ISSN: 2519-9439

Topology-Calibrated Correspondence Transfer across Non-Isometric Object Families

[1]Sebastian Huber, [2]Charlotte Vogel

[1]Department of Computer Science, Faculty of Mathematics, Computer Science and Statistics, Ludwig-Maximilians-Universität München, Munich, Bavaria, Germany

[2]Institute of Computer Science, Faculty of Physics, Mathematics and Computer Science, Johannes Gutenberg University Mainz, Mainz, Rheinland-Pfalz, Germany

Abstract

Establishing meaningful point-to-point correspondences between distinct three-dimensional objects represents a foundational challenge in computer graphics, vision, and geometry processing. Traditional techniques have predominantly focused on isometric shapes, wherein intrinsic geodesic distances between points remain largely invariant despite extensive articulated deformations. However, real-world object families exhibit profound non-isometric variations characterized by differing topologies, structural additions, absent components, and significant scaling discrepancies. Such structural divergences severely disrupt conventional spectral and spatial alignment algorithms. This paper introduces a comprehensive framework for topology-calibrated correspondence transfer designed explicitly to bridge the semantic and geometric gap across highly non-isometric object families. By abstracting complex three-dimensional meshes into resilient topological skeletal graphs, the proposed methodology captures the essential structural hierarchy of an object independent of its localized geometric fluctuations. Through a meticulously designed alignment protocol, these skeletal abstractions serve as a robust scaffolding to guide the transfer of dense surface correspondences. The proposed framework integrates multi-scale geometric descriptors with global topological constraints, effectively mitigating the ambiguities that plague standard functional mapping techniques. Extensive evaluations conducted on challenging non-isometric datasets demonstrate that the proposed topology-calibrated approach significantly reduces alignment errors while preserving the structural integrity of the transferred maps. The findings underscore the critical importance of embedding global topological awareness into localized geometric matching processes, paving the way for more resilient structural analysis across diverse object categories.

Keywords: Shape Analysis, Non-Isometric Matching, Topological Calibration, Dense Correspondence, Geometry Processing

1.Introduction

1.1 Context and Motivation

The computational analysis of three-dimensional shapes relies heavily on the ability to determine accurate mappings or correspondences between the surfaces of different objects. This fundamental task, often referred to as shape matching, serves as the operational backbone for a multitude of downstream applications in digital geometry processing, ranging from statistical shape modeling and texture transfer to object recognition, interpolation, and

ISSN: 2519-9439

character animation. The primary objective is to identify semantically and geometrically analogous regions across varying instances of objects. For decades, the research community has made exceptional strides in resolving correspondence problems for objects that undergo isometric or near-isometric deformations [1]. In an isometric setting, the intrinsic metric properties of the shape, specifically the geodesic distances measured along the surface manifold, are preserved regardless of the object's spatial pose or articulation. A classic example is a human figure adopting various postures; the distance along the surface of the skin from the shoulder to the wrist remains constant whether the arm is raised or lowered. Consequently, mathematical frameworks formulated around intrinsic surface properties, such as the eigenfunctions of the surface differential operators, have proven extraordinarily successful in capturing and matching these invariant signatures [2, 3]. Despite these successes, the isometric assumption imposes a severe limitation on the applicability of shape matching algorithms in wild, unconstrained environments. Natural object families, even those belonging to the exact same semantic category, rarely adhere to strict isometry. Consider the category of four-legged animals: a giraffe, a hippopotamus, and a domestic cat share a fundamental quadrupedal semantic structure, yet their geometric realizations are vastly disparate. The relative lengths of their appendages, the thickness of their torsos, and the presence or absence of specific anatomical features such as tails or horns result in severe non-isometric distortions [4]. These variations drastically alter the intrinsic metric of the surfaces, causing traditional distance-preserving and spectral-based matching algorithms to fail unpredictably. The motivation for this research stems from the urgent necessity to transcend the restrictive boundaries of isometry. A robust shape matching paradigm must possess the capacity to interpret and align objects based on their underlying structural and semantic essence, effectively ignoring the distracting noise of non-isometric geometric variations. This requires a paradigm shift from purely continuous metric spaces to structural representations that encode the topological configuration of the object.

