

LiteGE: Lightweight Geodesic Embedding for Efficient Geodesics Computation and Non-Isometric Shape Correspondence

Yohanes Yudhi Adikusuma¹, Qixing Huang¹, Ying He²

¹University of Texas at Austin

²Nanyang Technological University

quilin111@utexas.edu, huangqx@cs.utexas.edu, yhe@ntu.edu.sg

Abstract

Computing geodesic distances on 3D surfaces is fundamental to many tasks in 3D vision and geometry processing, with deep connections to tasks such as shape correspondence. Recent learning-based methods achieve strong performance but rely on large 3D backbones, leading to high memory usage and latency, which limit their use in interactive or resource-constrained settings. We introduce *LiteGE*, a lightweight approach that constructs compact, category-aware shape descriptors by applying PCA to unsigned distance field (UDFs) samples at informative voxels. This descriptor is efficient to compute and removes the need for high-capacity networks. LiteGE remains robust on sparse point clouds, supporting inputs with as few as 300 points, where prior methods fail. Extensive experiments show that LiteGE reduces memory usage and inference time by up to 300× compared to existing neural approaches. In addition, by exploiting the intrinsic relationship between geodesic distance and shape correspondence, LiteGE enables fast and accurate shape matching. Our method achieves up to 1000× speedup over state-of-the-art mesh-based approaches while maintaining comparable accuracy on non-isometric shape pairs, including evaluations on point-cloud inputs.

Code — <https://github.com/yya-111/LiteGE>

Introduction

Geodesic distances are a fundamental property of 3D shapes, with broad applications in deep 3D learning and digital geometry processing, such as shape matching (Eisenberger et al. 2021), surface reconstruction (Sprynski et al. 2008), parametrization (Lee, Tong, and Desbrun 2005), texture mapping (Sun et al. 2013), segmentation (Peyré et al. 2010), and building 3D convolutional networks (Masci et al. 2015). The problem of computing geodesic distances and paths on meshes has been extensively studied for decades (Mitchell, Mount, and Papadimitriou 1987), leading to a variety of methods, including exact algorithms (Mitchell, Mount, and Papadimitriou 1987; Chen and Han 1990; Surazhsky et al. 2005; Xin and Wang 2009; Ying, Xin, and He 2014; Xu et al. 2015; Qin et al. 2016) and approximate ones (Sethian 1996; Weber et al. 2008; Xin, Ying, and He 2012; Crane,

Weischedel, and Wardetzky 2013; Ying, Wang, and He 2013; Adikusuma, Fang, and He 2020; Meng et al. 2022; Adikusuma et al. 2022). Despite their effectiveness, these classical methods typically incur substantial computational costs, often taking minutes or even hours to compute geodesics on large-scale models.

Recently, deep learning-based approaches, such as NeuroGF (Zhang et al. 2023) and GeGNN (Pang et al. 2023), have demonstrated substantial efficiency improvements in geodesic computation. These methods use neural networks trained on collections of shapes to learn geodesic distances. Given a shape and query points, the network processes them to generate high-dimensional embeddings from which a lightweight multilayer perceptron (MLP) predicts the geodesic distance within milliseconds. However, the effectiveness of these methods still depends heavily on large backbone network architectures to process shapes that consume a huge amount of GPU memory. This limits their practicality in memory-constrained or interactive applications, particularly those that require geodesics in multiple shapes to be inferred at once. As such, we propose a lightweight and memory-efficient solution that supports interactive computation. Our approach is inspired by the embedding-based formulations used in NeuroGF and GeGNN. Rather than relying on large 3D neural networks, we canonicalize shapes efficiently, construct compact embeddings from them using PCA (Abdi and Williams 2010) on the unsigned distance field (UDF) of the shapes (Chibane, Pons-Moll et al. 2020), and learn geodesic distances using the embeddings. Our network, called LiteGE, preserves the generalization capability of existing methods for shapes from several categories while significantly reducing computational overhead for more than 300× in memory consumption and runtime.

We further demonstrate the utility of LiteGE in addressing the challenging task of non-isometric shape matching, where significant geometric and topological differences complicate correspondence estimation. Although previous work, such as NeuroMorph (Eisenberger et al. 2021), incorporates geodesic distances primarily as a regularization term that is not used on non-isometric shapes, our approach uses geodesic distances directly as the primary supervisory signal. Specifically, given a query point on one shape (e.g., a dog), our network identifies the corresponding point on another shape (e.g., a cat) by predicting and minimizing

geodesic distances in a coarse-to-fine manner.

Shape matching in non-isometric settings is particularly challenging due to significant geometric and topological differences between shapes, especially when we need to deal with point clouds. Our approach trains a network to predict the geodesic distance between one point on the shape X and one point on the shape Y . During training, the network is supervised with intra-shape geodesic distance, which is the geodesic distance between the two points when both lie on the same shape. At inference, a query point on X is matched to the point on Y that minimizes the predicted geodesic distance. This matching procedure is performed in a coarse-to-fine manner for efficiency.

Due to the efficiency of our compact UDF representation, LiteGE supports rapid inference of correspondences for individual query points, making it particularly suitable for local region matching or sparse landmark identification on high-resolution meshes. Experiments conducted on the SMAL dataset of 4-legged animals demonstrate that LiteGE achieves an accuracy comparable to the state-of-the-art Spectral-Meets-Spatial (SMS) (Cao et al. 2024), which is a mesh-based method, while providing inference speeds that are 100 to 1000 \times faster. Notably, thanks to our novel shapes canonicalization pipeline, we can generalize the efficiency and accuracy of our method from meshes to point clouds.

Our main contributions are summarized as follows.

