GAUSSIANMORPHING: MESH-GUIDED 3D GAUSSIANS FOR SEMANTIC-AWARE OBJECT MORPHING

Mengtian Li[1,5]**, Yunshu Bai**[1]**, Yimin Chu**[1]**, Yijun Shen**[1]**,**
Zhongmei Li[3]**, Weifeng Ge**[4]**, Zhifeng Xie**[1,5]**, Chaofeng Chen**[2,†]

[1]Shanghai University, [2]Wuhan University, [3]East China University of Science and Technology,
[4]Fudan University, [5]Shanghai Engineering Research Center of Motion Picture Special Effects

ABSTRACT

We introduce **GaussianMorphing**, a novel framework for semantic-aware 3D shape and texture morphing from multi-view images. Previous approaches usually rely on point clouds or require pre-defined homeomorphic mappings for untextured data. Our method overcomes these limitations by leveraging mesh-guided 3D Gaussian Splatting (3DGS) for high-fidelity geometry and appearance modeling. The core of our framework is a unified deformation strategy that anchors 3D Gaussians to reconstructed mesh patches, ensuring geometrically consistent transformations while preserving texture fidelity through topology-aware constraints. In parallel, our framework establishes unsupervised semantic correspondence by using the mesh topology as a geometric prior and maintains structural integrity via physically plausible point trajectories. This integrated approach preserves both local detail and global semantic coherence throughout the morphing process without requiring labeled data. On our proposed TexMorph benchmark, GaussianMorphing substantially outperforms prior 2D/3D methods, reducing color consistency error (ΔE) by 22.2% and EI by 26.2%. Project page: https://baiyunshu.github.io/GAUSSIANMORPHING.github.io/

1 INTRODUCTION

Morphing (Gregory et al., 1998; Zhang et al., 2024a) has long been a foundational technique in shape transformation, enabling the generation of continuous interpolation sequences between source and target shapes. Serving as a bridge between computer vision and computer graphics, morphing has emerged as an indispensable tool for applications spanning computer animation, geometric modeling, and shape analysis. Its prominence in visual effects for film and media production further underscores its practical significance.

Existing morphing techniques can be broadly categorized into two paradigms: image-based methods (Aloraibi, 2023; Zhang et al., 2024a) and 3D geometric methods (Eisenberger et al., 2021; Yang et al., 2025; Cao et al., 2024). As summarized in Figure 1, these approaches exhibit fundamental trade-offs. Image-based pipelines, such as DiffMorpher (Zhang et al., 2024b) and FreeMorph (Cao et al., 2025), produce high-fidelity 2D outputs but lack 3D geometric reasoning and multi-view consistency. Extensions like MorphFlow (Tsai et al., 2022) leverage Neural Radiance Fields (NeRF) to address view consistency but are limited by the absence of explicit 3D geometric constraints, resulting in incomplete volumetric reconstructions (denoted as 2.5D* in Figure 1). In contrast, 3D-centric methods such as Neuromorph (Eisenberger et al., 2021) enable mesh-based deformation but require high-quality mesh inputs, neglect texture-aware processing, and struggle with topological complexity. These limitations highlight a critical gap: the lack of a unified framework that balances geometric robustness, textural coherence, and input accessibility without reliance on high-fidelity 3D data, which remains a key challenge for advancing morphing techniques toward practical and general-purpose applications.

[†]Corresponding author.

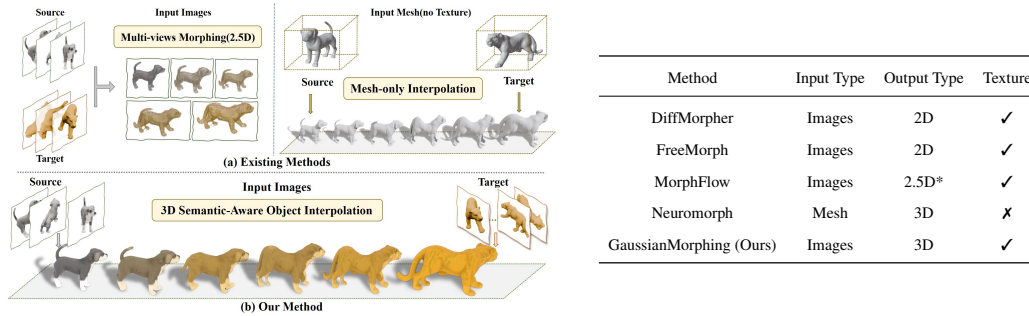| Method	Input Type	Output Type	Texture
DiffMorpher	Images	2D	✓
FreeMorph	Images	2D	✓
MorphFlow	Images	2.5D*	✓
Neuromorph	Mesh	3D	✗
GaussianMorphing (Ours)	Images	3D	✓

Figure 1: Our **GaussianMorphing** (left) takes input images of the source and target, reconstructs them into 3D Gaussian representations with surface meshes, and uses a **mesh-guided strategy** to generate intermediate shapes at timestamps $t \in [0, 1]$. Unlike prior approaches, our method achieves Semantic-Aware Object Morphing with textured colors without relying on 3D input data. The comparison table (right) shows that our method uniquely generates fully textured 3D outputs directly from images, offering complete geometric and textural fidelity.