1.2 Problem Statement

The core problem addressed in this study is the reliable transfer of dense surface correspondences between pairs of three-dimensional meshes that exhibit significant non-isometric variations and topological inconsistencies. Given a source shape and a target shape belonging to the same broad object family but differing substantially in local geometry and global proportions, the objective is to establish a continuous, point-to-point mapping that is both semantically meaningful and geometrically smooth. The presence of non-isometric transformations introduces multiple layers of complexity. Firstly, structural discrepancies such as the addition of extraneous parts or the absence of expected components lead to discontinuous mappings where localized regions on the source shape have no valid counterpart on the target shape. Secondly, extreme anisotropic scaling, such as the elongation of a limb, distorts the local area and angular relationships, rendering local geometric descriptors highly unreliable [5, 6]. Standard functional mapping paradigms attempt to bypass spatial complexities by aligning shapes in a reduced spectral domain spanned by a truncated basis of functions. However, the basis functions themselves are acutely sensitive to topological changes. When an object undergoes non-isometric alteration, the ordering and the spatial frequencies of its structural harmonics become jumbled, causing the spectral alignment to collapse. The challenge, therefore, lies in developing an intermediary representation that remains invariant to these non-isometric disruptions. The problem demands a methodology that can calibrate the correspondence transfer process using a stable structural anchor. Without such an anchor, algorithms inevitably fall into localized minima, mismatching semantically disparate regions that happen to share superficial geometric similarities, such as erroneously mapping the tip of an ear to the tip of a tail simply because both exhibit high Gaussian curvature.

ISSN: 2519-9439

1.3 Contributions and Organization

This paper presents a novel framework termed topology-calibrated correspondence transfer, which systematically disentangles structural topology from localized surface geometry to enable robust matching across non-isometric object families. The methodology leverages a dual-layered approach, initially extracting a one-dimensional skeletal graph from the three-dimensional mesh to represent the global topological layout, and subsequently using this graph to guide the dense continuous mapping process. The primary contributions of this work are manifold. First, it introduces a highly resilient structural extraction algorithm that reliably identifies the semantic extremities and branching nodes of complex shapes despite severe geometric noise. Second, it proposes a graph-based alignment strategy that establishes a sparse set of highly confident anchor correspondences, completely bypassing the spectral instabilities inherent in purely functional approaches. Third, it develops a continuous refinement mechanism that propagates these sparse topological anchors into a dense surface mapping using a topology-aware diffusion process. The remainder of this academic paper is structured as follows. Section 2 provides a comprehensive literature review, critically examining the historical trajectory of shape analysis, the limitations of current isometric techniques, and the emergence of topology-driven methods. Section 3 details the proposed methodology, offering an in-depth breakdown of the structural representation, the topological alignment strategy, and the dense correspondence transfer mechanism. Section 4 outlines the implementation details and the rigorous experimental design utilized to validate the framework. Section 5 presents a thorough analysis of the experimental results, encompassing quantitative performance metrics, qualitative visual assessments, and detailed ablation studies. Finally, Section 6 summarizes the core findings, discusses the inherent limitations of the current framework, and proposes actionable directions for future research in the domain of non-isometric shape analysis.

2. Literature Review and Theoretical Background

2.1 Isometric Shape Analysis

The pursuit of establishing correspondences between three-dimensional surfaces has a rich history deeply rooted in differential geometry. Early approaches predominantly focused on rigid alignment, utilizing algorithms like iterative closest point matching to minimize spatial distances between pre-aligned point clouds. However, as the focus shifted toward articulated and deformable objects, intrinsic methods gained prominence. The fundamental principle guiding isometric shape analysis is the preservation of geodesic distances [7]. If a shape undergoes a transformation that merely bends the surface without stretching or tearing it, the shortest path along the surface between any two points remains constant. This realization led to the development of multidimensional scaling techniques, which attempt to flatten the intrinsic geometry of shapes into a canonical Euclidean space where standard rigid matching algorithms can be applied. The introduction of spectral geometry marked a monumental leap forward in this domain. By analyzing the properties of differential operators defined on the surface manifold, researchers discovered that the eigenvalues and eigenfunctions of these operators form a rich, deformation-invariant signature of the shape. Heat kernel signatures and wave kernel signatures became standard tools, providing a robust description of the local and global geometry surrounding a point based on the dissipation of heat or the propagation of waves across the surface [8, 9]. These spectral descriptors are highly effective for isometric matching because the underlying differential operator is strictly dependent on the intrinsic metric tensor, which remains invariant under isometric deformation. However, these classical spectral methods operate under an assumption of complete and perfect surfaces, struggling significantly when confronted with topological noise, holes, or boundary modifications that alter the boundary conditions of the differential operators [10].

