

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

GraphWalks: Efficient Shape Agnostic Geodesic Shortest Path Estimation

Rolandos Alexandros Potamias¹, Alexandros Neofytou², Kyriaki Margarita Bintsi¹, Stefanos Zafeiriou¹ ¹Imperial College London, ²OCM Digital Media

{r.potamias,m.bintsi19,s.zafeiriou}@imperial.ac.uk, alex@orangeclickmedia.com

Abstract

Geodesic paths and distances are among the most popular intrinsic properties of 3D surfaces. Traditionally, geodesic paths on discrete polygon surfaces were computed using shortest path algorithms, such as Dijkstra. However, such algorithms have two major limitations. They are nondifferentiable which limits their direct usage in learnable pipelines and they are considerably time demanding. To address such limitations and alleviate the computational burden, we propose a learnable network to approximate geodesic paths. The proposed method is comprised by three major components: a graph neural network that encodes node positions in a high dimensional space, a path embedding that describes previously visited nodes and a point classifier that selects the next point in the path. The proposed method provides efficient approximations of the shortest paths and geodesic distances estimations. Given that all of the components of our method are fully differentiable, it can be directly plugged into any learnable pipeline as well as customized under any differentiable constraint. We extensively evaluate the proposed method with several qualitative and quantitative experiments.

1. Introduction

Geodesic paths arise in several fields across 3D computer vision, computer graphics and digital geometry with a broad amount of applications ranging from mesh morphing [41] to robotics and circuit design [7]. A geodesic path can be considered any "straight" curve that connects two points on a surface [10]. It is an intrinsic property of the 3D shape that remains unaffected by non-rigid deformations providing an expressive shape descriptor [25, 49, 55]. In a similar manner, geodesic distance can be defined as the distance of two points measured on the curved surface. Compared to Euclidean distance, geodesic distance is usually assumed to be a nonlinear function between the surface points. Traversal on different parts of a 3D object requires structural shape information that Euclidean distance usually fails to describe providing highly misleading results [39, 40]. The computation of geodesic paths and distances requires careful consideration of the intrinsic surface properties to embody the influence of curvature. Although several techniques have attempted to embed intrinsic properties of the points and directly calculate geodesic distances on the continuous surface domain, in most of the 3D computer vision and graphics tasks, usually, such computations are made on discretized polyhedral surfaces, i.e. meshes.

Given that 3D meshes can be considered as graphs embedded in the 3D Euclidean space, geodesic paths on discrete meshes are usually calculated using shortest path algorithms. Shortest path calculation between graph nodes lies in the core of many computer science algorithms and has been extensively studied with applications ranging from road networks and urban transportation [28, 54] to path planning [11]. Traditionally, algorithms for shortest paths calculation such as breadth-first-search (BFS), A^{*} [24], Dijkstra and Bellman-Ford suffer from diminishing performance when scaling to significantly larger graphs. Due to the extreme amount of applications requiring detection of shortest path several extensions of the aforementioned algorithms have been developed to reduce the computational and storage requirements. Additionally, several approximation methods have been proposed [21, 46] to address the shortcomings of the traditional methods. Recently, several methods [43,47] exploited the representation capabilities to estimate shortest path distances using deep learning. However, they rely on graph-specific node embeddings, such as Node2Vec [20] or DeepWalk [36], and can not be generalized to different graph topologies without re-training. Additionally, the aforementioned neural approximation methods only provide shortest path distances between graph nodes, thus the user can not directly infer the path connecting those nodes. In contrast, the shortest path estimation problem poses as a superordinate given that the shortest path distance can be easily computed using the weighted lengths of the predicted path.

In this study, we propose the first, to the best of our knowledge, neural-based geodesic path estimator that is fully differentiable with respect to its inputs. The proposed methods is equipped with a graph neural network (GNN) with a novel graph message passing operator to encode the topology priors of the surface. The path is predicted using an autoregressive fashion, where the current path state is encoded using a momentum paradigm [26]. The proposed method outperforms all baseline methods in geodesic distance estimation and enables a path prediction speedup of up to 9 times. Finally, we showcase that the proposed method can accurately estimate geodesic paths and distances on unstructured point clouds.