- We propose LiteGE, a lightweight geodesic prediction framework that uses PCA on unsigned distance field representations to produce compact shape embeddings.
- Our method eliminates the need for memory-intensive 3D networks, enabling interactive and scalable geodesic computation across point clouds and meshes.
- We show that LiteGE supports geodesic prediction on very sparse point clouds, outperforming existing methods in low-sample regimes.
- We demonstrate that LiteGE can be applied to non-isometric shape matching for meshes and point clouds with an accuracy similar to that of state-of-the-art mesh-based methods while being up to 1000 \times times faster.

Related Works

Geodesic Computation Traditional methods for computing geodesic distances on 3D surfaces have a long history, spanning computational geometry (Mitchell, Mount, and Papadimitriou 1987; Chen and Han 1990; Xin, He, and Fu 2012; Sharp and Crane 2020), partial differential equations (Sethian and Vladimirsky 2000; Belyaev and Fayolle 2015; Crane, Weischedel, and Wardetzky 2013), graph-based approaches (Ying, Wang, and He 2013; Adikusuma, Fang, and He 2020), and optimization-based formulations (Liu et al. 2017; Yuan et al. 2023). In contrast, deep learning-based geodesic computation is a relatively new research direction. Recent methods, such as NeuroGF (Zhang et al. 2023) and GeGNN (Pang et al. 2023), employ neural networks to predict geodesic distances by embedding shapes and query points in high-dimensional spaces. NeuroGF constructs embeddings by concatenating global shape descriptors—computed using large 3D backbone networks like

PointTransformer (Zhao et al. 2021) or MeshCNN (Hanocka et al. 2019)—with point coordinates, while GeGNN explicitly designs 3D graph neural networks specialized for geodesic tasks. However, these neural approaches typically require large memory footprints (often several gigabytes of GPU memory) even for a single geodesic query in one shape, significantly limiting their practical deployment. In this work, we observe that the shape descriptors used by NeuroGF can be constructed far more efficiently without heavy neural backbones, especially when we have a constrained set of shape categories (e.g., humans, animals, and faces). Specifically, we propose using PCA on UDF representations as a lightweight alternative. By canonicalizing shapes to a standard orientation, position, and scale, and selecting high-variance voxels near surfaces, we efficiently compute UDF-based PCA embeddings while retaining their expressiveness.

Shape Matching Shape matching aims to establish correspondences between points in pairs of shapes and is frequently tackled through dense correspondence frameworks such as functional maps (Ovsjanikov et al. 2012; Litany et al. 2017). Recent approaches increasingly integrate interpolation and correspondence tasks into unified neural architectures (Eisenberger et al. 2021; Cao, Roetzer, and Bernard 2023; Cao et al. 2024). However, many current methods struggle with non-isometric shape pairs, leading to increased matching errors due to substantial geometric and topological variations. Surprisingly, despite the intrinsic relationship between geodesics and shape correspondence, existing deep learning methods rarely leverage geodesic distances as a primary supervision signal. The only exception is NeuroMorph (Eisenberger et al. 2021), which includes geodesic distances as a regularization term for isometric shapes but disables it when training on non-isometric shapes. In contrast, we demonstrate that by employing geodesic distances as the main supervisory objective for shape-matching tasks, we can reduce sensitivity to non-isometric variations, enabling LiteGE to deliver superior efficiency and accuracy compared to existing state-of-the-art methods.

Method

Overview

This section outlines our method for constructing a lightweight yet effective geodesic representation based on PCA applied to UDF data. The objective is to generate compact shape embeddings that enable efficient geodesic distance inference and shape matching across shape categories. The overall pipeline is illustrated in Figure 1.

Our method consists of three main stages. First, we canonically transform each shape to ensure a consistent scale, orientation, and position across the dataset. Next, we construct compact shape embeddings by applying PCA to UDF data from informative points surrounding the shape. This typically results in 200- to 400-dimensional descriptors. Once the PCA-based embeddings have been computed, they are used in conjunction with a lightweight neural network to predict geodesic distances and perform fast shape matching.




Figure 1: Illustration of shape descriptor creation process in LiteGE. First, each input point cloud is aligned to a canonical upright orientation. Next, the shape is voxelized and the informative voxels near the shapes surface are selected. Finally, we compute the unsigned distance (UDF) from the informative voxels to the shape to construct the final shape descriptor. This is as shown on the rightmost figure where we depicted the cross-section slice of the UDF heatmap through the center of the shape.

Motivations

The key motivation behind LiteGE is that shapes within a constrained set of multiple categories share similar geometric structures. By applying shape canonicalization, which aligns scale, orientation, and position, we can significantly reduce variance across the dataset, thereby allowing the use of the PCA method to create compact yet effective shape representations. This simplification makes the learning task easier: training an MLP to separate high-dimensional, unnormalized data would require much more data and network capacity to learn complex decision boundaries. In contrast, working with normalized inputs allows the model to learn from more compact, lower-dimensional representations.

Unlike spectral methods that rely on invariant shape descriptors such as the Laplace–Beltrami Operator (LBO), which require costly eigenvalue decompositions, our approach leverages efficient normalization. This enables simpler and more flexible network architectures, leading to better scalability, faster runtime, and improved generalization on sparse or noisy inputs.

Shape Canonicalization

Our method supports geodesic inference on both meshes and point clouds. The only difference lies in the canonicalization step, which standardizes the input shape’s position, orientation, and scale. For geodesic regression, we first convert all inputs to point clouds. For shape matching, we also convert the input to point cloud, but the mesh is kept when we receive mesh as the test data. Shape matching involves two shapes, thereby increasing data variability. More precise normalization that reduces data variability can be done more efficiently using mesh data. Canonicalization is then performed in three steps: centering, scaling, and orientation.