To address this gap, this work introduces the first framework for joint 3D geometry and texture morphing using 3D Gaussians, where shape and appearance are intrinsically unified (Figure 1). The key challenge lies in achieving coherent deformation with Gaussian representations due to their unstructured nature and the complexity of maintaining geometry-texture alignment. Our solution integrates the rendering efficiency of 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) with the structural benefits of mesh-guided deformation. The approach establishes explicit bindings between 3DGS primitives and mesh elements, enabling smooth interpolation while preserving geometric and textural fidelity. Through mesh feature extraction and topological constraints, the method ensures stable morphing sequences that resist the geometric fragmentation typical of discrete point representations. A dual-domain optimization strategy employs geodesic-based geometric distortion loss and texture-aware color smoothness loss to govern deformation, ensuring temporal coherence from accessible 2D inputs without requiring specialized 3D assets.

The proposed framework bridges discrete 3DGS points with semantic-aware mesh structures, achieving significant improvements over state-of-the-art methods in geometric accuracy and texture preservation. Experiments demonstrate robust performance across diverse scenarios, including complex topologies and texture-rich objects, while reducing dependency on high-quality 3D data.

The main contributions are:

(1) A mesh-guided framework that integrates 3D Gaussian Splatting with semantic-aware morphing, enabling high-fidelity 3D interpolation from minimal inputs;

(2) Deformation mechanisms that are aware of both topology and semantics, preventing geometric fragmentation and ensuring stable, coherent morphing in Gaussian-based representations;

(3) A dual-domain optimization strategy combining geodesic-aware geometric constraints and texture-aware color interpolation that achieves seamless visual results.

2 RELATED WORK

2.1 IMAGE MORPHING

Image morphing is a long-standing problem in computer vision and graphics, aiming to generate smooth and perceptually natural transitions between images (Aloraibi, 2023; Zope & Zope, 2017; Wolberg, 1998). Traditional methods (Beier & Neely, 2023; Bhatt, 2011; Liao et al., 2014) rely on correspondence-driven warping and blending, which preserve visual consistency but struggle with content creation, often leading to artifacts. More recently, optimal transport has been applied to morphing simple 2D geometries (Benamou et al., 2015; Bonneel et al., 2011; Solomon et al., 2015), providing mathematically elegant transformations but lacking the texture richness of natural images. Diffusion-based approaches such as DiffMorpher (Zhang et al., 2024b), AID (He et al., 2024), and FreeMorph (Cao et al., 2025) leverage pre-trained generative models to enable flexible

morphing across diverse categories. In this work, we instead start from multi-view inputs, alleviating the need for large-scale pre-training and producing intermediate mesh-based representations that support shape-aware and texture-consistent 3D morphing.

2.2 SHAPE MATCHING

The problem of 3D shape correspondence aims to establish point-wise mappings between shapes and has been widely studied. Traditional methods rely on geometric constraints (Holzschuh et al., 2020; Roetzer et al., 2022) or non-rigid registration (Bernard et al., 2020; Eisenberger et al., 2019; Ezuz et al., 2019), but often require costly optimization and manual alignment, limiting scalability. Recent learning-based approaches have advanced the field by training networks to match vertices to a template (Monti et al., 2017; Boscaini et al., 2016; Masci et al., 2015), or by leveraging functional maps with learnable features (Litany et al., 2017; Ovsjanikov et al., 2012). Others integrate spectral and spatial cues (Cao et al., 2024; Attaiki & Ovsjanikov, 2023), use diffusion models for functional map prediction (Zhuravlev et al., 2025), or apply 2D correspondence priors to improve semantic consistency in 3D registration (Liu et al., 2025). Our method, with the assistance of neural networks, eliminates the need for costly 3D inputs and data annotations. By employing object reconstruction techniques, it derives geometric point-wise correspondences from images.

2.3 SHAPE INTERPOLATION

Shape interpolation addresses the fundamental challenge of smoothly transforming one shape into another by generating intermediate shapes at specified composition percentages. Traditional geometric methods (Brandt et al., 2016; Heeren et al., 2012; Wirth et al., 2011) formulate this as finding geodesic paths on high-dimensional manifolds, employing deformation metrics like As-Rigid-As-Possible (ARAP) (Sorkine & Alexa, 2007) and PriMo (Botsch et al., 2006) to minimize local distortions. Data-driven approaches alternatively navigate through collections of related shapes (Aydınlılar & Sahillioğlu, 2021; Gao et al., 2017), while physics-based methods model interpolation as constrained gradient flows (Eisenberger & Cremers, 2020; Eisenberger et al., 2019). MorphFlow exemplifies this approach by combining Wasserstein flow with rigidity constraints for multiview morphing (Tsai et al., 2022). Recent neural approaches have advanced unsupervised shape interpolation. NeuroMorph (Eisenberger et al., 2021) and Spectral Meets Spatial (Cao et al., 2024) demonstrate effective frameworks for shape matching and interpolation, with the latter incorporating spectral regularization for handling large non-isometric deformations. Other methods utilize 2D correspondence guidance (Liu et al., 2025) or diffusion priors for textured morphing (Yang et al., 2025). Our method combines geodesic distance measurements with ARAP constraints, utilizing a neural network-based interpolator to achieve smooth deformation from source to target shapes.

3 MESH-GUIDED GAUSSIAN MORPHING

Given source and target objects represented by multi-view images, we propose a semantic-aware 3D morphing framework that addresses a fundamental challenge: achieving geometrically consistent transformations while preserving photorealistic surface details. The core problem is that modern explicit representations present a trade-off: 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) lacks the topological connectivity needed for structured morphing, while traditional meshes struggle to model complex appearance.