2. Related Work

2.1. Shortest Path

Finding the shortest path between a pair of a nodes with the purpose of minimizing the sum of edge weights is a problem that has been studied a lot in graph theory. Most of the existing shortest-path algorithms can be divided in two main categories, namely single-source shortest-path (SSSP), and all-pairs shortest-path (APSP). The algorithms may generate exact or approximate paths. While exact paths would be the optimal case, they are not always feasible to find, especially in large graphs. The APSP problem aims to extract the shortest paths between every possible pair of source \mathcal{V} and destination \mathcal{V}' nodes in the graph. The most popular APSP algorithm for both directed and undirected graphs is Floyd-Warshall [15] with a complexity of $\mathcal{O}(|\mathcal{V}|^3)$, where $|\mathcal{V}|$ is the number of the nodes. Another generalization of the shortest path problem is the SSSP problem, which consists of finding the shortest paths between a given vertex \mathcal{V} and all other vertices in the graph. A standard choice for unweighted graphs is BFS [8], which, however, scales poorly and is proven impractical for today's massive graphs. Dijkstra's algorithm [12] is used to solve the SSSP problem in both directed and undirected graphs with non-negative edge weights with a complexity of $\mathcal{O}((|\mathcal{V}| + |\mathcal{E}|) \log |\mathcal{V}|)$ for the original algorithm, where $|\mathcal{V}|$ is the number of the nodes and $|\mathcal{E}|$ is the number of the edges, while the fastest to date version of the algorithm achieves a complexity of $\mathcal{O}(|\mathcal{E}| + |\mathcal{V}| \log |\mathcal{V}|)$ using a Fibonacci heap instead of the original min-priority queue [13, 17]. This is asymptotically the fastest known SSSP algorithm for the mentioned kind of graphs, with many implementations to use case-based optimizations which are out of the scope of this paper. An extension to the Dijkstra's algorithm is \mathcal{A}^* [24], a simple goal-directed algorithm that falls under the category of the best-first search algorithms. \mathcal{A}^* finds the shortest path faster than Dijkstra, by using an admissible heuristic function. While the estimation of the exact shortest paths would be the optimal case, they are not always feasible to find, especially in very large graphs. In many cases, the estimation of an approximate path between two nodes is adequate. For this purpose, methods that include heuristic strategies [46] and landmark strategies [21]

has been proposed. A Center Distance to Zone (CDZ) has been suggested in [52] that considers 10% of the nodes as central nodes and uses the Dijksta's algorithm in order to find the shortest path between central nodes. Although the CDZ algorithm manages to approximate the distance pretty accurately in some cases, it suffers from poor scaling with performance to drop on modern scale graphs. Several studies extended the notion of path approximation by modeling stochastic paths with edges that follow a probabilistic distributions instead of static costs [5, 6, 30, 31]. However, in many practical scenarios there is no guarantee that the obtained random variables of edges are independent of each other [56]. Many modifications of \mathcal{A}^* have been implemented that combine the technique introduced in the \mathcal{A}^* algorithm with the use of landmarks [19,22]. A fast approximate of the real distance was suggested in [38], that uses a subset of nodes as landmarks and estimates the distance of each node from each landmark. This however creates the problem of the correct choice of landmarks, which proved crucial for the accuracy of the estimation. More recently, machine learning techniques have been incorporated in the research of shortest path estimation. The authors in [14] propose an MLP to support and guide Dijkstra's algorithm to reduce its search space. For an extensive survey of the filed we refer the reader to [16, 18].

2.2. Geodesic Distances

Geodesic distance metrics prove a key ingredient to a vast amount of application ranging from representation learning [55] to graph clustering [37]. Calculation of geodesic distance can be achieved using two different approaches, based on the perspective of the 3D structure. Methods in the first approach consider 3D polygon meshes as a graph structure and calculate the geodesic distance as sum of the shortest path segments, using traditional shortest path algorithms as reviewed in Section 2.1. In particular, an extension to Dijkstra algorithm, Fast Marching Method, along with several modifications of it [2, 27, 48], remains one of the most popular choices for geodesic distance computation on discretized surfaces. On the contrary, methods on the second category calculate geodesic distances using algorithms from differentiable geometry under the assumption of a discretized differentiable surface. Although Euclidean distance seems a natural choice, the implicit assumption of linear geometric structure provides a misleading estimate for how far is a point in the nonlinear structure of 3D shapes [1]. In [29], the authors proposed a sphere piece fitting method to estimate local distances and calculate the geodesic distance through paths over spheres. Recently, Sharp et al. [50] proposed a modification of Delaunay flip algorithm to estimate geodesic paths and distances by flipping edges of the mesh. However, such algorithms do not always produce actual shortest paths, hence cannot guarantee accurate geodesic distances. Apart from optimization based methods, there has been an attempt to model shortest path distances using neural networks. The authors of [47] proposed to utilize Node2Vec [20] embeddings to approximate shortest path distances using training node pairs obtained from pre-computed actual shortest path distances on several graph landmarks as in [57]. In a similar manner, Qi *et al.* [43] trained a neural network to predict vertex embeddings along with a MLP to regress distance values. However, Node2Vec embeddings as well as the embedding layer proposed in [43], can not be generalized to different graph topologies and thus they are impractical in 3D real-world application. For additional resources on geodesic paths and distances we refer the reader to extensive surveys [1, 33].

