiments. The method appears particularly effective at handling the varied density distributions and complex feature arrangements found in real-world CAD models.

The primary advantage of LDS-Hilbert in this experiment appears to be its ability to maintain a balance between preserving local features (such as sharp edges and corners) while also ensuring sufficient point density in smooth regions. This balanced approach results in reconstructions that better capture both the overall structure and the fine details of the original models, making it particularly suitable for complex digital twin applications where fidelity across multiple scales is essential.

These findings align with the theoretical foundation of our method: the Hilbert curve's locality preservation combined with the Low-Discrepancy Sequence's uniform coverage creates a synergistic approach that offers particular advantages for complex, multi-feature objects. The strong performance across the diverse ModelNet40 dataset suggests that LDS-Hilbert provides a robust, generally applicable solution for high-fidelity down-sampling of 3D point clouds representing real-world objects.

4.6. Experiment 4: Mesh Reconstruction from Real-World Laser Scans (Stanford Dataset)

This experiment tested the down-sampling methods on challenging real-world data obtained from laser scans. These data featured complex organic shapes, intricate surface details, non-uniform point distributions, and potential scanning imperfections. The dataset consisted of 254 individual scans from four well-known models in the Stanford 3D Scanning Repository (Happy Buddha, Bunny, Dragon, Armadillo). The provided high-resolution meshes served as the ground truth, with their vertices forming the ground truth point clouds. Input point clouds were generated by adding Gaussian noise to these ground truth points. Each method down-sampled these clouds to appropriate ratios—20% for Armadillo and Dragon, and 10% for Happy Buddha and Bunny—to account for differences in model complexity. The resulting down-sampled point clouds were used to reconstruct Type C Mesh digital twins via the Ball Pivoting algorithm. Fidelity was assessed using the identical set of geometric metrics employed in previous experiments.

4.6.1. Visual Assessment

Figure 5 provides a detailed visual comparison of reconstruction quality using the Stanford Bunny as an example. Rather than showing the reconstructions separately, this visualization overlays the ground truth mesh (in green) with the reconstructed meshes (in gray) from each method. This approach clearly reveals where reconstructions deviate from the original model.

Several important observations can be made from this visualization:

1. The SRS, LDS-PCA, and LDS-Hilbert methods produce meshes that interlock with the ground truth mesh, indicating their overall dimensions closely match the original. In contrast, the FPS and Voxel methods generate meshes that completely envelop the ground truth, suggesting systematic overestimation of the model dimensions.
2. The yellow circles highlight two important detail regions: the ear and knee indentations of the bunny model. The LDS-Hilbert method successfully preserves these concave details, while the SRS, FPS, and Voxel methods tend to "smooth over" these features, losing important geometric characteristics. This observation confirms LDS-Hilbert's superior ability to maintain fine details in organic, complex models.
3. Red circles on the SRS and LDS-PCA reconstructions indicate significant surface defects and irregularities not present in the ground truth. These artifacts are largely absent in the LDS-Hilbert reconstruction, which achieves a better balance between detail preservation and surface smoothness.

The visual differences in detail preservation highlighted in Figure 5 are further substantiated by the quantitative metrics presented in Tables 7 and 8, which we will now discuss.

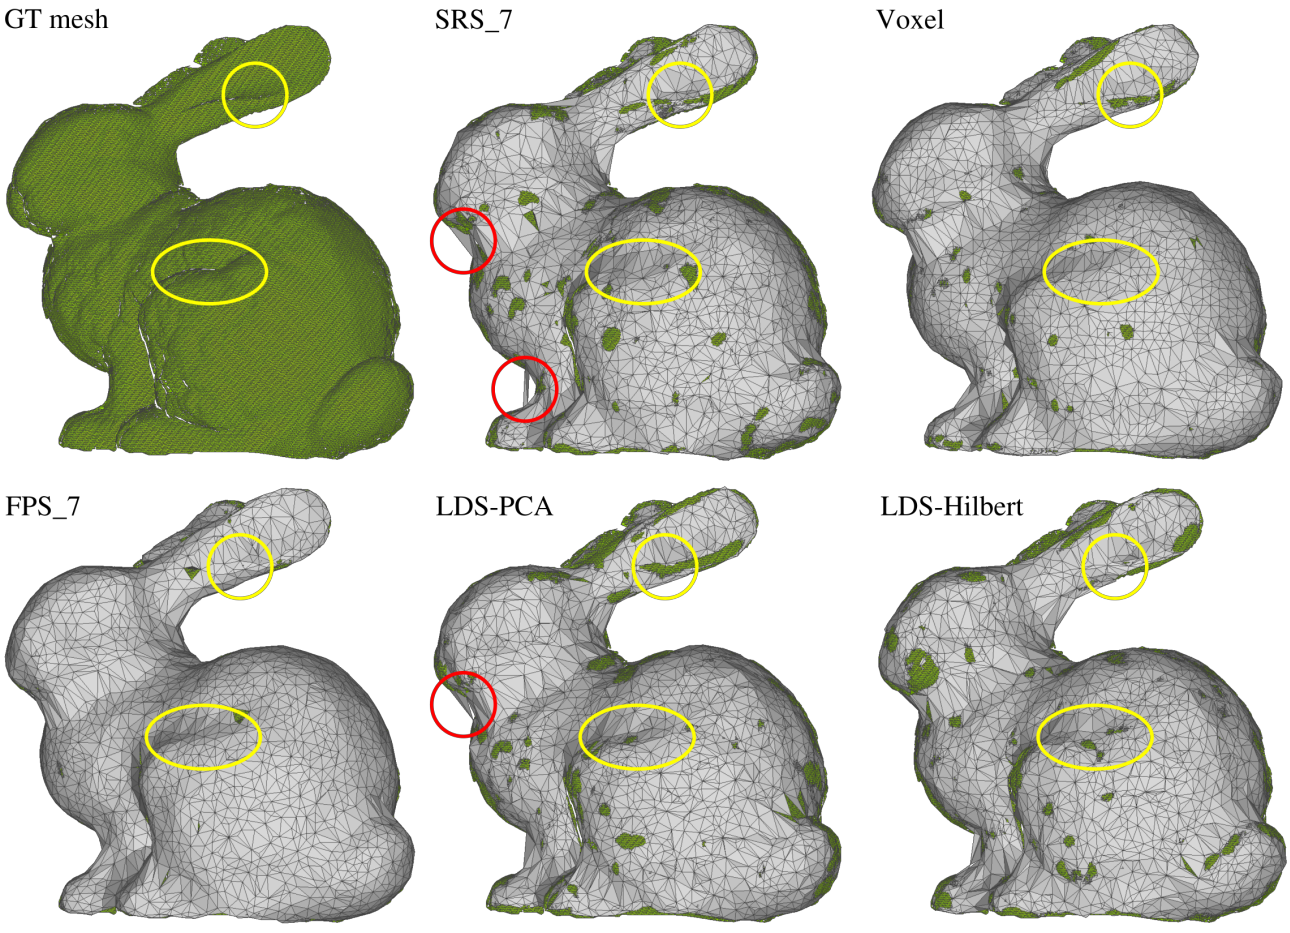


Figure 5. Comparative analysis of bunny model reconstruction. The yellow circles highlight two important detail regions and the red circles indicate significant surface defects.