As shown in Figure 2, Our framework, Mesh-Guided Gaussian Morphing, resolves this tension by introducing a novel hybrid paradigm. Our key insight is to impose an explicit triangular mesh as a *topological scaffold* to guide the transformation of unstructured Gaussians. By anchoring Gaussians to this mesh, we can leverage powerful mesh-based correspondence techniques to establish semantic connections. This allows us to compute a geometrically consistent morphing flow in the structured mesh domain while using the rich Gaussian representation for photorealistic rendering at any point in the transformation.
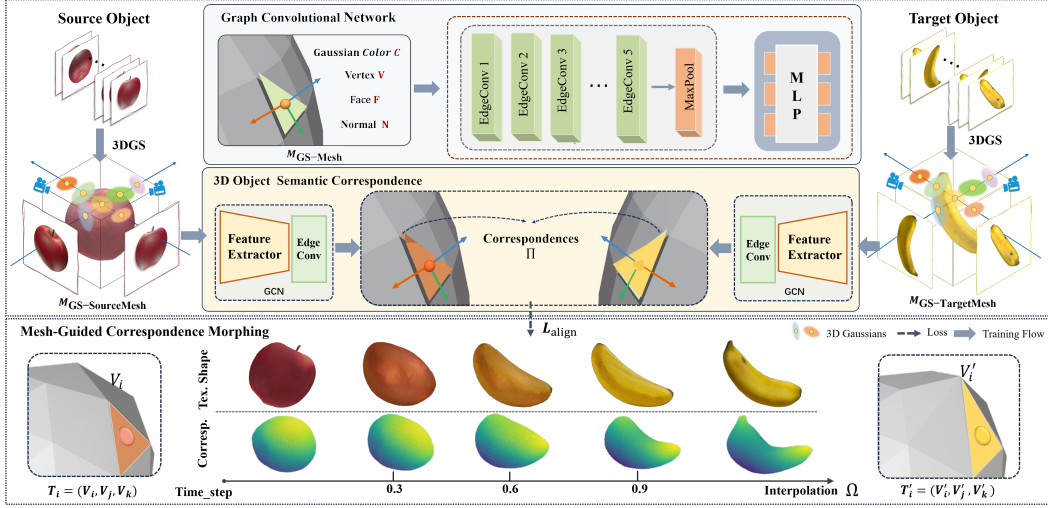
Figure 2: **Method Overview.** Our **GaussianMorphing** framework takes source $\mathcal{X}$ and target $\mathcal{Y}$ images as input. Surface meshes are extracted from 3D Gaussian Splatting (Sec. 3.1) and used with Gaussian points for geometry–texture alignment. Geometric features provide the correspondence matrix Π_{XY} (Sec. 3.2), and intermediate shapes are interpolated over time. Training relies on a joint loss (Sec. 3.3), yielding high-quality textured 3D morphing. (Up: Blender results; Down: correspondence visualization with Matplotlib.)

3.1 HYBRID MESH-GAUSSIAN REPRESENTATION FOR SEMANTIC MORPHING

The Connectivity Challenge in Gaussian-Based Morphing. 3DGS represents a scene as a set of anisotropic 3D Gaussians, which can be optimized to reproduce a set of input images, enabling photorealistic novel-view synthesis. Each Gaussian g is defined by its position $\mu_g \in \mathbb{R}^3$, covariance Σ_g, opacity α_g, and spherical harmonics (SH) coefficients $\mathbf{sh}_g$. While excellent for rendering, the discrete, unstructured nature of these Gaussians prevents the establishment of meaningful semantic correspondences between objects. A direct Gaussian-to-Gaussian matching would likely produce geometrically implausible results that tear or distort the structure of the object.

Mesh-Anchored Gaussian Binding. To overcome this limitation, we impose a topological structure by anchoring Gaussians to an explicit mesh. First, we extract a high-quality initial mesh from the optimized Gaussians. We follow recent methods like SuGaR (Guédon & Lepetit, 2024b) and FrostingGaussian (Guédon & Lepetit, 2024a), which use Poisson reconstruction (Kazhdan et al., 2006) alongside regularization terms to ensure the mesh surface accurately reflects the geometry captured by the Gaussians.

With this mesh scaffold, we establish an explicit binding between the Gaussians and the mesh faces. Each Gaussian is anchored to a specific triangular face $f = (V_1, V_2, V_3)$, with its position μ_g defined by barycentric coordinates (w_1, w_2, w_3) and a normal offset d:

$$\mu_g = w_1 V_1 + w_2 V_2 + w_3 V_3 + d \cdot \mathbf{n}_f, \tag{1}$$

where $\mathbf{n}_f$ is the face normal. This binding ensures that as the mesh vertices V_i deform over the course of the morph, the anchored Gaussians move cohesively with the surface, preserving the fine-grained geometric and appearance details they represent.

3.2 SEMANTIC CORRESPONDENCE THROUGH TOPOLOGICAL UNDERSTANDING

Semantic-Aware Mesh Correspondence. With the mesh structure established, we can tackle the core challenge of identifying which part of the source object should transform into which part of the target. We formulate this as a correspondence problem between the source mesh (V^S, F^S) and target mesh (V^T, F^T). The correspondence is encoded as a probabilistic matrix $\Pi \in \mathbb{R}^{n \times m}$:

$$\Pi_{ij} = P(V_j^T \mid V_i^S) = \frac{\exp(\sigma c_{ij})}{\sum_{k=1}^{m} \exp(\sigma c_{ik})}, \tag{2}$$

where c_{ij} is the cosine similarity between learned feature vectors for source vertex V_i^S and target vertex V_j^T. To learn semantically rich features, we use a 5-layer Graph Convolutional Network (GCN) that processes mesh connectivity, allowing it to capture local geometric context without relying on hand-engineered descriptors.