3. Method

Given a mesh $\mathcal{M}=(\mathcal{V}, \mathcal{E}, \mathcal{F})$ with vertices \mathcal{V} , edges \mathcal{E} and faces \mathcal{F} , the proposed method aims to accurately estimate the geodesic path between two points $\mathbf{v}_{str}, \mathbf{v}_{tgt} \in \mathcal{V}$ on the surface \mathcal{S} . It takes as input the current point \mathbf{p}^t , its neighborhood $\mathcal{N}(\mathbf{p}^t)$ along with the current path descriptor \mathbf{g}^t and predicts at every time step t the next point in the path sequence using an autoregressive manner:

$$\mathbf{p}^{t+1}, \mathbf{g}^{t+1} = \mathcal{G}(\mathbf{p}^t, \mathcal{N}(\mathbf{p}^t), \mathbf{g}^t)$$
(1)

In the following section we describe the three main modules of the proposed method, namely: Vertex Embedding Layer, Path Descriptor and Point Selector along with the training objective utilized.

3.1. Vertex Embedding Layer

Based on the observation that points close in space will share similar features [45], we propose the use of graph neural network to encode the nodes of the 3D shape, that consumes relative position features rather than raw features as traditionally used in networks such as PointNet [42]. Using a GNN we exploit topology priors that are essential for the task of shortest path estimation. Specifically, assuming two points on two distinct fingertips that are considered close in the Euclidean space, these points are actually distant on the surface curve and consequently in the graph domain. Thus, the scope of Vertex Embedding Layer (VEL) is to assign topology aware features node, compared to topology agnostic networks such as PointNet [42]. We model VEL using a stack of graph convolution layers as defined bellow.

3.1.1 Relative Multi-Aggregation Graph Convolution

Given that node embeddings are solely dependent on the expressive power of the GNN, we attempt to structure a graph operator that can effectively describe curved surfaces such as 3D meshes. The limitation of GNNs to expressively describe neighborhood information using a single aggregation function was initially explored in [9]. Based on such observations, we propose an expressive graph convolution operator, that is able to model neighborhood structure without being affected from outliers and local extremas. In particular, motivated by the works of [9, 23, 44], we introduce Relative Multi-Aggregation Graph Convolution (RMAGC) layer, with a node feature update rule defined as:

$$\mathbf{f}_{i}^{(l)} = \left[|| \Box_{k \in \mathcal{N}_{i}} \mathbf{W}_{\mathbf{n}} \left(\mathbf{f}_{i}^{(l-1)} - \mathbf{f}_{j}^{(l-1)} \right) \right]$$
(2)

where $\mathbf{f}_i^{(l)}$ is the feature of node \mathbf{v}_i in layer l, \mathbf{W}_n layer parameters, \mathcal{N}_i the neighbors of node \mathbf{v}_i , \Box the aggregation function and || the concatenation operator. Note that by using Cartesian point position features the vertex embedding layer can only learn node descriptors aligned with the training set shapes. In contrast, the proposed relative feature formulation enables vertex embeddings decomposition from the shape topology. In the proposed setting, we utilized four stacked RMAGC layers with four aggregation functions: [mean, min, max, std].

3.2. Path Descriptor

Apart from the general graph structure and shape, we need to proper encode the path direction to avoid predictions that are not aligned with the current predicted path. In particular, loops around nodes along with misleading direction vectors can derail the path from the target point. To avoid that, we encode the prefix of the path up to the current node $\mathbf{p}^t = \mathbf{v}_i$ using the momentum formulation inspired from [26]:

$$\beta = \operatorname{sigmoid}([\mathbf{g}^{t-1} || \mathbf{f}_i^t]) \tag{3}$$

$$\mathbf{g}^{t} = \beta \mathbf{f}_{i}^{t} + (1 - \beta) \mathbf{g}^{t-1}$$
(4)

where \mathbf{g}^t is the path descriptor in time step t and \mathbf{f}_i^t the encoding of node \mathbf{v}_i in time step t.

3.3. Point Selector

The point selector module acts as a classification model that selects the next point in the path from the neighborhood of the current point \mathbf{p}^t . Given the vertex embeddings \mathbf{F} , the path descriptor \mathbf{g}^t , we formulate an expressive way to represent the current point's neighborhood so that to aid the point selector module to decide the best fit in the path. In order to do so, we utilize a set of relative to the target point features. In particular, we define the relative features of point \mathbf{v}_j in the neighborhood of \mathbf{v}_i as:

$$\mathbf{r}_{j} = [\mathbf{v}_{tgt} - \mathbf{v}_{i} || \operatorname{cosine}(\mathbf{v}_{tgt} - \mathbf{v}_{j}, \mathbf{v}_{j} - \mathbf{v}_{i}) || \\ \operatorname{cosine}(\mathbf{v}_{tgt} - \mathbf{v}_{j}, \mathbf{v}_{tgt} - \mathbf{v}_{str}) || d(\mathbf{v}_{j}, \mathbf{v}_{tgt})]$$
(5)