For centering, we shift each shape so that the point cloud sampled from it has its mean coordinate at the origin. For scaling, we adopt two strategies depending on the task: (i) Geodesic regression: We scale the point cloud such that its bounding box has a fixed surface area (e.g., 1.7 in our implementation). This method is simple and computationally efficient. (ii) Shape matching: When meshes are available, we scale each shape to have a unit surface area. For point cloud inputs, we replicate our strategy to scale meshes by making the average nearest-neighbor distance among K Poisson

disk samples (Corsini, Cignoni, and Scopigno 2012) match that of a unit-area mesh.

Similarly, we adopt two orientation strategies: (i) Geodesic regression: We adopt PCA to find the principal components of the point cloud and align them to the global x , y , and z axes. To resolve the sign ambiguity of each principal axis, we flip the axis direction (if necessary) so that it points toward the side of the shape with more point samples. The third axis is determined via the cross product of the first two to ensure a right-handed coordinate system. (ii) Shape matching: We adopt T-Net alignment since orientation consistency is important for reducing data variability when dealing with two shapes. We train a T-Net (Qi et al. 2017) to predict two orthogonal 3D vectors, which are converted into a rotation matrix using Gram–Schmidt orthogonalization process (Zhou et al. 2019). The network is supervised using two terms: a geodesic loss that measures the angular difference between the predicted and ground truth rotation matrices (Alvarez-Tunon, Brodskiy, and Kayacan 2023), and an $L2$ loss that aligns the transformed point coordinates. The loss is given by:

$$L_{\text{align}} = \cos^{-1} \left(\frac{\text{Tr}(\hat{R}^T R) - 1}{2} \right) + \frac{2}{N} \sum_{i=1}^N \|\hat{p}_i - p_i\|_2,$$

where N is the number of input points, \hat{R} and R are the predicted and ground-truth rotation matrices, and \hat{p}_i and p_i are the aligned and ground-truth points coordinates.

PCA Representation Using Informative Voxels

After canonicalizing each input shape, we convert it into a voxel-based representation. Specifically, we used a binary voxel grid with a resolution of 128^3 . Each voxel is assigned a value of 1 if it lies within the shape and 0 otherwise (Figure 1, middle). In our application domain – non-isometric shape matching involving animals and humans – most voxels show zero or near-zero variance across the dataset, suggesting they carry little discriminative information.

To improve efficiency and focus on meaningful structures, we retain only a subset of *informative voxels* whose occupancy variance exceeds a threshold ε . These voxels typically lie near the surface of the shape and capture the geometric variation needed to distinguish shapes effectively.

As further note, voxelization can be done regardless whether the input is a point cloud or a mesh. But, in our




Figure 2: Cumulative variance captured by principal components of UDFs on the SMAL dataset. “Point Cloud UDF” denotes shape canonicalization using bounding-box scaling and PCA-based axis alignment, while “Unit Area Mesh UDF” refers to T-Net alignment with unit-area scaling. In both settings, the first 50 principal components account for over 95% of the total UDF variance, highlighting the effectiveness of PCA in creating compact shape descriptors.

experiments, the informative voxels are precomputed from available meshes dataset and frozen thereafter.

For each shape, we then compute the unsigned distance (UDF) from the center of each informative voxel to the surface, implemented efficiently via nearest-neighbor queries. The resulting UDF values form a high-dimensional vector that encodes the shape’s geometry. Even with a conservative variance threshold (e.g., $\epsilon \leq 0.0625$), each shape typically retains around 20K to 30K informative voxels, resulting in a high-dimensional representation that limits practical usage.

To reduce dimensionality while preserving geometric expressiveness, we apply principal component analysis to UDF vectors and retain only the top 50–400 components. As shown in Figure 2, the first 50 principal components already capture over 95% of the variance, confirming the strong redundancy in the full representation and justifying the effectiveness of the resulting UDF-PCA shape descriptor.

As shown in the figure, the mesh-based normalization used in shape matching achieves slightly higher explained variance than the point cloud-based normalization for geodesic regression. This is expected since mesh-based normalization employs unit-area scaling and TNet alignment, which are more consistent than PCA axis alignment and bounding-box scaling for point clouds. PCA alignment is sensitive to axis flipping, and bounding-box scaling is less precise. This supports the observation that better normalization reduces data variability and effective dimensionality.

Network Architecture

Geodesic regression The LiteGE architecture employs a series of lightweight multilayer perceptrons designed to predict geodesic distances between point pairs. The key idea is to embed each query point into a high-dimensional space such that the difference between two-point embeddings reflects their geodesic distance. Compared to previous methods such as GeGNN (Pang et al. 2023) and NeuroGF (Zhang et al. 2023), our setup is significantly simpler and more memory-efficient due to the use of compact UDF-PCA shape descriptors instead of large 3D backbone networks.

As illustrated in Figure 3, the UDF-PCA shape descriptor is first passed through an MLP to produce a global shape em-

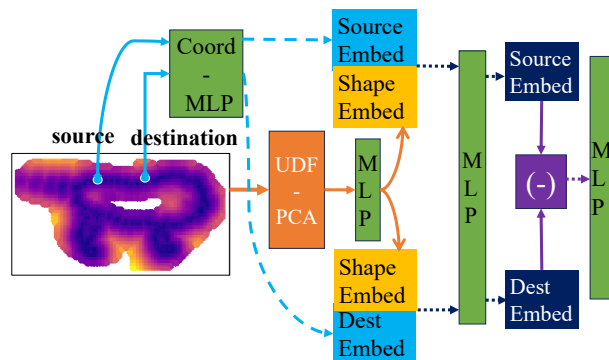


Figure 3: Architecture of LiteGE for geodesic regression. The UDF-PCA shape descriptor and the coordinates of source and destination points are embedded using separate MLPs. These embeddings are combined to produce two shape-specific point embeddings. The geodesic distance is predicted from the difference between the point embeddings.

bedding. The 3D coordinates of the source and destination points are processed through a shared MLP (Coord-MLP) to generate point-wise coordinate embeddings. Each of these is concatenated with the shape embedding and passed through another shared MLP to obtain shape-specific embeddings for the source and destination points. Their difference is then inputted into a final MLP that predicts the geodesic distance.