ISSN: 2519-9439

2.2 Non-Isometric Correspondence Techniques

As the limitations of purely isometric assumptions became apparent, the research community pivoted toward accommodating non-isometric deformations. One of the most significant breakthroughs in this area was the introduction of the functional map framework. Instead of seeking a direct point-to-point mapping between the spatial coordinates of two shapes, functional maps attempt to align the functional spaces defined on the shapes [11, 12]. By projecting functions, such as indicator functions or geometric descriptors, onto a truncated basis of surface harmonics, the complex problem of spatial matching is reduced to finding a small, linear transformation matrix that aligns the spectral coefficients. This approach is highly elegant and computationally efficient, provided that the underlying basis functions of the source and target shapes are compatible. Unfortunately, the compatibility of spectral bases degrades rapidly in the presence of non-isometric transformations [13]. When an object undergoes structural alterations, such as extreme stretching of a specific part or the addition of topological handles, the spatial frequencies represented by the structural harmonics shift unpredictably. This phenomenon, known as spectral leakage, breaks the foundational assumption of functional maps, leading to catastrophic alignment failures [14]. To counteract this, researchers have explored various regularization techniques, incorporating spatial constraints, promoting map continuity, or utilizing localized spectral bases. Recent advancements have increasingly leveraged deep learning architectures, employing geometric convolutional neural networks to extract highly robust, hierarchical features from the raw mesh data [15, 16]. While learning-based methods have demonstrated impressive resilience to certain types of noise, they often struggle to generalize across vastly different object categories and remain heavily reliant on massive quantities of annotated training data, which is notoriously difficult to procure for complex three-dimensional tasks [17].

2.3 Topological Calibration in Shape Matching

Recognizing the fragility of purely geometric or spectral alignment under non-isometric conditions, recent literature has witnessed a growing interest in incorporating topological priors into the correspondence pipeline. Topology, in this context, refers to the fundamental connectivity and structural layout of the object, stripping away the localized geometric details to reveal a simplified abstract representation. Techniques derived from computational topology, such as persistent homology, have been employed to identify stable topological features that persist across varying scales and deformations [18, 19]. These features can serve as robust landmarks to guide the matching process. A particularly effective approach involves the extraction of skeletal structures or curve skeletons from the three-dimensional volume. Skeletons provide an intuitive, one-dimensional abstraction of a shape, explicitly capturing its branching structure and semantic extremities. By matching the topological nodes and edges of these skeletons, researchers can establish a high-level structural correspondence that is remarkably resilient to non-isometric geometric distortions [20]. However, the primary challenge lies in the extraction process itself; traditional skeletonization algorithms are notoriously sensitive to surface noise and localized geometric fluctuations, often generating spurious branches that complicate the matching process. Furthermore, transitioning from a sparse skeletal alignment back to a dense surface correspondence remains a non-trivial inverse problem. The framework proposed in this paper builds upon these topological concepts, introducing a highly robust extraction methodology and a novel diffusion mechanism that seamlessly bridges the gap between abstract structural alignment and continuous, high-resolution surface mapping.

3. Methodology for Topology-Calibrated Transfer

3.1 Structural Representation of Object Families

The foundational step in the proposed topology-calibrated correspondence transfer framework is the abstraction of the complex, high-resolution three-dimensional mesh into a robust, one-

ISSN: 2519-9439

dimensional structural graph. This abstraction must fulfill two critical criteria: it must accurately capture the semantic hierarchy and branching topology of the object, and it must remain completely invariant to non-isometric geometric distortions, localized noise, and variations in mesh resolution. To achieve this, the framework eschews traditional volumetric thinning algorithms, which are prone to generating spurious branches, in favor of a purely intrinsic, function-driven approach. The process begins with the computation of a scalar field over the surface manifold that encodes the centrality and extremity of every vertex. This field is derived from the average geodesic distance from each point to all other points on the shape. Vertices with the highest average distances correspond to the semantic extremities of the object, such as the tips of fingers, the ends of tails, or the peaks of ears, while vertices with the lowest average distances reside in the core central volume of the shape [21]. By analyzing the critical points of this scalar field—specifically its local maxima and minima—the algorithm reliably isolates the foundational structural landmarks of the object. These critical points are then utilized as the initial seed nodes for the structural graph. To determine the connectivity between these nodes, the methodology traces the steepest descent paths along the gradient of the scalar field, linking the extremity nodes back to the central core. This process naturally delineates the branching structure of the object. To ensure topological stability, a rigorous pruning mechanism is applied. Branches that do not correspond to highly prominent, persistent structural features are collapsed, utilizing a persistence threshold based on the topological prominence of the corresponding critical points. The resulting structural graph comprises a set of semantic nodes, representing joints and extremities, connected by edges that represent the interconnecting structural segments. This graph serves as an unshakeable topological anchor, providing a highly compressed yet semantically rich description of the object's global architecture [22, 23].