Figure 1. Architecture of the proposed method. a) At each time t, the Point Selector module decides the next point in the path sequence given the node embeddings (Sec. 3.1), the path descriptor (Sec. 3.2) and the neighborhood relative point features (Sec. 3.3). b) Node embeddings produced by the Vertex Embedding Layer (green) introduce a bias to the neighborhood features (purple) and assist the Point Selector decide the next point in the path sequence.

where $\mathbf{v}_{str}, \mathbf{v}_{tgt}$ are the starting and the target points, $cosine(\cdot)$ denotes the cosine similarity function and $d(\cdot)$ the Euclidean distance operator. Similar to the vertex embedding layer, we use relative features to enable generalization to different shapes and topologies. We additionally utilize cosine similarity features to capture the path directions and their alignment with the directional vector between starting and ending points.

Finally, the probability that point v_j , in the neighborhood of v_i , will be the next point in the path sequence is be defined as:

$$\mathcal{P}(\mathbf{p}^{t+1} = \mathbf{v}_j | p^t, \mathcal{M}) = \texttt{softmax}(\texttt{MLP}([r_j || f_j || g^t]))$$
 (6)

where \mathcal{M} is the mesh that \mathbf{v}_i lies and MLP denotes a multilayer perceptron. In this paper we used a 5 layer MLP accompanied with ReLU activations.

3.4. Loss Functions

To train the proposed method we utilize a combination of two objective functions, namely the point selector cross entropy loss and a probabilistic geodesic distance loss.

3.4.1 Point Selection Classification Loss

The main scope of the point selector classification loss is to maximize the probability that, in every step of the path, the correct next point is selected by the Point Selector. In order to do so, we make use of a binary cross entropy loss \mathcal{L}_c to maximize the raw point probabilities of the Point Selector module according to the ground truth points in the path. Additionally, to enforce training efficiency we follow the teacher forcing strategy, where we feed the model with the ground truth path points after each prediction.

3.4.2 Probabilistic Geodesic Distance Loss

Given that there is no guarantee for a unique shortest path between two points, we utilize an additional loss function to ensure that, in every step, the geodesic distance is minimum. With this formulation, we reinforce the network to make predictions that minimize the distance between the current point and the target vertex. In mathematical parlance, we define the probabilistic geodesic loss as:

$$\mathcal{L}_d = \sum_{x,y,z} (\sum_{j \in \mathcal{N}_i} p_j \mathbf{v}_{ij} - \mathbf{v}_t)^2 \tag{7}$$

where \mathbf{v}_t is the ground truth vertex in step t, \mathbf{v}_{ij} the *j*-th neighbor of current vertex \mathbf{v}_i and p_j its corresponding selection probability as assigned by the Point Selector module.

3.5. Beam Search Inference

Trying to further aid the model in abstaining from commiting "short-sighted" errors during inference, we observed that the selection of the next node in the shortest path be compared with the selection of the next token in Natural Language Generation (NLG). Thus, inspired by the de facto algorithm used for decoding in NLG problems, we adjust the beam search algorithm to operate with node probabilities. Beam search is a form of pruned breadth-first search where the breadth is limited to $k \in \mathbb{Z} + [32]$. It is a timesynchronous approximate search algorithm [51] which increases total prediction step complexity by a constant multiplicative factor equal to the chosen breadth k. Experimentally we choose a breadth of k = 3 beams as the optimal accuracy-complexity tradeoff.

4. Experiments

In this section we evaluate the shortest paths generated by the proposed framework. We initially examine the generated path convergence and geodesic distances produced by our method. Additionally, we evaluate performance speedup of our module compared to traditional methods and showcase the contribution of each model module using an ablation study.

Dataset. To train our model we utilized the TOSCA dataset [3] that contains 80 unique shapes under different non-rigid transformations. Motivated by the diversity of the shapes and the enormous amount of rough regions they include we considered TOSCA dataset as a perfect fit for the task of geodesic path modeling. We randomly selected 60 shapes for training and 20 unique shapes for testing. To generate the dataset, we randomly selected 500 pairs of vertices for each shape and generated the ground truth paths using Bellman-Ford algorithm.

Baselines. For comparison, we selected several shortest path estimation baselines with different properties. Initially, we implemented two baseline methods that follow the same inference structure as the proposed method, without however having any learnable module. In particular, at each time step of the shortest path estimation a point from the neighborhood of the current point is selected as the next point in the path based on its distance to the target end point (Euclidean Distance) or its cosine similarity directional vector (Cosine Similarity). To be more precise, the first method selects the point v_i on the neighborhood of v_i that is closer, in terms of euclidean distance, to the target point v_{tqt} , whereas the second method selects the point that maximizes the cosine similarity, $\cos(v_i - v_i, v_{tqt} - v_j)$. Additionally, we compare the proposed method with the geodesic path estimation algorithm that is based on edge flips proposed in [49], the probabilistic shortest path estimation algorithm as described in [30], along with the recently proposed learnable geodesic distance estimation method [43]. We avoid comparing with [47], since it requires training Node2Vec embeddings for every single mesh topology, which cannot be obtained in an end-to-end framework.