Extension to shape matching To adapt the network for shape matching, we estimate the geodesic distance between two points that lie on different shapes (e.g., shapes X and Y), treating them as if they reside on the same surface. We modify the architecture by feeding both UDF-PCA descriptors—one from shape X and one from shape Y —into the shared shape MLP. The corresponding query points for each shape are embedded via the Coord-MLP. Each point embedding is then concatenated with its respective shape embedding and passed through the shared-point MLP to produce shape-specified embeddings, as in the single-shape case. The final MLP predicts the average geodesic distance between the two points, measured separately in both the shape X and the shape Y .

Implementation The Coord-MLP, which encodes 3D point coordinates, consists of two layers, each with 200 neurons. The MLP for generating the global shape embedding from the UDF-PCA descriptor also has three layers, each containing 200 neurons. The shared MLP for producing shape-specific point embeddings contains three layers, with 200 neurons per layer. Finally, the prediction MLP, which estimates the geodesic distance from the difference between source and destination embeddings, consists of two layers with 400 neurons each. In total, LiteGE has approximately 600K parameters, only 10% of GeGNN, signifying its compact and efficient design.

Training Ground-truth geodesic distances are generated using FastDGG (Adikusuma, Fang, and He 2020), with accuracy thresholds of 0.3% for training and 0.1% for testing. Here, $L1$ loss is used to train LiteGE networks.

Inference

During inference, each shape is first canonicalized, as described previously. Then, our network requires two inputs: (1) the UDF-PCA representation of a canonicalized shape, and (2) the 3D coordinates of the query points (e.g., source and destination).

To construct the UDF-PCA descriptor, we first compute the unsigned distances by measuring the distance from each informative voxel center to the shape surface using nearest-neighbor queries. This UDF vector is then projected onto the principal component basis, resulting in a compact representation. Specifically, the i -th component of the UDF-PCA vector is computed as:

$$z_{\text{UDF-PCA},i} = \langle \mathbf{x}_{\text{UDF}} - \bar{\mathbf{x}}_{\text{UDF}}, \mathbf{p}_{\text{UDF},i} \rangle,$$

where \mathbf{x}_{UDF} is the UDF vector of the current shape, $\bar{\mathbf{x}}_{\text{UDF}}$ is the mean UDF vector computed from the training set, and $\mathbf{p}_{\text{UDF},i}$ is the i -th principal component.

When handling multiple queries on a single shape, we can precompute the shape embedding for reuse. If certain query points appear multiple times, their shape-specific embeddings can also be cached. The same optimization applies in shape matching, where embeddings of both input shapes can be precomputed once and reused across multiple queries.

Shape Correspondence

We perform shape matching using a coarse-to-fine strategy, as detailed in Algorithms 1 and 2. Given a query point Q_X on shape X , the algorithm efficiently identifies its best match on shape Y by iteratively refining the predictions, starting from a sparse set of candidate points.

The process begins by constructing a multi-tier nearest-neighbor cache on shape Y (Algorithm 1). We define K tiers of point sets, where Tier 1 is the sparsest (e.g., N_1 points), and Tier K includes all mesh vertices. For each point in Tier i , we store its M_i nearest neighbors in Tier $i+1$.

Using the cache, Algorithm 2 finds a match for Q_X by sequentially querying the geodesic prediction network G at increasing resolution. In each tier, we select the best match from the current candidate set and use the cache to retrieve a set of neighbors for the next tier. This iterative process continues until Tier K , where the final match is returned.

This coarse-to-fine strategy significantly reduces the number of geodesic queries required, enabling efficient matching even on high-resolution meshes. In our implementation, we use three cache tiers, each created by randomly sampling mesh vertices. Both NN-cache construction and query operations are parallelized using `pynanoflann` and `Pytorch`.

Geodesic Path Tracing

To recover a geodesic path between two points, we use a gradient-based tracing method similar to GeGNN (Pang et al. 2023). The key idea is to iteratively update the destination point along the negative gradient of the predicted geodesic distance until convergence.

Let \mathbf{s} denote the source point, $\mathbf{d}^{(t)}$ the destination point at iteration t , and $\mathbf{z}_{\text{UDF-PCA}}$ the shape descriptor. The geodesic

Algorithm 1: Multi-Tiered Nearest Neighbor Caching

Require: Shape \mathcal{S} , multi-tiered point sets $\{P_i(\mathcal{S})\}_{i=1}^K$, where $P_i(\mathcal{S})$ contains N_i sampled points and $N_1 < \dots < N_K$ (with N_K being all vertices); neighbor counts M_1, \dots, M_{K-1}

Ensure: K -tiered Nearest Neighbor Cache \mathcal{C} for shape \mathcal{S}

- 1: Initialize $\mathcal{C}(\mathcal{S}) \leftarrow \emptyset$
- 2: **for** $i = 1$ **to** $K - 1$ **do** ▷ Construct cache between adjacent tiers
- 3: **for** each point $p \in P_i(\mathcal{S})$ **do**
- 4: Find M_i nearest neighbors of p in $P_{i+1}(\mathcal{S})$ using Euclidean distance.
- 5: Store the mapping $p \rightarrow \mathcal{N}_{p,i+1}$ in $\mathcal{C}(\mathcal{S})$
- 6: **return** $\mathcal{C}(\mathcal{S})$