Neural Morphing Flow. Rather than simple linear interpolation, we learn a continuous, non-linear deformation field. We employ a neural network, the Correspondence Morphing Flow (Ψ), to predict the morphing trajectory. At any time $t \in [0, 1]$, the morphed source vertices $V^S(t)$ are given by:

$$V^S(t) = V^S + \Psi(V^S, \Pi V^T - V^S, t). \tag{3}$$

Here, the term $\Pi V^T - V^S$ represents the *semantically-aligned displacement field* that maps each source vertex to its corresponding target location. The network Ψ learns to smoothly interpolate this displacement over time.

Consistent Gaussian Updates. As the mesh vertices $V^S(t)$ deform, the positions of the bound Gaussians $\mu_g(t)$ are updated consistently via the barycentric relationship established in Eq. 1:

$$\mu_g(t) = \sum_{i=1}^{3} w_i V_{f_i}(t), \tag{4}$$

where $V_{f_i}(t)$ are the deformed positions of the vertices of the triangle f to which Gaussian g is bound. This maintains the tight coupling between the mesh and the Gaussians throughout the entire morphing sequence.

3.3 Multi-Objective Optimization for Plausible Morphing

We optimize the correspondence matrix Π and the morphing flow network Ψ using a comprehensive loss function that balances geometric structure, appearance consistency, and semantic alignment.

Geometric Consistency. To prevent unnatural stretching and distortion, we enforce that the intrinsic geometric structure of the surfaces is preserved. We measure this using geodesic distances on the mesh. To compute the geodesic distance $D_g(i, j)$ between any two vertices, we run Dijkstra's algorithm on a hybrid graph formed by the union of the mesh adjacency graph G_{adj} (preserving topology) and a KNN graph G_{knn} (adding shortcuts to better approximate Euclidean distances). Further details are provided in Appendix A.2. The geodesic distortion loss is then:

$$\mathcal{L}_{geo} = \left\| \Pi D_g^T \Pi^\top - D_g^S \right\|_F^2, \tag{5}$$

where D_g^S and D_g^T are the geodesic distance matrices for the source and target meshes, and $\|\cdot\|_F$ is the Frobenius norm. This loss encourages the correspondence Π to map regions of the target mesh back to the source mesh in a way that respects their intrinsic geometry.

To further encourage local rigidity, we add an As-Rigid-As-Possible (ARAP) energy term (Sorkine & Alexa, 2007), which penalizes non-rigid deformations. We evaluate this over sampled timesteps during the morph:

$$\mathcal{L}_{arap} = \mathbb{E}_{t \sim U[0,1]} \left[E_{arap}(\mathbf{X}(t), \mathbf{X}(t + \delta t)) \right], \tag{6}$$

where $\mathbf{X}(t)$ is the mesh state at time t and δt is a small perturbation.

Appearance Consistency. To ensure smooth visual transitions, we introduce a geodesic-aware smoothness loss on the vertex colors. We first initialize the color of each vertex by averaging the RGB colors of its bound Gaussians (with SH coefficients evaluated from a canonical viewing direction). The loss then penalizes color differences between adjacent vertices, weighted inversely by their geodesic distance:

$$\mathcal{L}_{smooth} = \sum_{(i,j) \in E_{adj}} \frac{1}{D_g(i, j) + \epsilon} \cdot \left\| C_{morph}^i(t) - C_{morph}^j(t) \right\|_2^2, \tag{7}$$

where E_{adj} is the set of edges in the mesh adjacency graph. This encourages smooth color fields while allowing for sharp transitions across distant parts of the object.
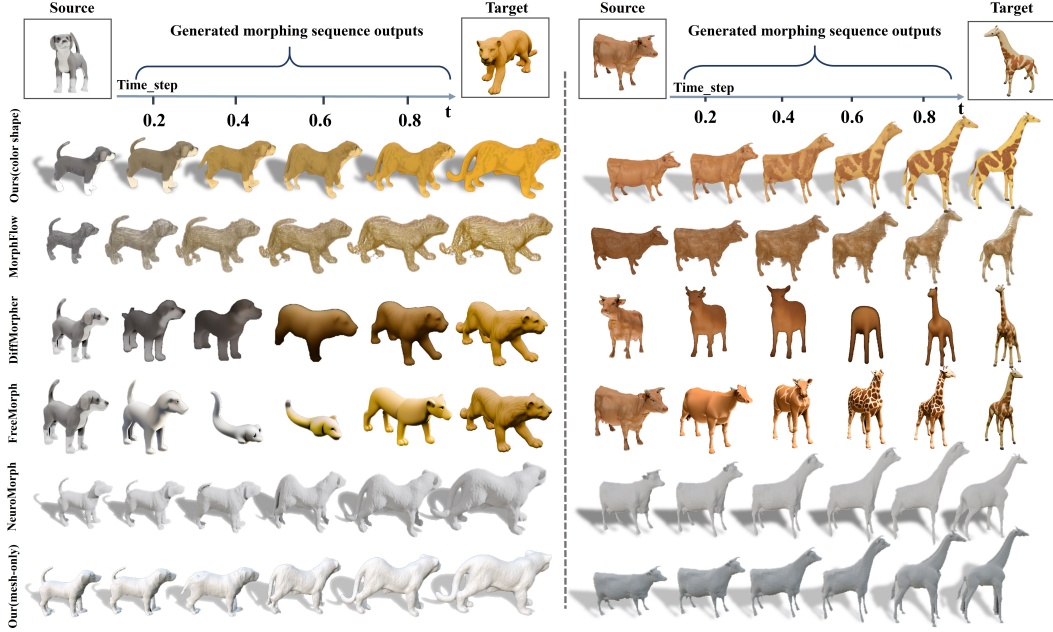
Figure 3: Qualitative comparison of morphing methods on the benchmark dataset. Baselines include DiffMorpher (Zhang et al., 2024b) and FreeMorph (Cao et al., 2025) for image morphing, Neuro-Morph (Eisenberger et al., 2021) for texture-free 3D shape morphing, and MorphFlow (Tsai et al., 2022) for textured multi-view morphing without true geometry. Our method generates textured 3D morphing with geometric details directly from image inputs.