4.1. Geodesic Paths and Distances

To assess the performance of the proposed method, we measure the geodesic distance error (GDE) on the test set shapes of TOSCA dataset. Note that the topologies we used to evaluate the model are not present in the training set. We measure the GDE between the predicted and the ground truth paths as:

$$GDE = \frac{|d_{pr} - d_{gt}|}{d_{qt}} \tag{8}$$

where d_{pr}, d_{gt} are the predicted and the ground truth geodesic distances respectively.

Table 1. Quantitative evaluation of the proposed and baseline methods.

Method	Convergence (%)	GDE
Cosine Similarity	38	0.39
Euclidean Distance	42	0.35
Sharp <i>et al</i> . [50]	24	0.37
Lim <i>et al.</i> [30]	79	0.38
Qi et al. [43]	-	0.88
Proposed	82	0.17

We also assess the convergence of each method in terms of steps required to reach the target point. In particular, for each method that predicts shortest path approximation, we measured the percentage of paths converged to the target point with up to 20% more steps than the shortest path.

As shown in Table 1, the proposed method attains to achieve an increased convergence percentage along with a relatively small geodesic distance error, compared to the baseline methods. Significantly, apart from the method of Lim et al. [30] all other baseline methods fail to consistently estimate paths that require similar steps to the ground truth shortest path. Additionally, as expected, methods that directly regress geodesic distances can not be easily adapted to unique mesh topologies and thus produce increased geodesic error. Such findings can be also qualitatively validated in Figure 2, where the proposed method is contrasted with the baselines under different shape topologies. It can be easily observed that the proposed method achieves to estimate accurate shortest path even in challenging shapes. In particular, in the case of the armadillo model (second row) as well as the Simpsons Homer model (fourth row), naive estimators such as the Cosine Similarity and the Euclidean distance get trapped in local minima and fail to converge. In contrast, the proposed method attains minimum geodesic distance error by selecting paths driven by the VEL and the path descriptor. Note the method of Sharp et al. [50], breaks the mesh connectivity and thus produces paths that are not aligned with the original mesh topology.

Additionally, we estimated the shortest paths from a single point to all points of the mesh and measured their geodesic error compared to the ground truth paths. In Figure 3, we qualitatively assess the geodesic error of the generated geodesic paths. As can be easily seen, the proposed method generates paths that respect the intrinsic property of the mesh and achieves the smallest geodesic error across the whole mesh surface.

4.2. Runtime

A main challenge of modern applications of shortest path and geodesic distance calculations is real-time execution



Figure 2. Qualitative comparison of the proposed and baseline methods. The ground truth shortest path is shown in black (left column) and the generated ones in red. Black dots denote start and end points.



Figure 3. Geodesic distance error color-coded on the mesh surface. Blue color denotes small GDE whereas red color denotes increased error.

of the algorithm. Apart from accurate and plausible solutions, a key property of a geodesic path estimation method

is to diminish the runtime of traditional greedy algorithms. Taking advantage of the progress of modern GPUs, neural networks have become vastly popular for their precision as real-time approximators. To compare the proposed and the baseline methods we constructed an experiment with a total of 50 paths and measured the total runtime each method required to generate the solutions. Note that we have not segregated CPU from GPU methods since we are interested solely in the total runtime required. In Figure 4 we report the average runtimes on 20 executions along with their respective standard deviations. Regarding the implementation



Figure 4. Runtime comparison between the proposed and baseline methods.

of choice for the traditional shortest path algorithms we selected for our baselines i.e. Bellman-Ford and Dijkstra, past research has aimed to optimize the algorithms' runtime on modern GPUs, as per [4, 34, 35]. However, such optimized prototypes have not found their spot in modern data science workflows due to the difficulty of incorporating them into the codebases of modern data science libraries, a fact often attributed to the absence of code for reproduction of their implementation and experiments. Therefore, we use the immensely popular library SciPy [53] as the source for our primitive algorithm baselines. While there is no explicit support of SciPy for GPU computing, its functions are highly optimized to run on modern CPUs as it uses Python wrappers around algorithms coded in C or C++. As seen in Figure 4, the proposed method runs on par with or faster than all compared baselines, both differentiable and purely algorithmic. Compared with traditional shortest path algorithms, our method scales down the runtime by up to 9 times, setting itself ready performance-wise for integration in broader learning frameworks. Furthermore, the proposed method attains to achieve efficient performance both in terms of GDE and convergence percentage as well as runtime execution compared to Euclidean Distance and Sharp et al. [50] methods.