Algorithm 2: Coarse-to-Fine Shape Matching

Require: Query point q_X on shape X , nearest neighbor cache $\mathcal{C}(Y)$ for shape Y , geodesic inference network G , UDF-PCA descriptors $\mathbf{z}_{\text{pca}}(X)$ and $\mathbf{z}_{\text{pca}}(Y)$, multi-tiered point sets $\{P_i(Y)\}_{i=1}^K$

Ensure: Corresponding point q_Y on shape Y

- 1: Initialize candidate set $\mathcal{C}_{\text{cand}} \leftarrow P_1(Y)$
- 2: **for** $i = 1$ **to** K **do** ▷ Iterate through tiers
- 3: $q_i(Y) \leftarrow \arg \min_{p \in \mathcal{C}_{\text{cand}}} G(q_X, p, \mathbf{z}_{\text{pca}}(X), \mathbf{z}_{\text{pca}}(Y))$
▷ Query for best match in the current candidates
- 4: **if** $i < K$ **then** ▷ If not the finest tier
- 5: $\mathcal{C}_{\text{cand}} \leftarrow$ the M_i nearest neighbors of $q_i(Y)$ from $P_{i+1}(Y)$ using $\mathcal{C}(Y)$ ▷ Refine search space
- 6: **else** ▷ Reach the finest tier
- 7: $q_Y \leftarrow q_K(Y)$ ▷ Final match found
- 8: **return** q_Y

prediction $G(\mathbf{s}, \mathbf{d}^{(t)}, \mathbf{z}_{\text{UDF-PCA}})$ is differentiated with respect to $\mathbf{d}^{(t)}$ to obtain the gradient, which guides the update:

$$\mathbf{d}^{(t+1)} = \text{Proj} \left(\mathbf{d}^{(t)} - \eta \nabla_{\mathbf{d}} G(\mathbf{s}, \mathbf{d}^{(t)}, \mathbf{z}_{\text{UDF-PCA}}) \right),$$

where η is the learning rate and $\text{Proj}(\cdot)$ denotes projection onto the shape (e.g., via nearest-neighbor search in the point cloud). The process terminates when the predicted distance ceases to decrease or falls below a threshold ε .

If a mesh is available, the resulting path, represented by a sequence of intermediate points, can be projected onto dense mesh samples to ensure that it lies on the surface.

Results and Discussions

All experiments were conducted on a dedicated GPU server equipped with an NVIDIA RTX A4000 GPU, an Intel Xeon Gold 5315Y CPU, and 45 GB of RAM. We evaluate the efficiency and effectiveness of LiteGE across three key tasks: geodesic distance prediction, non-isometric shape matching, and geodesic path tracing. We begin by assessing LiteGE’s core capability in predicting geodesic distances, then demonstrate its application to downstream tasks. Due to space constraints, additional implementation details, ablation, and extended experiments are provided in the appendix.

N_Q	Memory (MB)		Inference Time (ms)	
	NeuroGF	LiteGE	NeuroGF	LiteGE
1	162	1 (162×)	88	1.9 (47×)
4	647	2.7 (240×)	107	2 (54×)
8	1294	4.5 (288×)	145	2.6 (56×)
32	5172	17.25 (296×)	763	7.3 (104×)
64	10344	34.5 (300×)	1529	13.4 (114×)
96	OOM	51.8	OOM	19.1

Table 1: Comparison of LiteGE with NeuroGF in terms of memory consumption and runtime performance for varying numbers of queries. Reported runtimes include all necessary preprocessing for both methods. Each point cloud receives a single query, and N_Q denotes the total number of distinct queries evaluated. All tests are conducted on 2K point clouds. Improvements of LiteGE over NeuroGF are highlighted in orange. “OOM” indicates out-of-memory failure.

Lightweight Geodesic Distance Inference

We evaluate LiteGE against NeuroGF (Zhang et al. 2023) using the official codes and pretrained model provided by the authors. Both methods are tested on point cloud inputs, with shape canonicalization performed as described in the Methods section. NeuroGF uses a DGCNN (Wang et al. 2019) backbone, a deep graph convolutional network with a large multi-layer architecture (5–6 layers), which results in high memory usage and inference costs. In contrast, LiteGE employs a compact MLP-based architecture and uses UDF-PCA shape descriptors, significantly reducing computational costs. Since GeGNN (Pang et al. 2023) does not support point cloud input, we exclude it from comparisons. However, we note that GeGNN and NeuroGF exhibit similar accuracy and runtime performance when evaluated on mesh inputs, as reported in their respective papers.

For benchmarking, we use the SMAL PCA model (Zuffi et al. 2017) to generate 12,000 diverse quadruped animal meshes. All reported errors are normalized such that the ground-truth geodesic distances have a mean value of 100.0. We exclude the memory cost of storing the UDF-PCA basis, as it remains constant (10–20 MB) regardless of the number of queried point clouds.

As shown in Table 1, LiteGE is up to 300× more memory-efficient and over 100× faster than NeuroGF. Moreover, NeuroGF shows significantly lower accuracy across our SMAL test cases (Table 2), indicating limited capacity to generalize to unseen shape categories. In contrast, LiteGE remains robust even on sparse point clouds with only 300 samples (Figure 4), where NeuroGF fails with a 75 L1 error. This is likely because unsigned distances from informative voxels in LiteGE remain stable across sampling densities.