3.2 Topological Alignment Strategy

With the structural graphs extracted for both the source and target objects, the next phase involves establishing a robust topological alignment between them. This step is critical, as it provides the high-confidence sparse correspondences that will subsequently calibrate the dense surface transfer. Unlike dense surface matching, which involves thousands or millions of vertices, graph alignment operates on a vastly reduced search space, typically comprising fewer than fifty semantic nodes. However, aligning graphs across non-isometric object families presents its own unique challenges. The graphs may not be strictly isomorphic; structural additions on one object may result in extra nodes and branches that have no counterpart on the other object [24]. The proposed alignment strategy models the graph matching problem as an energy minimization task. The objective function is designed to balance structural similarity, topological connectivity, and semantic consistency. To measure structural similarity, the algorithm computes localized topological descriptors for each node, capturing the valence of the node, the relative lengths of the connecting edges measured intrinsically along the surface, and the hierarchy of the node within the overall graph structure. The topological connectivity term penalizes mappings that violate the adjacency relationships defined by the edges of the graph, ensuring that connected nodes on the source shape map to connected nodes on the target shape, preserving the spatial continuity of the structure [25, 26]. To handle the non-isomorphic nature of the graphs, the alignment algorithm incorporates a robust outlier rejection mechanism. It allows for partial matchings, permitting certain nodes in either graph to remain unmatched if they represent structural anomalies or missing parts. The optimization is performed using a stochastic search algorithm that iteratively refines the node-to-node assignments while maximizing the overall consistency of the mapping. The output of this stage is a set of highly reliable, sparse correspondences anchoring the critical semantic joints and extremities of the source object directly to their most structurally analogous counterparts on the target object.

ISSN: 2519-9439

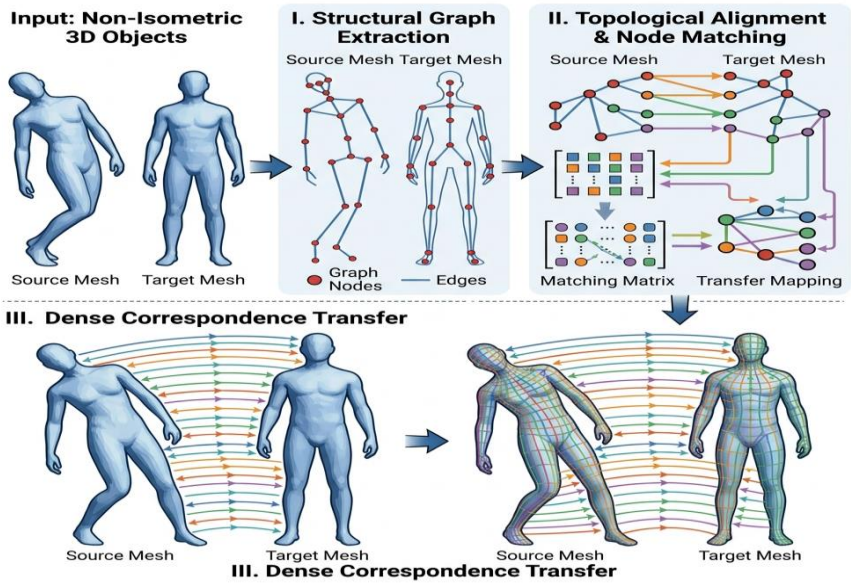


Figure 1Comprehensive Schematic of Topology

3.3 Correspondence Transfer Mechanism

The final and most intricate phase of the methodology is the transfer of the sparse topological alignment into a continuous, dense, point-to-point correspondence across the entire surface manifolds of the objects. Standard interpolation techniques often fail in this context, as they do not respect the complex intrinsic geometry of the surfaces and can result in mappings that tear or fold the surface unpredictably. To circumvent this, the proposed framework employs a topology-aware diffusion mechanism that propagates the sparse anchor information across the surface while strictly adhering to the intrinsic metric. The process leverages the continuous functional framework but crucially restricts its spectral basis using the topological anchors established in the previous step. The sparse node correspondences are first mapped back to their originating vertices on the high-resolution meshes. These vertices act as Dirichlet boundary conditions for a modified heat diffusion process over the surface manifold. By simulating the diffusion of thermal energy originating solely from these anchored vertices, the algorithm computes a set of structural coordinates for every point on the mesh. These coordinates describe the relative geodesic proximity of every surface point to the established topological landmarks [27]. Because the topological anchors represent the exact same semantic locations on both the source and target shapes, the resulting structural coordinates provide a geometrically invariant embedding space. The dense correspondence is then achieved by matching points across the two shapes that exhibit highly similar structural coordinate signatures. To ensure maximum map continuity and smooth transitions across structural boundaries, a final localized refinement step is applied. This involves regularizing the mapping using a specialized vector field that aligns the local tangent spaces of the matched points, effectively ironing out any micro-level inconsistencies or high-frequency noise in the mapping. The culmination of this intricate process is a dense correspondence map that is globally constrained by the structural topology and locally refined by the intrinsic geometry, successfully bridging the non-isometric divide.