4.3. Shortest Path Estimation on Point Clouds

As an attempt to evaluate the performance of the proposed method in extreme cases where the explicit 3D shape topology is absent, we attempted to estimate shortest paths on unstructured point clouds. In particular, we fed to the model point clouds where the connectivity is attained using k-nearest neighbors (k-NN) and we compared the predicted paths with the ground truth paths, obtained from the original mesh connectivity. We structure the neighborhood of each point by using five nearest neighbors based on their Euclidean distance. As can be seen in Figure 5, the proposed



Figure 5. Generalization of the proposed method on an unstructured point cloud. Predicted geodesic paths (green) achieve a GDE of 0.21 on the armadillo 3D shape compared to the ground truth paths (red).

method converges to the target point with a small amount of extra steps and achieves remarkable GDE compared to the ground truth paths.

Additionally, we examined the case of point neighborhood estimation on the armadillo point cloud. In particular, using the k-NN connectivity as described earlier, we estimated the geodesic distances from a fingertip to all other points of the point cloud and retained the ones that are less than a given threshold. In Figure 6, we compare the geodesic neighborhood of a finger tip point estimated by the proposed method and the k-NN method, that is the common practice in point cloud analysis task. It can be easily observed that using the proposed method we achieved geodesic neighborhood estimates that are very similar to the ground truth neighborhood of the point. On the contrary, k-NN method that is based purely on Euclidean point distances, fails to distinguish different fingers, which is particularly useful to tasks such as part segmentation.

4.4. Ablation Study

We performed an ablation study to illustrate the effectiveness of the key components of the proposed method. In



Figure 6. Color-coded neighborhood points estimation. From left to right we visualize points that are 0.1, 0.2 and 0.5 away from the finger tip. Results are illustrated on top of the structured mesh for better visualization.

particular, we assessed the importance of both three major modules of the proposed network along with the impact of each loss functions to the method's performance. Experimental results on the test set are summarized in Table 2. The importance of embedding node features and graph topology to estimate paths on the surface can be quantified with the removal of vertex embedding layer (VEL). As expected, the absence of VEL has a large impact on the path convergence leading to a significant decrease in performance. Similarly, a performance drop can be observed when we substitute relative features, as introduced in 3.3, with (x, y, z)-coordinate features. Using directly (x, y, z)-coordinate features the model can not generalize well to unseen figures and paths, whereas utilizing relative feature we model a, much easier, space of deformations, that aids network generalization capabilities. Experimental results also reveal that path descriptor is also important in the process of path estimation. Specifically, the momentum strategy introduced in the path descriptor boost the point selector to select points of the neighborhood that follow the pattern of the prefix path avoiding loops and derailed paths.

Finally, the contribution of each loss function was examined by training the proposed network using each loss individually. The two losses can be considered as a direct and indirect way of training the network, respectively. Classification loss impacts the point selector generated probabilities directly based on the ground truth points whereas the geodesic loss indirectly guides network predictions to minimize the distance to the target point. Using only the geodesic loss without the classification loss, the model tends to imitate the Euclidean Distance baseline where decisions are made purely by the distance to the end goal. How-

Table 2. Ablation Study: Quantitative evaluation of the proposed method key components.

Method	Convergence (%)	GDE
- w/o VEL	57	0.35
- w/o Path Descriptor	71	0.23
- w/o Relative Features	66	0.29
- w/o Classification Loss	53	0.34
- w/o Geodesic Loss	78	0.24
Proposed - Full	82	0.17

ever, its contribution is not negligible given that it enforces training of certain parts of the network such as the path descriptor.

5. Limitations and Conclusion

In this paper, we presented the first, to the best of our knowledge, geodesic path estimator that is based on neural networks. Motivated by the lack of differentiable and realtime estimation methods along with the success of neural networks as approximators, we implemented a framework that learns to approximate the shortest path problem. In particular, we model the shortest path problem using an autoregressive model that given the current point, its neighboring points and the prefix of the path, decides the next points in the path sequence. The proposed method is comprised by three main modules namely the vertex embedding layer that generates node embeddings given the mesh topology, a path descriptor that models the prefix path encodings and the point selector that given such information chooses the next point in the geodesic path. We evaluate the proposed method under a test set of diverse shapes and we showcase its superiority compared to other baseline approximation methods. Additionally, we explore the generalization capabilities of the proposed method on unstructured point clouds and illustrate that it can be directly used to point clouds without being significantly affected by the lack of point connectivity.

However, the proposed method is strongly attached to the approximation quality of the neural network, which sometimes limits the generated shortest paths. Specifically, in extreme scenarios that have not been explored in the training set and in cases where the shortest path is not obvious, the approximation quality of neural networks might be strictly bounded. To tackle such limitations, we plan to extend the proposed method to generate more robust and expressive features that will break the approximation burdens of neural networks. Finally, we believe that the rationale underlying the proposed method and the presented findings of this paper will not only benefit the 3D computer vision but also the geometric deep learning community.