Table 7. Performance Ratio on Stanford scans fidelity.

Metrics	SRS	FPS	Voxel	LDS-PCA
D_{p2m}	1	1.47	1.16	1
D_{CD}	1.02	1.17	1.02	1.02
$D_{\text{H,95, s2g}}$	1	1.28	1.13	1
$D_{\text{H,95, g2s}}$	1.08	0.87	0.83	1.07
D_{VFH}	2.6	3.19	3.25	2.41

Table 8. Win Percentage on Stanford scans fidelity.

Metrics	SRS	FPS	Voxel	LDS-PCA
D_{p2m}	41.67%	97.92%	85.42%	60.42%
D_{CD}	89.58%	100.00%	43.75%	85.42%
$D_{\text{H,95, s2g}}$	39.58%	100.00%	100.00%	45.83%
$D_{\text{H,95, g2s}}$	100.00%	0.00%	0.00%	97.92%
D_{VFH}	95.83%	97.92%	97.92%	91.67%

4.6.2. Quantitative Results

Tables 7 and 8 present the Performance Ratio and Win Percentage metrics for each evaluation criterion. The Performance Ratio (Table 7) quantifies the relative error of each comparison method to LDS-Hilbert, with values greater than 1.0 indicating LDS-Hilbert's superior performance. The Win Percentage (Table 8) shows the proportion of test cases where LDS-Hilbert outperformed each comparison method.

4.6.3. Key Findings

1. **Superior shape preservation**: The most striking result appears in the VFH Distance (D_{VFH}) metric, where LDS-Hilbert dramatically outperforms all comparison methods with Performance Ratios of 2.60, 3.19, 3.25, and 2.41 against SRS, FPS, Voxel, and LDS-PCA, respectively. These substantial differences indicate that LDS-Hilbert preserves the essential morphological characteristics and overall shape of these complex organic models 2–3 times better than other methods. The Win Percentages for this metric range from 91.67% to 97.92%, demonstrating consistent superiority across nearly all test cases.
2. **Balanced geometric fidelity**: The Point-to-Mesh Distance (D_{p2m}) shows LDS-Hilbert achieving Performance Ratios of 1.00, 1.47, 1.16, and 1.00 against SRS, FPS, Voxel, and LDS-PCA, respectively. While the advantages over SRS and LDS-PCA appear modest when averaged, the Win Percentages of 41.67% and 60.42% reveal that LDS-Hilbert outperforms these methods on many individual scans. Against FPS and Voxel, the advantages are more pronounced, with Win Percentages of 97.92% and 85.42%.
3. **Comprehensive distance metrics**: For the Chamfer Distance (D_{CD}), which considers both how well sampled points represent the surface and how well the surface is covered by sampled points, LDS-Hilbert demonstrates Performance Ratios of 1.02, 1.17, 1.02, and 1.02 against SRS, FPS, Voxel, and LDS-PCA, respectively. The Win Percentages for this metric are 89.58%, 100%, 43.75%, and 85.42%, indicating particularly consistent advantages over FPS and strong performance against SRS and LDS-PCA.
4. **Trade-off in coverage metrics**: For the Hausdorff Distance from ground truth to sample ($D_{\text{H,95,g2s}}$), the FPS and Voxel methods achieve lower error rates than LDS-Hilbert, with Performance Ratios of 0.87 and 0.83, respectively. This reflects these methods' design focus on ensuring uniform coverage at the expense of other quality aspects. The 0% Win Percentage against these methods confirms this is a systematic pattern. However, LDS-Hilbert performs significantly better on the Hausdorff Distance from

sample to ground truth ($D_{H,95,s2g}$), with Performance Ratios of 1.00, 1.28, 1.13, and 1.00 against SRS, FPS, Voxel, and LDS-PCA, respectively.

5. **Value of Hilbert ordering**: The comparison with LDS-PCA isolates the contribution of the Hilbert curve ordering. While LDS-PCA performs reasonably well across most metrics, LDS-Hilbert maintains an advantage in several key areas, particularly in the VFH Distance (Performance Ratio of 2.41) and Win Percentage (91.67%). This confirms that the Hilbert curve's spatial coherence properties specifically enhance shape preservation for complex organic models.

The significance analysis is provided in Section 4.7.

4.6.4. Implications and Synthesis

These results on the Stanford dataset—widely recognized as a benchmark for 3D reconstruction algorithms—provide compelling evidence that the LDS-Hilbert method's advantages extend to real-world scanned data. The method's ability to maintain fine details while preserving overall shape integrity makes it particularly valuable for digital twin applications involving complex, organic geometries with varying feature scales and density distributions.

The Stanford dataset results are especially meaningful because these models represent genuine scanned data with all the complexities and imperfections inherent in real-world acquisition processes. LDS-Hilbert's robust performance on this data suggests it would be well suited for practical applications in cultural heritage preservation, medical modeling, and other fields where high-fidelity digital reproduction of complex organic forms is essential.

The most significant advantage of LDS-Hilbert appears to be in preserving the overall morphological characteristics of the models, as evidenced by the dramatic improvements in the VFH Distance. This suggests that the method is particularly valuable for applications where the visual appearance and geometric features of the digital twin must closely match the physical original.

Combined with our findings from previous experiments, these results establish LDS-Hilbert as a consistently superior approach for point cloud down-sampling across a wide range of geometry types, from simple parametric surfaces to complex real-world scans. The method's advantages become increasingly pronounced as object complexity increases, making it especially relevant for challenging digital twin applications where both computational efficiency and geometric fidelity are critical requirements.

4.7. Statistical Significance

We conducted paired *t*-tests comparing LDS-Hilbert against each alternative method across all experiments. The quantities to be tested are the distance metrics in each experiment. The computed t-values are summarized in Tables 9–12. The corresponding significance levels are also indicated, where * denotes $p < 0.05$ and ** denotes $p < 0.01$.

The paired *t*-tests reveal that LDS-Hilbert's performance advantage becomes more pronounced as the point cloud size N increases, particularly evident in Experiments 1 and 2. This suggests that LDS-Hilbert's ability to maintain geometric fidelity is more effective for larger datasets. In the more complex point clouds of Experiments 3 and 4, LDS-Hilbert consistently outperforms other methods, with particularly notable improvements in Experiment 4, which features intricate surfaces. Across all mesh reconstruction tasks, LDS-Hilbert shows significant reductions in the VFH Distance, indicating superior preservation of overall shape and structural details. Additionally, LDS-Hilbert demonstrates balanced performance in other metrics, such as D_{p2m} and D_{CD}, suggesting its effectiveness in maintaining both local and global fidelity. These results highlight LDS-Hilbert's robust-

ness in handling diverse datasets and its potential for enhancing digital twin fidelity in various applications.