Also, LiteGE is robust to shapes with noise and missing regions. With the ground-truth distances normalized to a mean of 100, given 2K-point clouds sampled from unit area meshes in the SMAL dataset, LiteGE has a mean L1 error of 4.3 when 0.01-scale Gaussian noise is added to the point clouds. Moreover, LiteGE achieves mean L1 errors of 3.7 and 4.5 when 10% and 15% of the points are randomly removed by two slicing planes.




Figure 4: Visual comparison of geodesic distance predictions on a sparse (300-point) pig model, with the source point located on the left front paw. Distances are visualized using a heatmap: warmer colors indicate larger distances, while cooler colors indicate smaller ones. The mean prediction errors are 24 for NeuroGF and 2.5 for LiteGE, with ground-truth distances normalized to a mean of 100. This highlights LiteGE’s accuracy on sparse point clouds.

LiteGE Train Data	Test Data	#Samples	ϵ_{LGE}	ϵ_{NGF}
SMAL	SMAL	300	3.4	75
SMAL	SMAL	2000	2.3	13
SMAL & SURREAL	FAUST	700	2.6	NA
Objaverse-XL	Objaverse-XL	2000	3.9	NA

Table 2: Comparison of mean L1 errors of LiteGE (ϵ_{LGE}) and NeuroGF (ϵ_{NGF}) on the SMAL dataset under varying point-cloud sample densities, as well as LiteGE’s generalization ability (measured as median L1 errors) on unseen (FAUST) or diverse (Objaverse) datasets. Notably, NeuroGF fails on sparse point clouds with 300 samples. Ground-truth distances are normalized to have a mean of 100.

Finally, in Table 2, LiteGE demonstrates strong generalization to unseen and diverse datasets. When trained on a combined dataset of 8K meshes from the SMAL model (Zuffi et al. 2017) and 8K from the SURREAL human dataset (Groueix et al. 2018), LiteGE generalizes to the FAUST dataset (Bogo et al. 2014), even when tested using sparse inputs (700 points). Furthermore, when trained on 3K models from the Objaverse-XL (Deitke et al. 2023) dataset and tested on a separate validation set, LiteGE generalizes across diverse shape categories.

Fast and Accurate Shape Matching

Table 3 presents the results of our shape matching experiments on the 12K SMAL meshes used for geodesic inference evaluation. For both training and testing, we used randomly rotated meshes that were then canonicalized using T-Net and our mesh normalization pipeline. To ensure reliable supervision, we only included meshes with alignment error of less than 0.66 radians during training. To stabilize training, we repeated the first training epoch eight times and selected the best model checkpoint before continuing training.

We compare our method with Spectral-Meets-Spatial (SMS) (Cao et al. 2024), a recent state-of-the-art approach. Due to disk space constraints for storing LBO parameters, we trained SMS on 6K meshes. However, as shown in the appendix, the evaluation accuracy of SMS saturates after using more than 1200 meshes. SMS was trained using its official training schedule with the non-isometric setting en-

Mesh Type	V	SMS			LiteGE		
		T (ms)	AUC	Error	T (ms)	AUC	Error
Template	4K	21300	79.4	2.2	185	79.3	2.5
Remeshed	5K	27000	73.1	7.44	119	74	7.2
Anisotropic	10K	38500	74.9	6.8	120	74.3	7
Broken Mesh	5K	25700	33.2	28	185	69.5	8
Point Cloud	8K	NA			200	71.5	7.6
Point Cloud	500	NA			126	69.2	8.1

Table 3: Non-isometric shape correspondence results. We query 1K points matches using LiteGE and report mean geodesic error scaled by 100 on unit-area meshes. For template models, we use a threshold of 0.1 to compute the AUC; for all other settings, a threshold of 0.2 is used. Moreover, for template models, ensemble of 2 LiteGE is used to test, while no ensemble is used on other tests. Here, SMS runtime includes both the LBO precomputation and final query time. The LiteGE runtime includes UDF-PCA generation, NN-cache construction, and final query time.

abled, sampling 6400 random shape pairs per epoch.

For LiteGE, we used 1000 query points to ensure good coverage of the target mesh, particularly for high-resolution inputs. The configuration of our multi-tiered NN-cache is detailed in the appendix. We evaluated both methods on a variety of test inputs, including (1) SMAL template meshes, (2) isotropic and anisotropic remeshed models, (3) broken meshes with 40% of faces randomly removed, and (4) point clouds with either 8K or 500 samples. For remeshed models, ground-truth correspondences were obtained via barycentric projection to template meshes. Note that we evaluate only across distinct shapes; matching identical shapes (once uniformly scaled, centered, and PCA-aligned) is trivial.

For broken meshes and point clouds, we adopt a canonicalization strategy that scales the shape so that the average nearest-neighbor distance of K Poisson-disk samples matches that of a unit area mesh.

As shown in Table 3, LiteGE achieves an accuracy comparable to SMS on clean meshes, while being up to $300\times$ faster when matching 1K points. On the 5K-vertex remeshed model, LiteGE computes the UDF-PCA descriptors and NN-cache in 13 ms before performing per-point matching. Thus, it is up to $1000\times$ faster than SMS when matching 250 points or less in this setting. On remeshed models as well, arguably a more realistic test case, LiteGE generalizes better thanks to its UDF-based representation, outperforming SMS in accuracy, as shown in Figure 5. We also show in the appendix that augmenting SMS training with remeshed shapes does not improve its generalization. Furthermore, SMS fails in broken meshes (with error > 25), while LiteGE maintains strong performance. LiteGE also generalizes well to dense and sparse point clouds, achieving mesh-level accuracy even on sparse inputs with just 500 points (Figure 5).