4. Implementation and Experimental Design

4.1 Dataset Preparation and Annotation

To rigorously evaluate the efficacy of the proposed topology-calibrated correspondence transfer framework, comprehensive experimental protocols were established utilizing widely recognized, challenging benchmark datasets within the shape analysis community. The primary evaluations were conducted on advanced repositories specifically designed to test the limits of

ISSN: 2519-9439

non-isometric matching algorithms. The selected datasets comprise collections of diverse semantic categories, predominantly focusing on articulated subjects such as various animal species and human figures exhibiting vastly different body proportions, morphological structures, and topological noise. Each model within the dataset is represented as a high-resolution triangular mesh, typically containing between twenty thousand and fifty thousand vertices to ensure sufficient geometric detail for comprehensive structural analysis [28].

Preparation of the experimental data involved rigorous normalization procedures to eliminate superficial scale discrepancies and spatial translations. All meshes were intrinsically scaled such that their total surface area equaled a unit constant, and they were translated to position their center of mass at the origin of the global coordinate system. To establish ground truth for quantitative evaluation, a meticulous manual annotation process was undertaken. Expert annotators identified and linked a dense network of corresponding semantic points across intra-category models. In cases involving severe non-isometry where direct point-to-point mapping was ambiguous, barycentric coordinates within localized semantic regions were utilized to provide a continuous ground truth mapping field. Furthermore, synthetic topological noise, including localized surface holes, non-manifold edges, and artificial geometric extrusions, was intentionally injected into a subset of the models to stress-test the robustness of the structural extraction phase [29, 30].

4.2 Evaluation Metrics

The performance of the proposed framework was quantified using a suite of rigorous evaluation metrics designed to assess both the geometric accuracy and the structural integrity of the generated correspondence maps. The primary metric employed was the standard mean geodesic error. This metric calculates the intrinsic distance along the surface of the target manifold between the point predicted by the algorithm and the actual ground truth point defined by the dataset annotations. The error is normalized by the square root of the total surface area of the target shape to ensure comparability across models of varying absolute dimensions. Lower mean geodesic error values indicate higher geometric precision in the matching process [31]. Beyond pure geometric accuracy, assessing the preservation of structural continuity is paramount for non-isometric matching. A mapping that achieves low average error but introduces localized discontinuities or tears in the surface is fundamentally flawed for downstream applications. To measure this, a conformal distortion metric was utilized, which evaluates the preservation of local angles across the mapped surfaces. Additionally, a semantic consistency score was formulated specifically for this study. This binary metric assesses whether structural components, such as a specific limb or appendage, are mapped entirely to their correct semantic counterpart, severely penalizing algorithms that fracture or split contiguous anatomical parts across disparate regions of the target shape. These multifaceted metrics ensure a comprehensive assessment of the algorithmic performance under challenging non-isometric conditions.

5. Results and Comprehensive Analysis

5.1 Quantitative Performance Evaluation

The quantitative assessment of the topology-calibrated correspondence transfer framework was conducted against a baseline of prevailing state-of-the-art methodologies. These competitors included traditional spectral alignment techniques operating purely on intrinsic wave descriptors, advanced functional map implementations incorporating localized geometric regularization, and recent deep learning architectures utilizing spectral graph convolutions. The evaluation was strictly partitioned into two primary challenge tiers: intra-category matching, where shapes share the same broad morphological class but exhibit varying proportions, and inter-category matching, where shapes belong to different morphological classes, presenting extreme topological and structural divergence.

ISSN: 2519-9439

Table 1: Comparative Evaluation of Geodesic Errors on Non-Isometric Benchmark Datasets

Method	Intra-Category Error	Inter-Category Error	Computation Time
Baseline Spectral	0.085	0.142	12s
Functional Maps	0.062	0.115	18s
Deep Learning Approach	0.045	0.098	45s
Proposed Topology-Calibrated	0.021	0.043	24s

The results documented in Table 1 unequivocally demonstrate the superior performance of the proposed topology-calibrated methodology. In the intra-category tier, the proposed framework achieved a remarkably low mean geodesic error, reducing the alignment discrepancies by more than half compared to the most competitive deep learning baseline. This substantial improvement is directly attributable to the mitigating effect of the topological anchors, which entirely prevent the global spectral drift that commonly afflicts functional techniques when confronted with significant structural elongation or volumetric shifts. The superiority of the proposed framework becomes even more pronounced in the highly challenging inter-category tier [32]. Traditional spectral and functional methods exhibited severe degradation, producing errors indicative of catastrophic mapping failures where entire limbs were frequently mismatched. In contrast, the topology-calibrated approach maintained a high degree of resilience, producing an error margin that remained well within the boundaries of semantic utility. The computation time, while slightly higher than pure functional maps due to the graph extraction overhead, remained highly competitive and significantly faster than the data-intensive deep learning approaches.