Table 9. Paired *t*-test on D_{L1} in Experiment 1.

N	Ratio	SRS	FPS	Voxel	LDS-PCA
9180	0.1	2.2626	2.3988	2.3613	1.4792
	0.2	1.8992	2.0358	2.3981	2.1543
37,240	0.1	1.4637	2.2552	2.2581	0.1516
	0.2	3.4722 *	2.4084	2.1334	1.4663
95,760	0.1	4.7777 **	2.5374 *	2.6974 *	3.3146 *
	0.2	3.9261 **	2.6270 *	2.4392	4.0564 **

*: $p < 0.05$ and **: $p < 0.01$.

Table 10. Paired *t*-test on Experiment 2.

N	Metric	SRS	FPS	Voxel	LDS-PCA
9180	$D_{\mathrm{H,95,\ s2g}}$	−0.7048	15.2714 **	1.4661	−1.5009
	$D_{\mathrm{H,95,\ g2s}}$	17.4325 **	−13.5211	−9.4246	14.6754 **
	D_{CD}	21.6308 **	5.5173 **	−6.7199	13.4815 **
	D_{p2m}	0.3618	15.6020 **	−4.8813	−0.0572
	D_{VFH}	1.0269	−0.5189	2.0438	1.0983
37,240	$D_{\mathrm{H,95,\ s2g}}$	−1.9041	59.0391 **	13.6449 **	−1.3476
	$D_{\mathrm{H,95,\ g2s}}$	−2.071	−11.0762	−13.8813	−2.7027
	D_{CD}	−2.8076	25.7126 **	−0.9972	−2.7011
	D_{p2m}	−1.2783	82.7670 **	8.7161 **	0.0613
	D_{VFH}	1.6435	0.9146	−0.0662	1.5344
95,760	$D_{\mathrm{H,95,\ s2g}}$	0.6534	116.4929 **	28.0707 **	2.4064 *
	$D_{\mathrm{H,95,\ g2s}}$	20.8601 **	20.9876 **	−3.5097	19.7301 **
	D_{CD}	21.0916 **	91.6320 **	18.7912 **	16.9903 **
	D_{p2m}	0.9368	140.4269 **	23.6502 **	−0.1805
	D_{VFH}	2.9560 *	3.1374 **	2.7974 *	3.1259 **

*: $p < 0.05$ and **: $p < 0.01$.

Table 11. Paired *t*-test on Experiment 3.

Metric	SRS	FPS	Voxel	LDS-PCA
$D_{\mathrm{H,95,\ s2g}}$	−0.4106	25.3279 **	1.3399	−0.219
$D_{\mathrm{H,95,\ g2s}}$	3.9070 **	−1.3709	−1.3842	4.7496 **
D_{CD}	4.0613 **	0.7314	1.2208	4.5332 **
D_{p2m}	0.9432	25.1428 **	7.0922 **	0.7867
D_{VFH}	3.6487 **	2.1050 *	2.4367 *	3.2510 **

*: $p < 0.05$ and **: $p < 0.01$.

Table 12. Paired *t*-test on Experiment 4.

Metric	SRS	FPS	Voxel	LDS-PCA
$D_{\mathrm{H,95,\ s2g}}$	−1.2756	53.5738 **	15.8549 **	−0.5698
$D_{\mathrm{H,95,\ g2s}}$	8.2243 **	−9.69	−11.5884	8.3965 **
D_{CD}	7.6155 **	28.5556 **	−0.0294	7.5771 **
D_{p2m}	−1.9121	3.9056 **	2.3743 *	0.5398
D_{VFH}	13.2551 **	18.6591 **	18.6035 **	9.1917 **

*: $p < 0.05$ and **: $p < 0.01$.

Summary of Computational Efficiency

Beyond fidelity, computational efficiency is a practical consideration for down-sampling methods. A comparative analysis of the average down-sampling time and its coefficient of variation was conducted (details in Appendix E, Table A4, Figure A3). The results show that the proposed LDS-Hilbert method exhibits practical execution times comparable to other established techniques like Farthest Point Sampling (FPS) and the LDS-PCA variant used in our ablation study. While not as computationally inexpensive as Simple Random Sampling (SRS) or optimized Voxel Grid implementations (which benefit from simple random selection or grid hashing, respectively), LDS-Hilbert demonstrates highly stable performance with low variance in execution time across different runs and geometries. This predictability is a notable advantage over the Voxel method, whose timing and output size can be sensitive to the chosen cell size parameter. This balance of consistently high fidelity (demonstrated across Experiments 1–4) and predictable, reasonable computational cost further supports the suitability of LDS-Hilbert for practical digital twin applications where both accuracy and performance stability are valued.

5. Conclusions and Limitations

5.1. Conclusions

This research addresses the challenge of point cloud down-sampling through the development of a novel LDS-Hilbert approach. Our method leverages the complementary mathematical properties of Low-Discrepancy Sequences and Hilbert curve ordering to create a synergistic approach that preserves both global structure and local details.

Across four diverse experiments, LDS-Hilbert consistently outperformed traditional approaches:

1. For parametric fitting (Type A), our method achieved parameter recovery with up to 40% lower error.
2. For mesh reconstruction (Types B and C), it demonstrated superior preservation of critical features while maintaining better overall shape fidelity, with VFH Distances up to 50% better.
3. On real-world laser scans, LDS-Hilbert preserved fine details that other methods smoothed over or distorted.

The integration of Low-Discrepancy Sequences with Hilbert curve ordering represents a significant advancement in point cloud down-sampling technology. The method achieves these improvements without requiring feature-specific calculations, extensive pre-processing, or task-specific training data.

Beyond computational efficiency, this work has significant implications across various domains. It can improve environmental perception for autonomous vehicles, enhance structural assessment in infrastructure digital twins, and potentially boost diagnostic accuracy in medical imaging by better preserving crucial details. Importantly, by maintaining the original density distributions, our method respects the inherent information captured during scanning, often reflecting focused attention on complex or critical regions.

Building upon these findings, we offer practical guidelines for method selection in Section 5.1 and outline several avenues for future research in Section 5.2 to further advance this work.

Practical Guidelines for Method Selection

Based on our experimental results, we offer the following guidelines for practitioners:

1. **For large-scale point clouds (>100 K points)**: LDS-Hilbert consistently offers superior performance across all tested geometry types and is strongly recommended, particularly when down-sampling ratios below 20% are required.
2. **For time-critical applications**: If processing speed is the primary concern, SRS provides the fastest execution but with quality compromises. Voxel Grid Filtering offers a reasonable compromise between speed and quality for simpler geometries without sharp features.
3. **For feature-rich models**: When the point cloud contains important sharp features, edges, or corners, LDS-Hilbert significantly outperforms traditional methods, with differences most pronounced for complex CAD models and real-world scans.
4. **For parametric surface fitting**: When the goal is parameter recovery rather than visual fidelity, the advantage of LDS-Hilbert over FPS and Voxel methods is particularly substantial, with error reductions exceeding 40%.
5. **Implementation considerations**: The LDS-Hilbert method requires minimal parameter tuning beyond specifying the desired output size, making it more robust across different point cloud characteristics compared to methods like FPS (neighborhood size) or Voxel (cell size) that require careful parameter selection.