Semantic Alignment Constraint. To ensure the morphing sequence reaches its destination, we add a terminal constraint that drives the deformed source mesh to the target configuration at the final timestep:

$$\mathcal{L}_{\text{align}} = \left\| V^S(t=1) - \Pi V^T \right\|_F^2. \tag{8}$$

This loss acts as a boundary condition, ensuring that the morph respects the learned semantic correspondences.

Unified Loss Function. Our final objective function is a weighted sum of these components:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{geo}} \mathcal{L}_{\text{geo}} + \lambda_{\text{arap}} \mathcal{L}_{\text{arap}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}} + \lambda_{\text{align}} \mathcal{L}_{\text{align}}, \tag{9}$$

where the λ hyperparameters balance the competing objectives of geometric fidelity, structural rigidity, appearance consistency, and semantic alignment.

4   EXPERIMENTS

We conduct comprehensive experiments to validate the ability of our method to produce high-quality, semantically consistent 3D morphs. We introduce a new benchmark, **TexMorph**, designed specifically for this task. Our evaluation protocol includes quantitative comparisons against state-of-the-art 2D and 3D methods using novel metrics, qualitative analysis of the generated morphing sequences, and an ablation study to analyze the contributions of our proposed key components.

4.1   EXPERIMENTAL SETUP

TexMorph Benchmark. To rigorously evaluate 3D morphing from multi-view images, we created a new benchmark named **TexMorph** (**Tex**ture-rich, **M**orphing-focused). The benchmark is comprised of challenging source-target pairs designed to test geometric and appearance transformations. It includes: (1) high-fidelity synthetic models with complex textures rendered from multiple viewpoints; (2) real-world objects captured via 3D scanning; and (3) objects captured in-the-wild using

Figure 4: Qualitative morphing results with non-isometric deformations, demonstrating robust interpolation under challenging geometric conditions (Up: synthetic datas; Middle: real-world scanned objects from GSO (Downs et al., 2022); Bottom: real-world photos).

standard mobile phone cameras. The dataset features over ten object categories, including animals, fruits, and vehicles, providing diverse topological and textural challenges. Further details are provided in Appendix A.1.

Evaluation Metrics. Standard metrics for novel-view synthesis are inadequate for evaluating morphing sequences. We thus propose three metrics to assess the spatio-temporal quality of the transformation from source ($t = 0$) to target ($t = 1$):

- **Structural Stability (MSE-SSIM):** Measures geometric consistency by computing the Mean Squared Error of the temporal SSIM scores against an ideal linear trajectory.

$$\mathcal{E} = \frac{1}{N} \sum_{t \in T} \Big(\text{SSIM}_{\text{ideal}}(A, G_t) - \text{SSIM}_{\text{actual}}(A, G_t) \Big)^2 + \Big(\text{SSIM}_{\text{ideal}}(G_t, B) - \text{SSIM}_{\text{actual}}(G_t, B) \Big)^2.$$
(10)

  A lower value indicates a more stable transformation with fewer structural artifacts.

- **Color Consistency (ΔE):** Assesses appearance smoothness by averaging the perceptual color difference (ΔE_{ab}^*) between corresponding surface points throughout the morph.

$$\Delta E_{ab}^* = \sqrt{(L_1^* - L_2^*)^2 + (a_1^* - a_2^*)^2 + (b_1^* - b_2^*)^2}.$$
(11)

  A lower ΔE signifies a smoother transition without color bleeding.

- **Edge Integrity (EI):** Evaluates silhouette continuity by measuring the temporal stability of the rendered edge map of object.

$$EI = N_{Edges}(Canny(I, T_{low}, T_{high})) - 1.$$
(12)

  A lower score indicates less fragmented edges, suggesting more stable structural transition in the morphing sequence.

Detailed formulations are available in Appendix A.3.

Implementation Details. All experiments were conducted on a single NVIDIA RTX A6000 GPU. For a typical object pair with a mesh of approximately 12,000 faces, the initial hybrid mesh-Gaussian representation is generated in about 1 hour. The optimization of our morphing framework takes between 500 and 1000 iterations, depending on mesh complexity. Once trained, generating a full, high-resolution morphing sequence takes approximately 2 minutes.

4.2 EVALUATION

We perform a comprehensive evaluation of GaussianMorphing against several state-of-the-art 2D and 3D morphing methods. For 2D baselines, we compare against DiffMorpher (Zhang et al., 2024b), a diffusion-based method, and FreeMorph (Cao et al., 2025), a tuning-free approach. For 3D baselines, we include MorphFlow (Tsai et al., 2022), which leverages optimal transport for multi-view transitions, and NeuroMorph (Eisenberger et al., 2021), which computes topology-aligned shape correspondences.