Geodesic Paths Tracing

In addition to geodesic distance inference, LiteGE also supports geodesic path tracing via gradient-based backtracking. As shown in Table 4, LiteGE enables accurate path recon-




Figure 5: Visual results on shape matching. Left: LiteGE accurately performs point cloud matching on an 8K-sample point cloud derived from a unit-area mesh, achieving a mean geodesic error of 0.022. Matched point pairs are connected by lines. Right: Shape matching performance of LiteGE and SMS on remeshed 5K-vertex models and their broken variants (with 40% of faces randomly removed). Average errors across all test shape pairs are visualized on a dog model.

#PC Samples	Error	Memory (MB)	Learning rate
800	4.73	4.4	0.66
1,000	4.27	4.22	0.66
8,000	3.77	5.1	0.35
15,000	3.66	6.4	0.25

Table 4: Geodesic path tracing results on a batch of eight shapes randomly chosen from the SMAL dataset. We report the median prediction error, with ground truth distances scaled such that their mean is 100. The learning rate used for each example is listed in the last column.

struction on a batch of eight sparse point clouds, each containing fewer than 1,000 points. Remarkably, this is achieved with a memory footprint of under 5 MB.

Conclusion and Future Directions

We introduced LiteGE, a lightweight neural framework for geodesic distance prediction on 3D shapes, based on PCA embeddings of unsigned distance fields. Unlike prior neural methods that rely on large, memory-intensive backbones, LiteGE offers a highly efficient and scalable solution that supports both meshes and point clouds, including sparse, noisy or incomplete inputs. Extensive experiments show that LiteGE maintains an accuracy comparable to or superior to that of state-of-the-art methods while achieving over $300\times$ reductions in memory and runtime. Moreover, we demonstrated its effectiveness in non-isometric shape matching, where LiteGE achieves up to $1000\times$ speedup over existing methods and remains robust under sparse inputs.

Future work includes extending LiteGE to produce geodesic-aware shape descriptors for other downstream tasks, such as shape segmentation, analysis, and retrieval. Furthermore, incorporating temporal consistency for dynamic or time-varying shapes would further expand the practical utility of LiteGE in real-world 4D applications. Finally, exploring self-supervised or few-shot variants could reduce dependence on labeled data, improving generalization across diverse shape categories.

Acknowledgments

This project was partially supported by the NSF-IIS-2413161 grant and the Ministry of Education, Singapore, under its Academic Research Fund Grant (RT19/22).

References

- Abdi, H.; and Williams, L. J. 2010. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4): 433–459.
- Adikusuma, Y. Y.; Du, J.; Fang, Z.; and He, Y. 2022. An Accuracy Controllable and Memory Efficient Method for Computing High-Quality Geodesic Distances on Triangle Meshes. *Computer-Aided Design*, 150: 103333.
- Adikusuma, Y. Y.; Fang, Z.; and He, Y. 2020. Fast Construction of Discrete Geodesic Graphs. *ACM Trans. Graph.*, 39(2): 14:1–14:14.
- Alvarez-Tunon, O.; Brodskiy, Y.; and Kayacan, E. 2023. Loss it right: Euclidean and riemannian metrics in learning-based visual odometry. In *ISR Europe 2023: 56th International Symposium on Robotics*, 107–111. VDE.
- Belyaev, A. G.; and Fayolle, P.-A. 2015. On Variational and PDE-Based Distance Function Approximations. *CGF*, 34(8): 104–118.
- Bogo, F.; Romero, J.; Loper, M.; and Black, M. J. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3794–3801.
- Cao, D.; Eisenberger, M.; El Amrani, N.; Cremers, D.; and Bernard, F. 2024. Spectral meets spatial: Harmonising 3d shape matching and interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3658–3668.
- Cao, D.; Roetzer, P.; and Bernard, F. 2023. Unsupervised learning of robust spectral shape matching. *arXiv preprint arXiv:2304.14419*.
- Chen, J.; and Han, Y. 1990. Shortest paths on a polyhedron. In *SoCG*, 360–369.
- Chibane, J.; Pons-Moll, G.; et al. 2020. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems*, 33: 21638–21652.
- Corsini, M.; Cignoni, P.; and Scopigno, R. 2012. Efficient and flexible sampling with blue noise properties of triangular meshes. *IEEE transactions on visualization and computer graphics*, 18(6): 914–924.
- Crane, K.; Weischedel, C.; and Wardetzky, M. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph.*, 32(5): 152.
- Deitke, M.; Liu, R.; Wallingford, M.; Ngo, H.; Michel, O.; Kusupati, A.; Fan, A.; Laforte, C.; Voleti, V.; Gadre, S. Y.; et al. 2023. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36: 35799–35813.
- Eisenberger, M.; Novotny, D.; Kerchenbaum, G.; Labatut, P.; Neverova, N.; Cremers, D.; and Vedaldi, A. 2021. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7473–7483.
- Groueix, T.; Fisher, M.; Kim, V. G.; Russell, B. C.; and Aubry, M. 2018. 3d-coded: 3d correspondences by deep deformation. In *Proceedings of the european conference on computer vision (ECCV)*, 230–246.
- Hanocka, R.; Hertz, A.; Fish, N.; Giryas, R.; Fleishman, S.; and Cohen-Or, D. 2019. Meshcnn: a network with an edge. *ACM Trans. Graph.*, 38(4): 1–12.
- Lee, H.; Tong, Y.; and Desbrun, M. 2005. Geodesics-Based One-to-One Parameterization of 3D Triangle Meshes. *IEEE Multim.*, 12(1): 27–33.
- Litany, O.; Remez, T.; Rodola, E.; Bronstein, A.; and Bronstein, M. 2017. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision*, 5659–5667.
- Liu, B.; Chen, S.; Xin, S.; He, Y.; Liu, Z.; and Zhao, J. 2017. An optimization-driven approach for computing geodesic paths on triangle meshes. *Comput. Aided Des.*, 90: 105–112.
- Masci, J.; Boscaini, D.; Bronstein, M.; and Vandergheynst, P. 2015. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, 37–45.
- Meng, W.; Xin, S.; Tu, C.; Chen, S.; He, Y.; and Wang, W. 2022. Geodesic Tracks: Computing Discrete Geodesics With Track-Based Steiner Point Propagation. *IEEE Trans. Vis. Comput. Graph.*, 28(12): 4887–4901.
- Mitchell, J. S.; Mount, D. M.; and Papadimitriou, C. H. 1987. The discrete geodesic problem. *SICOMP*, 16(4): 647–668.
- Ovsjanikov, M.; Ben-Chen, M.; Solomon, J.; Butscher, A.; and Guibas, L. J. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4): 30:1–30:11.
- Pang, B.; Zheng, Z.; Wang, G.; and Wang, P.-S. 2023. Learning the geodesic embedding with graph neural networks. *ACM Transactions on Graphics (TOG)*, 42(6): 1–12.
- Peyré, G.; Péchaud, M.; Keriven, R.; Cohen, L. D.; et al. 2010. Geodesic methods in computer vision and graphics. *Foundations and Trends® in Computer Graphics and Vision*, 5(3–4): 197–397.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Qin, Y.; Han, X.; Yu, H.; Yu, Y.; and Zhang, J. 2016. Fast and Exact Discrete Geodesic Computation Based on Triangle-Oriented Wavefront Propagation. *ACM Trans. Graph.*, 35(4).
- Sethian, J. A. 1996. A fast marching level set method for monotonically advancing fronts. *Proceedings of National Academy of Sciences*, 93: 1591–1595.
- Sethian, J. A.; and Vladimirsky, A. 2000. Fast methods for the Eikonal and related Hamilton-Jacobi equations on unstructured meshes. *Proceedings of National Academy of Sciences*, 97: 5699–5703.