5.2. Limitations and Future Work

Despite its advantages, several limitations warrant acknowledgment:

1. The current implementation has been validated primarily on 3D point clouds; extension to higher-dimensional feature spaces represents an important area for future investigation.
2. The approach treats all regions with equal importance; future work could explore adaptive variants incorporating importance weighting.
3. Integration with learning-based methods and adaptation for streaming or online processing could enhance applicability.
4. The current study provides average processing times, but a dedicated analysis of real-time performance (e.g., latency, throughput) on large-scale point clouds, crucial for certain applications, remains an area for future investigation.
5. Future work includes establishing formal theoretical bounds on LDS-Hilbert's approximation error. This involves analyzing how Hilbert curve locality and LDS discrepancy jointly determine the sampled subset's representativeness (geometric fidelity and uniformity), offering stronger guarantees for critical applications.

By addressing these limitations, the method's utility can be further expanded to benefit a wider range of applications in fields such as autonomous navigation, industrial inspection, cultural heritage preservation, and medical modeling.

Author Contributions: Conceptualization, Y.M. and L.G.; methodology, Y.M.; validation, M.L.; formal analysis, Y.M.; investigation, M.L.; resources, L.G.; data curation, Y.M.; writing—original draft preparation, Y.M.; writing—review and editing, L.G.; supervision, L.G.; project administration, L.G.; funding acquisition, L.G. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported by the Qilu Young Scholars Program Fund of Shandong University.

Institutional Review Board Statement: Not Applicable.

Informed Consent Statement: Not Applicable.

Data Availability Statement: The key scripts and data supporting the experiments in this paper are openly available in the `LDS-Hilbert-Point-Cloud` repository at https://github.com/Yuening-Ma/LDS-Hilbert-Point-Cloud (accessed on 9 June 2025).

Acknowledgments: The authors thank Miao Zhang for her valuable suggestions on the development of the low-discrepancy sequence-based down-sampling method. During the preparation of this manuscript, the authors used Claude 3.7 for the purposes of proofreading. Claude 3.7 also assisted the authors in the Introduction and Literature Review sections. The authors have thoroughly reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The funders had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A. Illustration of Proposed Method

To demonstrate our method's effectiveness, we present a 2D point cloud example consisting of 1230 points arranged in a circular contour with Gaussian noise ($\sigma = 0.02$) added. Due to the sampling methodology, the point distribution is denser near the Y-axis and sparser away from it.

Figure A1 illustrates the sampling process in LDS-Hilbert. Figure A2 compares our method with benchmark methods (SRS, FPS, Voxel) at a 10% sampling ratio.

Both FPS and Voxel produce spatially uniform results but alter the original density variations. The LDS-Hilbert method preserves the original spatial distribution, including density differences. While SRS generally adheres to the true distribution, its stochastic nature can lead to local gaps or clusters.

For a more precise quantitative comparison of how well each down-sampling method preserves the point cloud's morphology, we calculated the following metrics comparing the down-sampled cloud $\mathcal{P}'$ to the ground truth cloud $\mathcal{P}_{\text{true}}$: $D_{\text{pc}}(\mathcal{P}', \mathcal{P}_{\text{true}})$ (average distance from sampled points to the nearest true point), $D_{\text{pc}}(\mathcal{P}_{\text{true}}, \mathcal{P}')$ (average distance from true points to the nearest sampled point), and the Chamfer Distance $D_{\text{Chamfer}}(\mathcal{P}_{\text{true}}, \mathcal{P}')$ (the sum of the previous two metrics). As Viewpoint Feature Histograms (VFH) are not applicable to 2D data, we used the Kullback–Leibler (KL) divergence between distribution histograms of $\mathcal{P}'$ and $\mathcal{P}_{\text{true}}$ to compare the preservation of the overall point cloud structure. Table A1 presents these results, in which a ratio >1 means LDS-Hilbert achieves better performance.

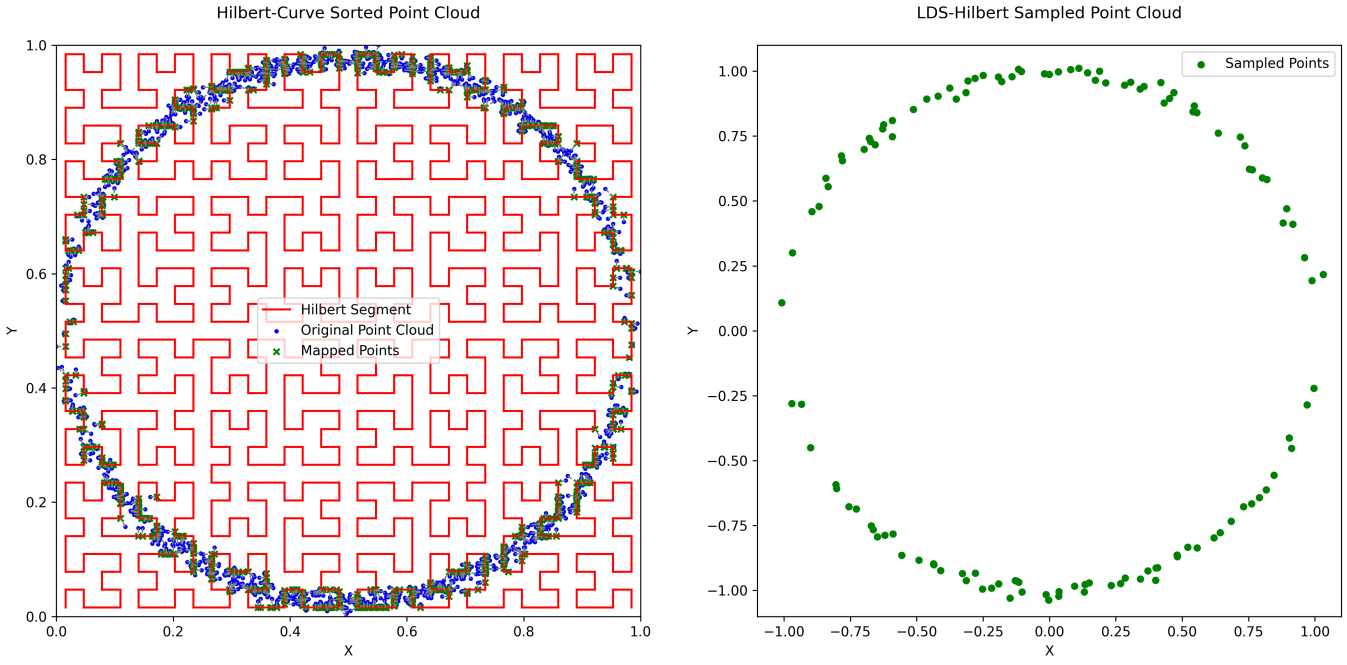


Figure A1. Visualization of the down-sampled point cloud using LDS-Hilbert.