Method	MSE(SSIM)↓	ΔE↓	EI↓
DiffMorpher	0.19	105	97
MorphFlow	0.17	8.23	33.6
Neuromorph	0.13	/	13.0
FreeMorph	0.20	13.0	21.6
Our	**0.11**	**6.40**	**9.0**

Table 1: Quantitative comparison of morphing methods evaluates structural similarity using the MSE of SSIM, color consistency with ΔE, and edge continuity through EI.
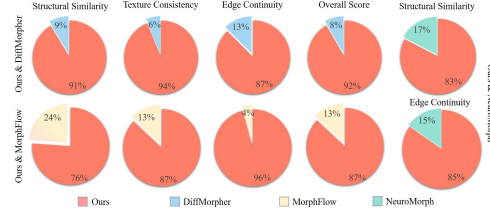

Figure 5: User study results: Comparing our method with the baseline methods in terms of color consistency, structural similarity, and edge continuity. A higher percentage of participants preferred our results across all metrics.

Textured Morphing Analysis. Our method excels at producing smooth, high-fidelity texture transitions, as qualitatively demonstrated in Figure 3. The linear color interpolation of MorphFlow is inadequate for high-dimensional color spaces, leading to oversmoothed transitions and loss of detail. For example, during the "dog→lion" transformation, it reduces the morph to a simple color shift, failing to preserve the intricate fur patterns of the lion or the distinct white patches of the dog. The 2D methods perform poorly on challenging cross-category pairs; DiffMorpher fails in both geometric and color alignment, while the 2D SOTA, FreeMorph, introduces severe structural artifacts (*e.g.*, lizard-like textures) and color oversaturation. In contrast, our approach achieves superior color fidelity, corroborated by lower ΔE values (Table 1), and maintains fine-grained texture details throughout the transformation. Furthermore, as shown in Figure 4, our method produces smooth and plausible morphing results even in the presence of significant non-isometric deformations. The interpolated sequences remain visually coherent, demonstrating the robustness of our approach under challenging geometric conditions. By covering synthetic models, real-world scanned objects from GSO (Downs et al., 2022), and photographs of everyday items, the results further highlight the generalization ability of Gaussian morphing across diverse data sources.

Geometric and Structural Analysis. As shown in Table 1, our method achieves state-of-the-art structural consistency and edge continuity, primarily due to $\mathcal{L}_{geo}$, which preserves local geometric details. For a fair comparison with NeuroMorph, we use the same input meshes for both methods. NeuroMorph relies on mesh connectivity for geodesic computation making it brittle when handling fragmented or coarse geometries. Our hybrid graph representation bypasses this dependency, yielding a more robust and efficient solution. Furthermore, our semantic-aware mechanism produces more plausible deformations, correctly preserving features like the tail in "dog→lion" morphs and neck details in giraffe morphs, where NeuroMorph falters. MorphFlow suffers from a lack of constraints on mesh topology or semantic information, an absence that leads to noticeable edge fragmentation and silhouette tearing. By contrast, our topology-aware framework effectively avoids these issues by leveraging the mesh structure to ensure enhanced edge continuity.

User Study. To validate the perceptual quality of our results, we conducted a user study with 54 participants, who compared outputs of our method against those from DiffMorpher, MorphFlow, NeuroMorph and the Ablation study. The evaluation focused on four criteria: structural similarity, texture consistency, edge continuity, and overall preference.The criteria shown below:

- **Structural Similarity**: Preservation of structure in intermediate frames.
- **Texture Consistency**: Smooth and natural color transitions without abrupt jumps.
- **Edge Continuity**: Smooth and continuous edges without breaks or distortions.
- **Overall Score**: Comprehensive evaluation based on structural similarity, texture consistency, and edge continuity.

Full details are provided in Appendix A.4. The results show an overwhelming preference for our method across all metrics. Over 80% of users rated our morphs as superior overall, with particularly strong and consistent agreement on aspects such as texture consistency and edge continuity. This perceptual validation confirms that our method generates more visually coherent and high-quality morphs, aligning with our quantitative experiments.

4.3 ABLATION STUDY

We conducted an ablation study to isolate the contributions of our core components: the mesh-guided strategy and the geometric distortion loss.
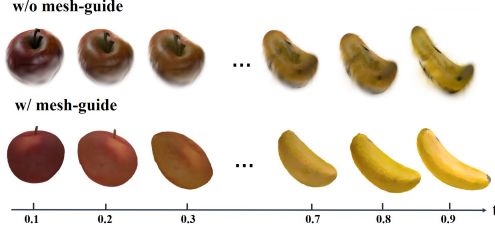

Figure 6: Ablation Study for mesh-guided strategy. Top: Morphing without the **mesh-guided strategy**. Bottom: Morphing with the strategy, demonstrating its role in achieving edge-continuous and smooth transitions.
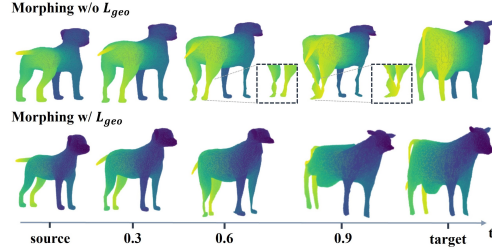

Figure 7: Ablation Study for **geometric distortion loss**. Comparison of morphing results without (up) and with (below) the geometric distortion loss.

Importance of Mesh Guidance.