- Sharp, N.; and Crane, K. 2020. You can find geodesic paths in triangle meshes by just flipping edges. *ACM Trans. Graph.*, 39(6): 249:1–249:15.
- Sprynski, N.; Szafran, N.; Lacolle, B.; and Biard, L. 2008. Surface reconstruction via geodesic interpolation. *Computer-Aided Design*, 40(4): 480–492.
- Sun, Q.; Zhang, L.; Zhang, M.; Ying, X.; Xin, S.; Xia, J.; and He, Y. 2013. Texture brush: an interactive surface texturing interface. In *Symposium on Interactive 3D Graphics and Games, I3D '13*, 153–160.
- Surazhsky, V.; Surazhsky, T.; Kirsanov, D.; Gortler, S. J.; and Hoppe, H. 2005. Fast Exact and Approximate Geodesics on Meshes. *ACM Trans. Graph.*, 24(3): 553–560.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (tog)*, 38(5): 1–12.
- Weber, O.; Devir, Y. S.; Bronstein, A. M.; Bronstein, M. M.; and Kimmel, R. 2008. Parallel Algorithms for Approximation of Distance Maps on Parametric Surfaces. *ACM Trans. Graph.*, 27(4): 104:1–104:16.
- Xin, S.-Q.; He, Y.; and Fu, C.-W. 2012. Efficiently Computing Exact Geodesic Loops within Finite Steps. *TVCG*, 18(6): 879–889.
- Xin, S.-Q.; and Wang, G.-J. 2009. Improving Chen and Han’s algorithm on the discrete geodesic problem. *ACM Trans. Graph.*, 28(4): 104.
- Xin, S.-Q.; Ying, X.; and He, Y. 2012. Constant-time all-pairs geodesic distance query on triangle meshes. In *I3D*, 31–38.
- Xu, C.-X.; Wang, T. Y.; Liu, Y.-J.; Liu, L.; and He, Y. 2015. Fast Wavefront Propagation (FWP) for Computing Exact Geodesic Distances on Meshes. *IEEE Transactions on Visualization and Computer Graphics*, 21(7): 822–834.
- Ying, X.; Wang, X.; and He, Y. 2013. Saddle vertex graph (SVG): A novel solution to the discrete geodesic problem. *ACM Trans. Graph.*, 32(6): 170:1–12.
- Ying, X.; Xin, S.-Q.; and He, Y. 2014. Parallel Chen-Han (PCH) Algorithm for Discrete Geodesics. *ACM Transactions on Graphics*, 33(1): 9:1–9:11.
- Yuan, N.; Wang, P.; Meng, W.; Chen, S.; Xu, J.; Xin, S.; He, Y.; and Wang, W. 2023. A Variational Framework for Curve Shortening in Various Geometric Domains. *IEEE Trans. Vis. Comput. Graph.*, 29(4): 1951–1963.
- Zhang, Q.; Hou, J.; Adikusuma, Y.; Wang, W.; and He, Y. 2023. NeuroGF: A neural representation for fast geodesic distance and path queries. *Advances in Neural Information Processing Systems*, 36: 19485–19501.
- Zhao, H.; Jiang, L.; Jia, J.; Torr, P. H.; and Koltun, V. 2021. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 16259–16268.
- Zhou, Y.; Barnes, C.; Lu, J.; Yang, J.; and Li, H. 2019. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5745–5753.
- Zuffi, S.; Kanazawa, A.; Jacobs, D.; and Black, M. J. 2017. 3D Menagerie: Modeling the 3D Shape and Pose of Animals. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.