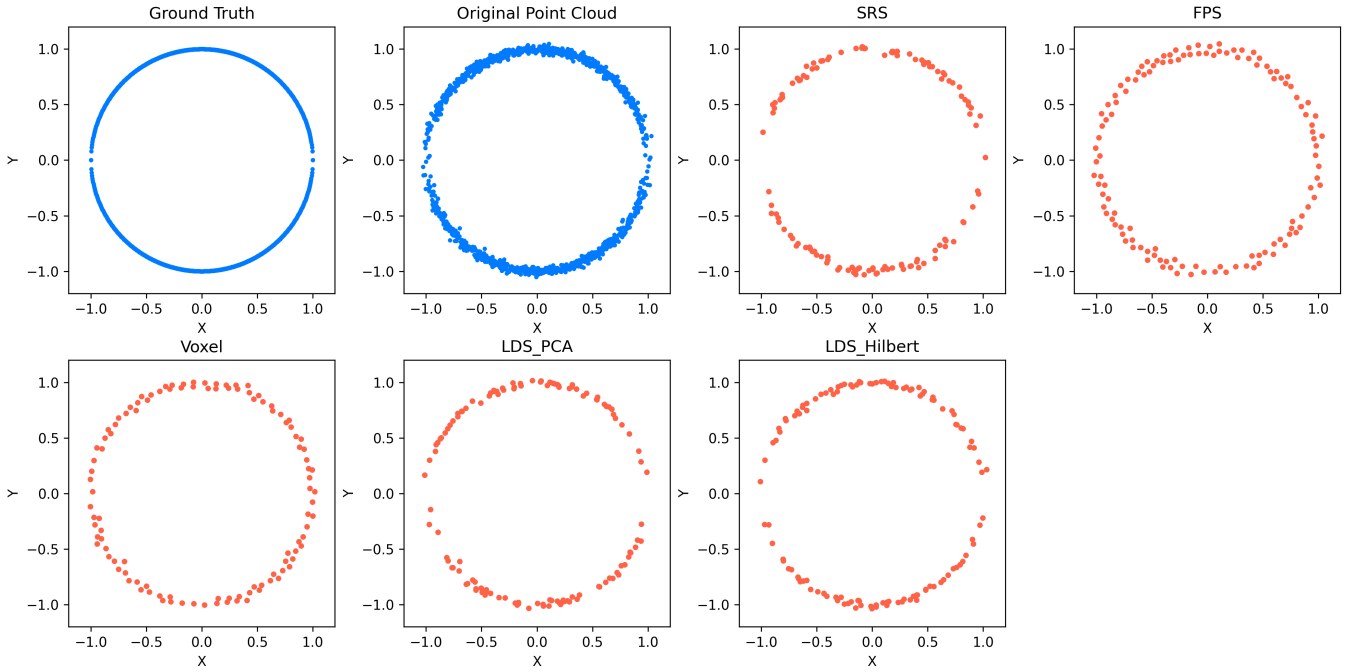


Figure A2. Comparison of down-sampling results using different methods. The point cloud down-sampled by the LDS-Hilbert method better preserves the distributions of both the ground truth and the original point cloud.

Table A1. Performance Ratio of down-sampling methods on the 2D example using fidelity metrics.

Metrics	SRS	FPS	Voxel	LDS-PCA
$D_{\text{pc}}(\mathcal{P}', \mathcal{P}_{\text{true}})$	1.2	1.99	1.26	0.98
$D_{\text{pc}}(\mathcal{P}_{\text{true}}, \mathcal{P}')$	1.23	1.23	1	1.12
D_{CD}	1.22	1.52	1.1	1.07
$D_{\text{KL}}(\cdot \Vert \cdot)$	1.86	2.39	2.64	2.11

Appendix B. Dataset Generation and Parameter Fitting in Experiment 1

Point clouds in Experiment 1 were generated from seven analytical surfaces with specific parameters in Table A2.

Gaussian noise ($\sigma = 0.05$ relative to object scale) was added to simulate measurement errors.

The parametric surface fitting is with least-squares optimization. The optimization objectives are summarized in Table A2 as well.

Plane fitting is implemented using the least squares API in NumPy, while sphere, ellipsoid, torus, cylinder, cone, and paraboloid fitting are implemented using the corresponding API in the SciPy library..

Table A2. Surface types, parameters, and optimization objectives.

No.	Surface Type	Model Equation	Parameters	Optimization Objective
1	Plane	$\begin{cases} z = ax + by + c \\ x \in [-2, 2], y \in [-2, 2] \end{cases}$	$\begin{cases} a = 1 \\ b = 0.5 \\ c = 1.5 \end{cases}$	$z - (ax + by + c) = 0$
2	Sphere	$\begin{cases} x = x_c + r \sin\phi \cos w \\ y = y_c + r \sin\phi \sin w \\ z = z_c + r \cos\phi \\ w \in [0, 2\pi], \phi \in [0, \pi] \end{cases}$	$\begin{cases} x_c = 0 \\ y_c = 0 \\ z_c = 0 \\ r = 1 \end{cases}$	$\sqrt{(x - x_c)^2 + (y - y_c)^2 + (z - z_c)^2} - r = 0$

Table A2. *Cont.*

No.	Surface Type	Model Equation	Parameters	Optimization Objective
3	Ellipsoid	$\begin{cases} x = x_c + a\sin\phi\cos w \\ y = y_c + b\sin\phi\sin w \\ z = z_c + c\cos\phi \\ w \in [0, 2\pi], \phi \in [0, \pi] \end{cases}$	$\begin{cases} x_c = 0 \\ y_c = 0 \\ z_c = 0 \\ a = 3 \\ b = 2 \\ c = 1 \end{cases}$	$\frac{(x-x_c)^2}{a^2} + \frac{(y-y_c)^2}{b^2} + \frac{(z-z_c)^2}{c^2} - 1 = 0$
4	Torus	$\begin{cases} x = x_c + (R + r\cos\phi)\cos w \\ y = y_c + (R + r\cos\phi)\sin w \\ z = z_c + r\sin\phi \\ w \in [0, 2\pi], \phi \in [0, 2\pi] \end{cases}$	$\begin{cases} x_c = 0 \\ y_c = 0 \\ z_c = 0 \\ R = 2 \\ r = 0.5 \end{cases}$	$\sqrt{(\sqrt{(x-x_c)^2 + (y-y_c)^2} - R)^2 + (z-z_c)^2} - r = 0$
5	Cylinder	$\begin{cases} x = x_c + r\cos w \\ y = y_c + r\sin w \\ z \in [-2, 2] \\ w \in [0, 2\pi] \end{cases}$	$\begin{cases} x_c = 0 \\ y_c = 0 \\ r = 1 \end{cases}$	$\sqrt{(x-x_c)^2 + (y-y_c)^2} - r = 0$
6	Cone	$\begin{cases} x = x_c + rz\cos w \\ y = y_c + rz\sin w \\ z \in [0, 2] \\ w \in [0, 2\pi] \end{cases}$	$\begin{cases} x_c = 0 \\ y_c = 0 \\ z_c = 0 \\ r = 1 \end{cases}$	$\sqrt{(x-x_c)^2 + (y-y_c)^2} - r = 0$
7	Paraboloid	$\begin{cases} x = x_c + r\cos w \\ y = y_c + r\sin w \\ z = z_c + ar^2 \\ r \in [0, 2], w \in [0, 2\pi] \end{cases}$	$\begin{cases} x_c = 0 \\ y_c = 0 \\ z_c = 0 \\ a = 0.5 \end{cases}$	$z - [a(x-x_c)^2 + a(y-y_c)^2 + z_c] = 0$

Appendix C. Dataset Generation for Closed Geometries

Experiment 2 used 15 geometric shapes with dimensions defined relative to a base scale factor (SCALE = 100), whose dimensions are shown in Table A3.