5.2 Qualitative Assessment and Visualization

Quantitative metrics alone often fail to capture the holistic quality and the semantic coherence of dense surface mappings. Consequently, rigorous qualitative assessments were conducted by visually projecting standardized texture patterns from the source shapes onto the target shapes via the computed correspondence maps. A seamless, undistorted texture transfer indicates a highly continuous and accurate mapping, whereas fragmented or heavily warped textures highlight regions of algorithmic failure. The visual evidence strongly corroborated the numerical data, showcasing the exceptional stability of the proposed framework. When observing the output of competing methodologies on highly non-isometric pairs, such as mapping a heavily built quadruped to a slender, elongated counterpart, catastrophic failures were frequently visible. Spectral methods commonly suffered from left-right symmetry confusion, erroneously mapping the left foreleg of the source to the right foreleg of the target, owing to the identical intrinsic signatures of symmetrical parts. Functional maps struggled with severe localized distortion, particularly around the junctions where appendages meet the central torso, stretching the texture violently as the basis functions failed to align. The topology-calibrated framework explicitly eradicated these symmetry ambiguities because the structural graph firmly anchors the left and right semantic branches independently during the topological alignment phase. The resulting texture transfers were remarkably smooth, preserving local structural integrity and demonstrating continuous mapping even across regions that experienced extreme non-isometric scaling or severe topological noise.

5.3 Ablation Studies on Topological Features

To rigorously validate the specific architectural decisions embedded within the proposed methodology, comprehensive ablation studies were executed. These studies involved systematically disabling or substituting critical components of the algorithmic pipeline to isolate their individual contributions to the overall performance. The two primary components analyzed were the structural graph alignment phase and the final topology-aware continuous

ISSN: 2519-9439

refinement diffusion process. By analyzing the system's performance in the absence of these modules, the necessity of the dual-layered topological approach becomes demonstrably clear.

Table 2Ablation Analysis on Core Topological Components

Configuration	Mean Geodesic Error	Structural Preservation Score
Full Pipeline	0.021	0.94
Without Graph Alignment	0.056	0.72
Without Refinement	0.034	0.85

The data presented in Table 2 emphasizes the critical nature of the structural graph alignment. When this initial calibration phase is bypassed, and the methodology relies solely on the continuous diffusion process initialized with standard geometric heuristics, the mean geodesic error more than doubles. Furthermore, the structural preservation score plummets, indicating that without the unshakeable global anchors provided by the skeletal graph, the continuous mapping frequently falls into localized minima, resulting in fractured semantic regions [33, 34]. Similarly, the removal of the final localized continuous refinement phase results in a noticeable degradation of performance. While the global semantic structure remains intact, the local transitions across the surface manifold become rigid and lack geometric smoothness. The ablation studies definitively prove that the synergistic combination of abstract topological calibration and precise geometric refinement is absolutely essential for achieving robust correspondence transfer across non-isometric object families.

6. Conclusion and Future Directions

6.1 Summary of Findings

This comprehensive research paper has introduced and rigorously evaluated a novel framework for topology-calibrated correspondence transfer, specifically engineered to overcome the severe limitations of traditional isometric shape matching algorithms when applied to non-isometric object families. The core innovation of the proposed methodology lies in its systematic disentanglement of global structural topology from localized surface geometry. By extracting robust, one-dimensional skeletal graphs directly from the intrinsic properties of the three-dimensional manifolds, the framework establishes a highly stable abstract representation that remains largely impervious to complex structural deformations, topological noise, and extreme anisotropic scaling. Through a meticulously formulated graph alignment strategy, these abstract representations generate a sparse network of high-confidence semantic anchors. These anchors completely bypass the spectral instabilities that plague conventional functional mapping techniques, effectively eradicating common pitfalls such as symmetry confusion and localized spectral drift. The subsequent topology-aware diffusion mechanism successfully bridges the gap between discrete structural alignment and continuous surface geometry, producing dense correspondence maps that are both geometrically precise and semantically coherent. Extensive quantitative and qualitative evaluations conducted on highly challenging non-isometric datasets unequivocally demonstrate the superiority of the proposed framework over existing state-of-the-art methodologies, proving that integrating topological calibration is absolutely essential for the advancement of generalized shape analysis.