To evaluate the critical role of mesh guidance in maintaining morphing coherence, we first conducted a comparative evaluation between two distinct approaches: (1) a variant of our method that removes mesh guidance, relying solely on point-based morphing, and (2) our full model, which incorporates mesh guidance by leveraging the complete mesh topology (including vertices, edges, faces, and normals) to establish a shared correspondence Π. As summarized in Table 2 and illustrated in Figure 6, the point-based variant fails to maintain structural coherence, resulting in significant tearing and discontinuities along object surfaces, particularly noticeable at edges. Quantitatively, this structural degradation is reflected in a higher Edge Continuity Index (EI) score of 34.3, indicating poorer performance. The absence of topological guidance also compromises texture quality, leading to blurry artifacts. In contrast, the mesh-guided approach effectively preserves structural integrity and ensures smoother transitions by enforcing spatial and textural consistency through the explicit use of mesh structure.

Role of Geometric Distortion Loss.

Next, we ablate the geometric distortion loss. Without this constraint, the morphing process introduces severe and unnatural deformations, such as the distorted leg geometry shown in Figure 7. These artifacts not only degrade visual quality but also disrupt the structural plausibility of the interpolated shapes, making the transitions appear unrealistic. By explicitly penalizing local shape changes, this loss serves as a key regularizer that preserves structural integrity, en-

Table 2: Mesh-Guided Strategy Ablation: Quantifying edge continuity (EI), user-rated transition quality, and texture preservation (MSE(SSIM)) to validate the importance of mesh guidance for smooth shape and texture morphing.

	Edge Continuity		Texture Quality
	EI↓	User↑	MSE(SSIM)↓
w/o Mesh-Guided	34.3	0.02	0.34
w/o $\mathcal{L}_{smooth}$	–	–	0.22
Ours	**9.0**	**0.98**	**0.11**

forces geometric continuity, and produces smoother, more plausible transformations. User feedback corroborates this finding, confirming a marked reduction in visual distortion when the loss is applied.

In summary, these studies demonstrate that the synergy between mesh guidance and geometric distortion loss is essential for achieving high-fidelity geometric and textural transformations, significantly improving geometric continuity and leading to more natural morphing results.

5 CONCLUSION

We introduced **GaussianMorphing**, a novel semantic-aware framework that unifies 3D shape and texture morphing from multi-view images. Our key innovation is a *mesh-guided Gaussian morphing*

strategy that anchors 3D Gaussians to semantic mesh patches. This approach bypasses the need for pre-aligned 3D assets and ensures that geometry and appearance are interpolated in a structurally consistent and texturally coherent manner. Through unsupervised learning guided by mesh topology, our method achieves state-of-the-art performance, outperforming existing 2D and 3D techniques in structural similarity, color consistency, and edge continuity. By generating efficient and visually faithful transformations, GaussianMorphing sets a new standard for 3D morphing and opens up new possibilities for applications in visual effects and digital content creation.

REFERENCES

Alyaa Qusay Aloraibi. Image morphing techniques: A review. *Technium*, 9, 2023.

Souhaib Attaiki and Maks Ovsjanikov. Shape non-rigid kinematics (snk): A zero-shot method for non-rigid shape matching via unsupervised functional map regularized reconstruction. *Advances in Neural Information Processing Systems*, 36:70012–70032, 2023.

Melike Aydınlılar and Yusuf Sahillioğlu. Part-based data-driven 3d shape interpolation. *Computer-Aided Design*, 136:103027, 2021.

Thaddeus Beier and Shawn Neely. Feature-based image metamorphosis. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 529–536. 2023.

Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.

Florian Bernard, Zeeshan Khan Suri, and Christian Theobalt. Mina: Convex mixed-integer programming for non-rigid shape alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13826–13835, 2020.

Bhumika G Bhatt. Comparative study of triangulation based and feature based image morphing. *Signal & Image Processing*, 2(4):235, 2011.

Nicolas Bonneel, Michiel Van De Panne, Sylvain Paris, and Wolfgang Heidrich. Displacement interpolation using lagrangian mass transport. In *Proceedings of the 2011 SIGGRAPH Asia conference*, pp. 1–12, 2011.

Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. *Advances in neural information processing systems*, 29, 2016.

Mario Botsch, Mark Pauly, Markus H Gross, and Leif Kobbelt. Primo: coupled prisms for intuitive surface modeling. In *Symposium on Geometry Processing*, pp. 11–20, 2006.

Christopher Brandt, Christoph von Tycowicz, and Klaus Hildebrandt. Geometric flows of curves in shape space for processing motion of deformable objects. In *Computer Graphics Forum*, volume 35, pp. 295–305. Wiley Online Library, 2016.

Dongliang Cao, Marvin Eisenberger, Nafie El Amrani, Daniel Cremers, and Florian Bernard. Spectral meets spatial: Harmonising 3d shape matching and interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3658–3668, 2024.

Yukang Cao, Chenyang Si, Jinghao Wang, and Ziwei Liu. Freemorph: Tuning-free generalized image morphing with diffusion model. 2025.

Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2553–2560. IEEE, 2022.

Marvin Eisenberger and Daniel Cremers. Hamiltonian dynamics for real-world shape interpolation. In *European conference on computer vision*, pp. 179–196. Springer, 2020.

Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. Divergence-free shape correspondence by deformation. In *Computer Graphics Forum*, volume 38, pp. 1–12. Wiley Online Library, 2019.

Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7473–7483, 2021.

Danielle Ezuz, Behrend Heeren, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. Elastic correspondence between triangle meshes. In *Computer Graphics Forum*, volume 38, pp. 121–134. Wiley Online Library, 2019.