Each shape was randomly rotated before sampling points using Poisson disk sampling. Gaussian noise ($\sigma = 1$) was added to simulate scanning errors.

Table A3. Geometric shapes and their parameters. Note that the numerical values need to be multiplied by a base scale factor (SCALE = 100).

ID	Shape	Param_1	Param_2	Param_3
1	Sphere	Radius $r = 0.5$	-	-
2	Cylinder	Radius $r = 0.3$	Height $h = 1$	-
3	Equilateral Triangular Prism	Base edge length $a = 0.6$	Height $h = 1$	-
4	Cube	Edge length $a = 1$	-	-
5	Rectangular Prism	Length $l = 0.4$	Width $w = 0.8$	Height $h = 1$
6	Pentagonal Prism	Base edge length $a = 0.4$	Height $h = 1$	-
7	Hexagonal Prism	Base edge length $a = 0.4$	Height $h = 1$	-
8	L-beam	Width $w = 0.6$	Height $h = 1$	Thickness $t = 0.1$
9	T-beam	Width $w = 0.6$	Height $h = 1$	Thickness $t = 0.1$
10	U-beam	Width $w = 0.6$	Height $h = 1$	Thickness $t = 0.1$
11	H-beam	Width $w = 0.6$	Height $h = 1$	Thickness $t = 0.1$
12	Cone	Radius $r = 0.5$	Height $h = 1$	-
13	Torus	Major radius $R = 0.35$	Minor radius $r = 0.15$	-
14	Tetrahedron	Circumscribed radius $r = 0.75$	-	-
15	Octahedron	Circumscribed radius $r = 0.5$	-	-

Appendix D. Parameter Settings for Ball Pivoting

In Experiments 2, 3, and 4, we utilized the Ball Pivoting algorithm for mesh reconstruction from point clouds. This algorithm constructs meshes by rolling a ball over the point cloud, with the radius of the ball influencing the scale of detected geometric features. To ensure consistent and effective mesh reconstruction across different point cloud densities, we dynamically adjusted the ball radius based on the average point spacing of the input data.

Specifically, we first calculated a base radius d that is inversely proportional to the average point spacing. Given that the point clouds represent hollow surface models with uniform heights and similar total surface areas, the average point spacing is inversely related to the square root of the number of points N. Based on this relationship, we estimated the base radius d using the following formula:

$$d = \text{round}\left(\sqrt{\frac{27{,}000}{N}}, 1\right)$$

where N is the number of points in the point cloud. The constant 27,000 was determined through preliminary experiments to ensure that point clouds down-sampled by different methods could all generate a sufficient number of triangles, thereby guaranteeing the fairness and effectiveness of the experiments.

For Experiment 2, the radii list used was $2d, 3d, 4d, 5d, 6d$. For Experiments 3 and 4, the radii list was adjusted $2d, 2.25d, 2.5d, 2.75d, 3d$.

These radii lists were passed to the ball-pivoting API in the Open3d library, which attempts to create triangles using each radius in ascending order. This approach allows the algorithm to capture geometric features at different scales, ensuring accurate and complete mesh reconstruction.

By employing this dynamic radius adjustment strategy, we effectively balanced the detail retention and computational efficiency of mesh reconstruction, providing high-quality mesh models for subsequent digital twin creation.

Appendix E. Time Consuming Analysis

To validate the efficiency of the LDS-Hilbert method, a down-sampling time analysis was conducted on the point clouds generated in Experiment 2. Specifically, the average down-sampling time and the coefficient of variation (i.e., the ratio of the standard deviation to the mean) for each method were calculated across 15 experimental objects with the same number of points. Table A4 presents the average time results, and Figure A3 shows the heatmap of the coefficients of variation.

From Table A4 and Figure A3, we can observe the following:

1. In terms of time consumption, SRS and Voxel are the most efficient methods, while FPS, LDS-PCA, and LDS-Hilbert have similar time costs. Among them, LDS-Hilbert is slightly slower, but the difference is not significant.
2. Except for the Voxel method, all down-sampling methods, including the proposed LDS-Hilbert, have low coefficients of variation in time, indicating stable efficiency. One reason for the instability of the Voxel method is that the number of output points from down-sampling is difficult to determine, often requiring adjustments to the Voxel size to achieve an output point count close to the desired number. This is also one of the drawbacks of the Voxel method.

Since the Voxel, FPS, and LDS-Hilbert algorithms all rely on well-established algorithm libraries in the industry, their time costs are influenced by the level of algorithm optimization. Therefore, this experiment is for reference only, to demonstrate that the proposed

LDS-Hilbert method can achieve efficiency and stability comparable to commonly used industry algorithms.

Table A4. Average down-sampling time on Experiment 2 data.

N	SRS	FPS	Voxel	LDS-PCA	LDS-Hilbert
9180	0.000302	0.017179	0.002178	0.008097	0.022676
37,240	0.001111	0.056052	0.003589	0.041993	0.067687
95,760	0.003005	0.159396	0.006999	0.120169	0.197373

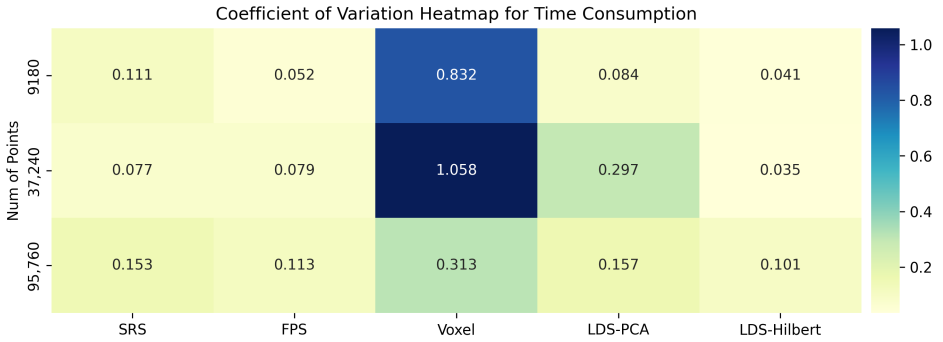


Figure A3. Coefficient of variation heatmap for time consumption.