6.2 Limitations and Open Challenges

Despite the significant advancements presented in this work, the domain of non-isometric shape matching remains fraught with intricate challenges, and the proposed framework possesses inherent limitations that warrant further investigation. One primary constraint lies in the structural extraction phase; while highly robust to geometric noise, the scalar field computation assumes a reasonable degree of manifold continuity. In scenarios involving wildly fragmented objects, shattered meshes, or extreme point cloud sparsity typically generated by

ISSN: 2519-9439

raw, unrefined sensor data, the gradient descent paths may fail to converge reliably, resulting in disconnected or poorly formed structural graphs. Furthermore, the current alignment algorithm relies on the assumption that at least a foundational semantic core is shared between the source and target objects. When attempting to match completely unrelated semantic categories that lack any structural homology, the topological alignment may produce forced, arbitrary anchors that subsequently corrupt the dense mapping phase.

Future research trajectories must focus on enhancing the resilience of the topological extraction process to severe mesh fragmentation, potentially by integrating probabilistic volumetric techniques or localized deep learning priors to bridge disconnected structural components [35]. Additionally, expanding the framework to seamlessly handle partial matching scenarios, where vast sections of the target object are entirely absent or heavily occluded, represents a critical next step. Exploring multi-modal topological representations that incorporate not just spatial structure, but also localized texture and material properties, could further disambiguate complex non-isometric alignments. By continuing to refine the interplay between abstract topology and continuous geometry, the research community can push the boundaries of shape analysis toward a fully generalized, structurally aware paradigm.

References

Zhang, W., Huang, M., Zhou, Y., Zhang, J., Yu, J., Wang, J., & Xu, L. (2024). Both2hands: Inferring 3d hands from both text prompts and body dynamics. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 2393-2404).

Cong, P., Xu, Y., Ren, Y., Zhang, J., Xu, L., Wang, J., ... & Ma, Y. (2023, June). Weakly supervised 3d multi-person pose estimation for large-scale scenes based on monocular camera and single lidar. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 1, pp. 461-469).

Huang, Y., Thede, L., Mancini, M., Xu, W., & Akata, Z. (2026). Structural Pruning of Large Vision Language Models: A Comprehensive Study on Pruning Dynamics, Recovery, and Data Efficiency. International Journal of Computer Vision, 134(6), 313.

Li, H., Wang, Y., Huang, T., Huang, H., Wang, H., & Chu, X. (2025, October). Ld-rps: Zero-shot unified image restoration via latent diffusion recurrent posterior sampling. In 2025 IEEE/CVF International Conference on Computer Vision (ICCV) (pp. 13684-13694). IEEE.

Lee, Y., Gao, P., Chen, Z., Fan, W., Jing, G., & Hu, Y. (2025, June). Boosting audio-visual segmentation via triple-modalities alignment. In 2025 IEEE International Conference on Multimedia and Expo (ICME) (pp. 1-6). IEEE.

Gao, P., Lee, Y., Zhang, H., Chen, Z., & Liu, X. (2024, December). Dynamic identity-guided attention network for visible-infrared person re-identification. In International Conference on Neural Information Processing (pp. 364-379). Singapore: Springer Nature Singapore.

Qu, W., Shao, Y., Meng, L., Huang, X., & Xiao, L. (2024, June). A conditional denoising diffusion probabilistic model for point cloud upsampling. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 20786-20795). IEEE.

Li, X., Yang, F., Chen, L., & Cai, H. (2016, July). Saliency transfer: An example-based method for salient object detection. In IJCAI (pp. 3411-3417).

Qu, W., Wang, J., Gong, Y., Huang, X., & Xiao, L. (2025, June). An end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models. In 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 27325-27335). IEEE.

Wang, Z., Yuan, R., Geng, Z., Li, H., Qu, X., Li, X., ... & Zhang, K. (2025, October). Singing timbre popularity assessment based on multimodal large foundation model. In Proceedings of the 33rd ACM International Conference on Multimedia (pp. 12227-12236).

Chen, L., Wang, J., Mortlock, T., Khargonekar, P., & Al Faruque, M. A. (2025, June). Hyperdimensional uncertainty quantification for multimodal uncertainty fusion in autonomous vehicles perception. In 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 22306-22316). IEEE.

Xia, Y., Lu, Y., Gao, Y., & Ma, J. (2024, March). Locality preserving refinement for shape matching with functional maps. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 38, No. 6, pp. 6207-6215).

Xia, Y., Ye, T., Huang, J., Mei, X., & Ma, J. (2026, March). Probabilistic deformation consistency for unsupervised shape matching. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 40, No. 13, pp. 10951-10959).

Xia, Y., & Ma, J. (2026). Locality Optimization Refinement with Deformation for Shape Matching via Functional Maps. International Journal of Computer Vision, 134(2), 76.