References

- Prosenjit Bose, Anil Maheshwari, Chang Shu, and Stefanie Wuhrer. A survey of geodesic paths on 3d surfaces. *Computational Geometry*, 44(9):486–498, 2011. 2, 3
- [2] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Efficient computation of isometry-invariant distances between surfaces. *SIAM Journal on Scientific Computing*, 28(5):1812–1836, 2006. 2
- [3] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 5
- [4] Amadou Chaibou, Oumarou Sie, et al. Improving global performance on gpu for algorithms with main loop containing a reduction operation: case of dijkstra's algorithm. *Journal of Computer and Communications*, 3(08):41, 2015. 7
- [5] Anthony Chen and Zhaowang Ji. Path finding under uncertainty. *Journal of advanced transportation*, 39(1):19–37, 2005.
- [6] Liyu Chen and Haipeng Luo. Finding the stochastic shortest path with low regret: The adversarial cost and unknown transition case. In *International Conference on Machine Learning*, pages 1651–1660. PMLR, 2021. 2
- [7] J. Cong, A.B. Kahng, and Kwok-Shing Leung. Efficient algorithms for the minimum shortest path steiner arborescence problem with applications to vlsi physical design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(1):24–39, 1998. 1
- [8] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, 2nd edition, 2001. 2
- [9] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*, 33:13260–13271, 2020. 3
- [10] Keenan Crane, Marco Livesu, Enrico Puppo, and Yipeng Qin. A survey of algorithms for geodesic paths and distances. arXiv preprint arXiv:2007.10430, 2020. 1
- [11] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2018. 1
- [12] E.W. DIJKSTRA. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. 2
- [13] James R. Driscoll, Harold N. Gabow, Ruth Shrairman, and Robert E. Tarjan. Relaxed heaps: An alternative to fibonacci heaps with applications to parallel computation. *Commun. ACM*, 31(11):1343–1354, nov 1988. 2
- [14] Willem Feijen and Guido Schäfer. Using machine learning predictions to speed-up dijkstra's shortest path algorithm. *arXiv preprint arXiv:2112.11927*, 2021. 2
- [15] Robert W. Floyd. Algorithm 97: Shortest path. Commun. ACM, 5(6):345, jun 1962. 2
- [16] Daniel Foead, Alifio Ghifari, Marchel Budi Kusuma, Novita Hanafiah, and Eric Gunawan. A systematic literature review of a* pathfinding. *Procedia Computer Science*, 179:507– 514, 2021. 2

- [17] Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. J. ACM, 34(3):596–615, jul 1987. 2
- [18] Giorgio Gallo and Stefano Pallottino. Shortest path algorithms. Annals of operations research, 13(1):1–79, 1988. 2
- [19] Andrew V. Goldberg and Renato F. Werneck. Computing point-to-point shortest paths from external memory. In *ALENEX/ANALCO*, 2005. 2
- [20] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016. 1, 3
- [21] Andrey Gubichev, Srikanta Bedathur, Stephan Seufert, and Gerhard Weikum. Fast and accurate estimation of shortest paths in large graphs. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, page 499–508, New York, NY, USA, 2010. Association for Computing Machinery. 1, 2
- [22] Ronald J. Gutman. Reach-based routing: A new approach to shortest path algorithms optimized for road networks. In Lars Arge, Giuseppe F. Italiano, and Robert Sedgewick, editors, ALENEX/ANALC, pages 100–111. SIAM, 2004. 2
- [23] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. Advances in neural information processing systems, 30, 2017. 3
- [24] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. 1, 2
- [25] Bruce Hopenfeld, Hiroshi Ashikaga, and Elliot R McVeigh. Geodesic based registration of sensor data and anatomical surface image data. *Annals of biomedical engineering*, 35(10):1771–1781, 2007. 1
- [26] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 2, 3
- [27] Ron Kimmel and James A Sethian. Computing geodesic paths on manifolds. *Proceedings of the national academy* of Sciences, 95(15):8431–8435, 1998. 2
- [28] Nicolas Lefebvre and Michael Balmer. Fast shortest path computation in time-dependent traffic networks. *Arbeitsberichte Verkehrs-und Raumplanung*, 439, 2007. 1
- [29] Didong Li and David B Dunson. Geodesic distance estimation with spherelets. arXiv preprint arXiv:1907.00296, 2019.
 2
- [30] Sejoon Lim, Christian Sommer, Evdokia Nikolova, and Daniela Rus. Practical route planning under delay uncertainty: Stochastic shortest path queries. In *Robotics: Science and Systems*, volume 8, pages 249–256. United States, 2013. 2, 5
- [31] Parimala Mani, Biju Vasudevan, and Murali Sivaraman. Shortest path algorithm of a network via picture fuzzy digraphs and its application. *Materials Today: Proceedings*, 45:3014–3018, 2021. 2
- [32] Clara Meister, Tim Vieira, and Ryan Cotterell. If beam search is the answer, what was the question? *arXiv preprint arXiv:2010.02650*, 2020. 4