References

1. Grieves, M. Digital twin: Manufacturing excellence through virtual factory replication. *White Pap.* **2014**, *1*, 1–7.
2. Tao, F.; Cheng, J.; Qi, Q.; Zhang, M.; Zhang, H.; Sui, F. Digital twin-driven product design, manufacturing and service with big data. *Int. J. Adv. Manuf. Technol.* **2018**, *94*, 3563–3576. [CrossRef]
3. Li, D.; Wei, Y.; Zhu, R. A comparative study on point cloud down-sampling strategies for deep learning-based crop organ segmentation. *Plant Methods* **2023**, *19*, 124. [CrossRef]
4. Lyu, W.; Ke, W.; Sheng, H.; Ma, X.; Zhang, H. Dynamic Downsampling Algorithm for 3D Point Cloud Map Based on Voxel Filtering. *Appl. Sci.* **2024**, *14*, 3160. [CrossRef]
5. Liu, J.; Li, J.; Wang, K.; Guo, H.; Yang, J.; Peng, J.; Xu, K.; Liu, X.; Guo, J. LTA-PCS: Learnable Task-Agnostic Point Cloud Sampling. In Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 16–22 June 2024; pp. 28035–28045. [CrossRef]
6. Mchirgui, N.; Quadar, N.; Kraiem, H.; Lakhssassi, A. The Applications and Challenges of Digital Twin Technology in Smart Grids: A Comprehensive Review. *Appl. Sci.* **2024**, *14*, 10933. [CrossRef]
7. Xu, M. BDSR: A Best Discrepancy Sequence-Based Point Cloud Resampling Framework for Geometric Digital Twinning. Master's Thesis, Shandong University, Jinan, China, 2021.
8. Chen, W.; Zhu, X.; Chen, G.; Yu, B. Efficient Point Cloud Analysis Using Hilbert Curve. In Proceedings of the Computer Vision—ECCV 2022, Tel Aviv, Israel, 23–27 October 2022; Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T., Eds.; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2022; Volume 13662, pp. 730–747. [CrossRef]
9. Chen, X.; Wu, Y.; Xu, W.; Li, J.; Dong, H.; Chen, Y. PointSCNet: Point Cloud Structure and Correlation Learning Based on Space-Filling Curve-Guided Sampling. *Symmetry* **2021**, *14*, 8. [CrossRef]
10. Hu, A.; Xu, K.; Yin, X.; Wang, D. LiDAR Point Cloud Simplification Strategy Utilizing Probabilistic Membership. *Front. Phys.* **2024**, *12*, 1471077. [CrossRef]
11. Chen, L.; Feng, C.; Ma, Y.; Wang, C. A review of rigid point cloud registration based on deep learning. *Front. Neurorobot.* **2024**, *17*, 1281332. [CrossRef]
12. Liu, X.; Zhang, Z. Effects of LiDAR Data Reduction and Breaklines on the Accuracy of Digital Elevation Model. *Surv. Rev.* **2011**, *43*, 614–628. [CrossRef]
13. Al-Rawabdeh, A.; He, F.; Habib, A. Automated Feature-Based Down-Sampling Approaches for Fine Registration of Irregular Point Clouds. *Remote Sens.* **2020**, *12*, 1224. [CrossRef]
14. Rusu, R.B.; Cousins, S. 3D is here: Point Cloud Library (PCL). In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China, 9–13 May 2011; pp. 1–4.
15. Schnabel, R.; Klein, R. Octree-based Point-Cloud Compression. In Proceedings of the Symposium on Point-Based Graphics, Boston, MA, USA, 29–30 June 2006; Botsch, M., Chen, B., Pauly, M., Zwicker, M., Eds.; The Eurographics Association: Eindhoven, The Netherlands, 2006. [CrossRef]

16. Meagher, D. Geometric modeling using octree encoding. *Comput. Graph. Image Process.* **1982**, *19*, 129–147. [CrossRef]
17. Hornung, A.; Wurm, K.M.; Bennewitz, M.; Stachniss, C.; Burgard, W. OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Auton. Robot.* **2013**, *34*, 189–206. [CrossRef]
18. Ester, M.; Kriegel, H.P.; Sander, J.; Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96), Oregon, Portland, 2–4 August 1996; pp. 226–231.
19. Pauly, M.; Gross, M.; Kobbelt, L.P. Efficient Simplification of Point-Sampled Surfaces. In Proceedings of the Conference on Visualization '02, Boston, MA, USA, 27 October–1 November 2002; pp. 163–170.
20. Qi, L.; Xu, S.; Xiao, J.; Wang, Y. Sharp feature-preserving 3d mesh reconstruction from point clouds based on primitive detection. *Remote Sens.* **2023**, *15*, 3155. [CrossRef]
21. Chen, S.; Wang, J.; Pan, W.; Gao, S.; Wang, M.; Lu, X. Towards Uniform Point Distribution in Feature-preserving Point Cloud Filtering. *Comput. Vis. Media* **2023**, *9*, 249–263. [CrossRef]
22. Pathak, S.; Baldwin-McDonald, T.; Sels, S.; Penne, R. GP-PCS: One-shot Feature-Preserving Point Cloud Simplification with Gaussian Processes on Riemannian Manifolds. In Proceedings of the Pattern Recognition, Chongqing, China, 13–16 June 2025; Antonacopoulos, A., Chaudhuri, S., Chellappa, R., Liu, C.L., Bhattacharya, S., Pal, U., Eds.; Springer: Cham, Switzerland, 2025; pp. 436–452.
23. Matsuzaki, K.; Nonaka, K. Point Cloud Sampling Preserving Local Geometry for Surface Reconstruction. In Proceedings of the 34th British Machine Vision Conference 2023, BMVC 2023, Aberdeen, UK, 20–24 November 2023.
24. Zhang, G.; Zhao, W.; Liu, J.; Liu, X. REPS: Reconstruction-based Point Cloud Sampling. *arXiv* **2024**, arXiv:2403.05047.
25. Zhang, R.; Wu, Y.; Jin, W.; Meng, X. Deep-Learning-Based Point Cloud Semantic Segmentation: A Survey. *Electronics* **2023**, *12*, 3642. [CrossRef]
26. Sugimoto, R.; Kanai, K.; Maki, S.; Katto, J. A Point Cloud Downsampling Method Balancing Global and Local Shapes for 3D Object Classification. In Proceedings of the 2025 IEEE International Conference on Consumer Electronics, Las Vegas, NV, USA, 11–14 January 2025; pp. 1–6. [CrossRef]
27. Deng, C.; Peng, Z.; Chen, Z.; Chen, R. Point cloud deep learning network based on balanced sampling and hybrid pooling. *Sensors* **2023**, *23*, 981. [CrossRef]
28. Niederreiter, H. *Random Number Generation and Quasi-Monte Carlo Methods*; CBMS-NSF Regional Conference Series in Applied Mathematics; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 1992; Volume 63. [CrossRef]
29. Glasserman, P. *Monte Carlo Methods in Financial Engineering*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2003; Volume 53.
30. Jäckel, P. *Monte Carlo Methods in Finance*; John Wiley & Sons: Hoboken, NJ, USA, 2002.
31. Kuipers, L.; Niederreiter, H. *Uniform Distribution of Sequences*; Wiley-Interscience: Hoboken, NJ, USA, 1974.
32. Halton, J.H. On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.* **1960**, *2*, 84–90. [CrossRef]
33. Sobol', I. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Comput. Math. Math. Phys.* **1967**, *7*, 86–112. [CrossRef]
34. Niederreiter, H. Low-discrepancy and low-dispersion sequences. *J. Number Theory* **1988**, *30*, 51–70. [CrossRef]
35. Sloan, I.H.; Woźniakowski, H. When Are Quasi-Monte Carlo Algorithms Efficient for High Dimensional Integrals? *J. Complex.* **1998**, *14*, 1–33. [CrossRef]
36. Shirley, P.; Edwards, D.; Boulos, S. Monte Carlo and Quasi-Monte Carlo Methods for Computer Graphics. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*; Springer: Berlin/Heidelberg, Germany, 2008; pp. 167–177. [CrossRef]
37. Pharr, M.; Jakob, W.; Humphreys, G. *Physically Based Rendering: From Theory to Implementation*, 3rd ed.; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2016.
38. Yan, D.; Guo, J.; Wang, B.; Zhang, X.; Wonka, P. A Survey of Blue-Noise Sampling and Its Applications. *J. Comput. Sci. Technol.* **2015**, *30*, 439–452. [CrossRef]
39. Fang, K.T. Some Applications of Quasi-Monte Carlo Methods in Statistics. In Proceedings of the Monte Carlo and Quasi-Monte Carlo Methods 2000, Hong Kong, China, 27 November–1 December 2000; pp. 10–26.
40. Morokoff, W.J.; Caflisch, R.E. Quasi-Monte Carlo integration. *J. Comput. Phys.* **1995**, *122*, 218–230. [CrossRef]
41. Dick, J.; Pillichshammer, F. *Digital Nets and Sequences: Discrepancy Theory and Quasi-Monte Carlo Integration*; Cambridge University Press: Cambridge, UK, 2010.
42. Zhang, M.; Li, M.; Guo, L.; Liu, J. A low-cost AI-empowered stethoscope and a lightweight model for detecting cardiac and respiratory diseases from lung and heart auscultation sounds. *Sensors* **2023**, *23*, 2591. [CrossRef]
43. Samet, H. *The Design and Analysis of Spatial Data Structures*; Addison-Wesley: New York, NY, USA, 1990.
44. Bader, M. *Space-Filling Curves: An Introduction with Applications in Scientific Computing*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2012.
45. Hilbert, D. Ueber die stetige Abbildung einer Linie auf ein Flächenstück. *Math. Ann.* **1891**, *38*, 459–460. [CrossRef]