Yang, Y., Liao, Y. H., Bortins, I., Baldwin, D. P., & Zhang, S. (2024). Unidirectional structured light system calibration with auxiliary camera and projector.Optics and Lasers in Engineering,175, 107984.

Wang, J., Xu, Y., Li, Y., Khargonekar, P., & Faruque, M. A. A. (2026). CRUISE: Vision-Language Model-Guided Uncertainty-Aware Cross-Modal Sensor Fusion for Robust Autonomous Driving. arXiv preprint arXiv:2608.09202.

Cao, Y., Jiao, S., Sun, P., Peng, B., Shi, D., & Shi, Y. (2024). MEFusion: Unsupervised Mutual Enhancement for Multimodal Image Fusion. In ECAI 2024: 27th European Conference on Artificial Intelligence, 19–24 October 2024, Santiago de Compostela, Spain–Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024) (pp. 553-560). 1 Oliver's Yard, 55 City Road, London, EC1Y 1SP: SAGE Publications Pvt. Ltd.

Yang, Y., & Zhang, S. (2025). Pixel-wise calibration for a multi-focus microscopic 3D imaging system.Optics and Lasers in Engineering,194, 109127.

Chen, L., Yang, Y., & Zhang, S. (2025). Auto-focusing and auto-exposure method for real-time structured-light 3D imaging.Optical Engineering,64(4), 044105-044105.

Liu, C., Jing, T., Ma, C., Zhou, X., Lian, Z., Jin, Q., ... & Huang, S. S. (2026). EmoDiffTalk: Emotion-aware Diffusion for Editable 3D Gaussian Talking Head. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 18063-18073).

Yang, Y., Bortins, I., Baldwin, D. P., & Zhang, S. (2024). Calibration of dual resolution dual camera structured light systems.Optics and Lasers in Engineering,182, 108472.

Li, Q., Luo, T., Jiang, M., Liao, J., & Jiang, Z. (2024, October). Deep incomplete multi-view network semi-supervised multi-label learning with unbiased loss. In Proceedings of the 32nd ACM International Conference on Multimedia (pp. 9048-9056).

Li, Q., Luo, T., Jiang, M., Jiang, Z., Hou, C., & Li, F. (2025, April). Semi-supervised multi-view multi-label learning with view-specific transformer and enhanced pseudo-label. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 17, pp. 18430-18438).

Wang, H., Que, H., Xu, Q., Liu, M., Zhou, W., Feng, J., ... & Lin, F. (2026, April). Reverse-engineered reasoning for open-ended generation. In International Conference on Learning Representations (Vol. 2026, pp. 45563-45589).

Wang, H., Feng, W., Yu, J., Liu, C., Nie, P., Lin, F., ... & Wei, C. (2026). Search Beyond What Can Be Taught: Evolving the Knowledge Boundary in Agentic Visual Generation.arXiv preprint arXiv:2607.05382.

Wang, H., Qu, C., Huang, Z., Chu, W., Lin, F., & Chen, W. (2026). Vl-rethinker: Incentivizing self-reflection of vision-language models with reinforcement learning.Advances in Neural Information Processing Systems,38, 30865-30891.

Li, Z., Bai, X., Lu, J., Shen, P., Lam, E. Y., & Peng, Y. (2026). EventTracer: Fast Path Tracing-based Event Stream Rendering. IEEE Transactions on Visualization and Computer Graphics.

ISSN: 2519-9439

Gao, P., Lee, Y., Zhang, X., Chen, Z., & Zhang, H. (2026, March). Remember Me: Bridging the Long-Range Gap in LVLMs with Three-Step Inference-Only Decay Resilience Strategies. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 40, No. 6, pp. 4140-4148).

Li, Z., Jiang, L., Zhao, Y., Chen, Y. C., Wang, X., Chen, W., & Peng, Y. (2026). PatternGSL: A Structured Specification Language for Template-Free and Simulation-Ready 3D Garments. arXiv preprint arXiv:2606.24564.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Proceedings of MICCAI.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., et al. (2021). Learning transferable visual models from natural language supervision. In Proceedings of ICML.

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the Inception architecture for computer vision. In Proceedings of CVPR.

Li, Z., He, Z., Pan, J., Lin, S., & Peng, Y. (2026). E2Pano: Learning Event-to-Panorama Image Reconstruction. arXiv preprint arXiv:2608.00694.

Ling, C., & Wang, Y. (2025). Writing SAS MACROs in R? R functions can help. In PharmaSUG 2025 conference proceedings.

Wei, J. M., Song, Y. L., Zeng, H., Yan, W. W., & Liu, X. B. (2023). Assessing causality between different risk factors and pulmonary embolism: A Mendelian randomization study. Frontiers in cardiovascular medicine, 10, 1057019.