- [33] Joseph SB Mitchell et al. Geometric shortest paths and network optimization. *Handbook of computational geometry*, 334:633–702, 2000. 3
- [34] U. Supriya Nayak, Rajeev Pandey, and Uday Chourasia. A comparative study on analysis of various shortest path algorithms on gpu using opencl. *International Journal of Computer Applications*, 179:34–38, 2018. 7
- [35] Hector Ortega-Arranz, Yuri Torres, Diego R Llanos, and Arturo Gonzalez-Escribano. A new gpu-based approach to the shortest path problem. In 2013 International Conference on High Performance Computing & Simulation (HPCS), pages 505–511. IEEE, 2013. 7
- [36] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deep-Walk. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, aug 2014. 1
- [37] Diego Ulisse Pizzagalli, Santiago Fernandez Gonzalez, and Rolf Krause. A trainable clustering algorithm based on shortest paths from density peaks. *Science advances*, 5(10):eaax3770, 2019. 2
- [38] Michalis Potamias, Francesco Bonchi, Carlos Castillo, and Aristides Gionis. Fast shortest path distance estimation in large networks. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, page 867–876, New York, NY, USA, 2009. Association for Computing Machinery. 2
- [39] Rolandos Alexandros Potamias, Giorgos Bouritsas, and Stefanos Zafeiriou. Revisiting point cloud simplification: A learnable feature preserving approach. arXiv preprint arXiv:2109.14982, 2021. 1
- [40] Rolandos Alexandros Potamias, Stylianos Ploumpis, and Stefanos Zafeiriou. Neural mesh simplification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2022. 1
- [41] Rolandos Alexandros Potamias, Jiali Zheng, Stylianos Ploumpis, Giorgos Bouritsas, Evangelos Ververas, and Stefanos Zafeiriou. Learning to generate customized dynamic 3d facial expressions. In *European Conference on Computer Vision*, pages 278–294. Springer, 2020. 1
- [42] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference* on computer vision and pattern recognition, pages 652–660, 2017. 3
- [43] Jianzhong Qi, Wei Wang, Rui Zhang, and Zhuowei Zhao. A learning based approach to predict shortest-path distances. 2020. 1, 3, 5
- [44] Ekagra Ranjan, Soumya Sanyal, and Partha Talukdar. Asap: Adaptive structure aware pooling for learning hierarchical graph representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5470–5477, 2020. 3
- [45] Ekagra Ranjan, Soumya Sanyal, and Partha Pratim Talukdar. ASAP: Adaptive structure aware pooling for learning hierarchical graph representations. *arXiv preprint arXiv:1911.07979*, 2019. 3
- [46] Matthew J. Rattigan, Marc Maier, and David Jensen. Using structure indices for efficient approximation of network

properties. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, page 357–366, New York, NY, USA, 2006. Association for Computing Machinery. 1, 2

- [47] Fatemeh Salehi Rizi, Joerg Schloetterer, and Michael Granitzer. Shortest path distance approximation using deep learning techniques. In 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pages 1007–1014. IEEE, 2018. 1, 3, 5
- [48] James Albert Sethian. Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science, volume 3. Cambridge university press, 1999. 2
- [49] Gil Shamai and Ron Kimmel. Geodesic distance descriptors. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 6410–6418, 2017. 1, 5
- [50] Nicholas Sharp and Keenan Crane. You can find geodesic paths in triangle meshes by just flipping edges. ACM Trans. Graph., 39(6), 2020. 2, 5, 7
- [51] Felix Stahlberg and Bill Byrne. On nmt search errors and model errors: Cat got your tongue? arXiv preprint arXiv:1908.10090, 2019. 4
- [52] Jin Tao Tang, Ting Wang, and Ji Wang. Shortest path approximate algorithm for complex network analysis. *Ruan Jian Xue Bao/Journal of Software*, 22(10):2279–2290, oct 2011.
 2
- [53] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020. 7
- [54] Lingkun Wu, Xiaokui Xiao, Dingxiong Deng, Gao Cong, Andy Diwen Zhu, and Shuigeng Zhou. Shortest path and distance queries on road networks: An experimental evaluation. arXiv preprint arXiv:1201.6564, 2012. 1
- [55] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 915–924, 2021. 1, 2
- [56] Bin Yang, Chenjuan Guo, Christian S Jensen, Manohar Kaul, and Shuo Shang. Stochastic skyline route planning under time-varying uncertainty. In 2014 IEEE 30th International Conference on Data Engineering, pages 136–147. IEEE, 2014. 2
- [57] Xiaohan Zhao and Haitao Zheng. Orion: shortest path estimation for large social graphs. In 3rd Workshop on Online Social Networks (WOSN 2010), 2010. 3