46. Moon, B.; Jagadish, H.; Faloutsos, C.; Saltz, J.H. Analysis of the clustering properties of the Hilbert space-filling curve. *IEEE Trans. Knowl. Data Eng.* **2001**, *13*, 124–141. [CrossRef]
47. Kamel, I.; Faloutsos, C. Hilbert R-tree: An Improved R-tree using Fractals. In Proceedings of the 20th International Conference on Very Large Data Bases, San Francisco, CA, USA, 12–15 September 1994; VLDB '94, pp. 500–509.
48. Pandhare, A.A. Point2Point: A Framework for Efficient Deep Learning on Hilbert sorted Point Clouds with applications in Spatio-Temporal Occupancy Prediction. *arXiv* **2023**, arXiv:2306.16306.
49. Chen, J.; Yu, L.; Wang, W. Hilbert Space Filling Curve Based Scan-Order for Point Cloud Attribute Compression. *IEEE Trans. Image Process.* **2022**, *31*, 4609–4621. [CrossRef]
50. Gotsman, C.; Lindenbaum, M. On the metric properties of discrete space-filling curves. *IEEE Trans. Image Process.* **1996**, *5*, 794–797. [CrossRef]
51. Skilling, J. Programming the Hilbert curve. *Aip Conf. Proc.* **2004**, *707*, 381–387. [CrossRef]
52. Hua, L.K.; Wang, Y. *Applications of Number Theory to Numerical Analysis*; Springer: Berlin/Heidelberg, Germany, 1981. [CrossRef]
53. Guo, L.; Liu, J.; Lu, R. Subsampling bias and the best-discrepancy systematic cross validation. *Sci. China Math.* **2021**, *64*, 197–210. [CrossRef]
54. Cignoni, P.; Scopigno, R. METRO: Measuring error on simplified surfaces. *Comput. Graph. Forum* **1998**, *17*, 167–174. [CrossRef]
55. Barrow, H.G.; Tenenbaum, J.M.; Bolles, R.C.; Wolf, H.C. Parametric Correspondence and Chamfer Matching: Two New Techniques for Image Matching. In Proceedings of the International Joint Conference on Artificial Intelligence, San Francisco, CA, USA, 22–25 August 1977.
56. Fan, H.; Su, H.; Guibas, L. A Point Set Generation Network for 3D Object Reconstruction from a Single Image. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 2463–2471.
57. Huttenlocher, D.P.; Klanderman, G.A.; Rucklidge, W.J. Comparing Images Using the Hausdorff Distance. *IEEE Trans. Pattern Anal. Mach. Intell.* **1993**, *15*, 850–863. [CrossRef]
58. Taha, A.A.; Hanbury, A. Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool. *BMC Med. Imaging* **2015**, *15*, 29. [CrossRef]
59. Rusu, R.B.; Blodow, N.; Beetz, M. Fast Point Feature Histograms (FPFH) for 3D registration. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Kobe, Japan, 12–17 May 2009; pp. 3212–3217.
60. Rusu, R.B.; Bradski, G.R.; Thibaux, R.; Hsu, J. Fast 3D Recognition and Pose Using the Viewpoint Feature Histogram. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Taipei, Taiwan, 18–22 October 2010; pp. 2155–2162.
61. Zhou, Q.Y.; Park, J.; Koltun, V. Open3D: A Modern Library for 3D Data Processing. *arXiv* **2018**, arXiv:1801.09847.
62. Harris, C.R.; Millman, K.J.; van der Walt, S.J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N.J.; et al. Array programming with NumPy. *Nature* **2020**, *585*, 357–362. [CrossRef] [PubMed]
63. Dalinky, L. fpsample: Python Efficient Farthest Point Sampling (FPS) Library. 2023. Available online: https://github.com/leonardodalinky/fpsample (accessed on 17 January 2025).
64. Princeton Laboratory for Intelligent Probabilistic Systems. Numpy-Hilbert-Curve: Numpy Implementation of Hilbert Curves in Arbitrary Dimensions. 2023. Available online: https://github.com/PrincetonLIPS/numpy-hilbert-curve (accessed on 9 February 2025).
65. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
66. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.