Lin Gao, Shu-Yu Chen, Yu-Kun Lai, and Shihong Xia. Data-driven shape interpolation and morphing editing. In *Computer Graphics Forum*, volume 36, pp. 19–31. Wiley Online Library, 2017.

Arthur Gregory, A State, Ming C Lin, Dinesh Manocha, and Mark A Livingston. Feature-based surface decomposition for correspondence and morphing between polyhedra. In *Proceedings Computer Animation'98 (Cat. No. 98EX169)*, pp. 64–71. IEEE, 1998.

Antoine Guédon and Vincent Lepetit. Gaussian frosting: Editable complex radiance fields with real-time rendering. In *European Conference on Computer Vision*, pp. 413–430. Springer, 2024a.

Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5354–5363, 2024b.

Qiyuan He, Jinghao Wang, Ziwei Liu, and Angela Yao. Aid: Attention interpolation of text-to-image diffusion. *arXiv preprint arXiv:2403.17924*, 2024.

Behrend Heeren, Martin Rumpf, Max Wardetzky, and Benedikt Wirth. Time-discrete geodesics in the space of shells. In *Computer Graphics Forum*, volume 31, pp. 1755–1764. Wiley Online Library, 2012.

Benjamin Holzschuh, Zorah Lähner, and Daniel Cremers. Simulated annealing for 3d shape correspondence. In *2020 International Conference on 3D Vision (3DV)*, pp. 252–260. IEEE, 2020.

Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

Jing Liao, Rodolfo S Lima, Diego Nehab, Hugues Hoppe, Pedro V Sander, and Jinhui Yu. Automating image morphing using structural similarity on a halfway domain. *ACM Transactions on Graphics (TOG)*, 33(5):1–12, 2014.

Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision*, pp. 5659–5667, 2017.

Haolin Liu, Xiaohang Zhan, Zizheng Yan, Zhongjin Luo, Yuxin Wen, and Xiaoguang Han. Stablescore: A stable registration-based framework for 3d shape correspondence. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 917–928, 2025.

Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pp. 37–45, 2015.

Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5115–5124, 2017.

Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)*, 31(4):1–11, 2012.

Paul Roetzer, Paul Swoboda, Daniel Cremers, and Florian Bernard. A scalable combinatorial solver for elastic geometrically consistent 3d shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 428–438, 2022.

Justin Solomon, Fernando De Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (ToG)*, 34(4):1–11, 2015.

Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pp. 109–116. Citeseer, 2007.

Chih-Jung Tsai, Cheng Sun, and Hwann-Tzong Chen. Multiview regenerative morphing with dual flows. In *European Conference on Computer Vision*, pp. 492–509. Springer, 2022.

Benedikt Wirth, Leah Bar, Martin Rumpf, and Guillermo Sapiro. A continuum mechanical approach to geodesics in shape space. *International Journal of Computer Vision*, 93(3):293–318, 2011.

George Wolberg. Image morphing: a survey. *The visual computer*, 14(8-9):360–372, 1998.

Songlin Yang, Yushi Lan, Honghua Chen, and Xingang Pan. Textured 3d regenerative morphing with 3d diffusion prior. *arXiv preprint arXiv:2502.14316*, 2025.

Kaiwen Zhang, Yifan Zhou, Xudong Xu, Bo Dai, and Xingang. Pan. Diffmorpher: Unleashing the capability of diffusion models for image morphing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7912–7921, 2024a.

Kaiwen Zhang, Yifan Zhou, Xudong Xu, Bo Dai, and Xingang Pan. Diffmorpher: Unleashing the capability of diffusion models for image morphing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7912–7921, 2024b.

Aleksei Zhuravlev, Zorah Lähner, and Vladislav Golyanik. Denoising functional maps: Diffusion models for shape correspondence. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 26899–26909, 2025.

Bhushan Zope and Soniya B Zope. A survey of morphing techniques. *International Journal of Advanced Engineering, Management and Science*, 3(2):239773, 2017.

A   APPENDIX

Statements The language of this manuscript was refined with the assistance of a large language model (LLM). The authors remain fully responsible for the originality, accuracy, and integrity of all academic content, analyses, and ideas presented in the paper.

A.1   TEXMORPH BENCHMARK

Our new morphing benchmark, TexMorph, leverages high-precision synthetic object models crafted by artists and 3D object models captured from real scenes, forming a diverse dataset that spans multiple object categories, all constructed in Blender.

We utilize the $BlenderNeRF$ plugin to define a spherical orbit path for an active camera(COS) around the object. Training frames are rendered by uniformly sampling random camera views oriented toward the center. The dataset includes over ten categories of objects, such as synthetic and real-world collected (scanning model and photo) fruits, animals, furniture, vehicles, and more(see in Figure 8). We utilize this benchmark to conduct both qualitative and quantitative tests on the baselines mentioned below, evaluating the superiority of our method.

Figure 8: Examples of objects from the texmorph dataset, emphasizing the diversity in texture and structural features for morphing assessment

A.2   GEODESIC DISTANCE APPROXIMATION

Computing exact geodesic distances on meshes is computationally expensive for large-scale morphing. We approximate them using a hybrid graph representation that balances accuracy and efficiency. We construct two complementary graphs: the adjacency graph G_{adj} ensures topological consistency by connecting face-adjacent vertices, while the KNN graph G_{knn} provides local geometric awareness for improved distance approximation in sparse regions. The adjacency graph G_{adj} encodes face-sharing connectivity:

$$G_{\text{adj}} = \big\{(v_i, v_j) \mid v_i, v_j \text{ share a face in } F\big\}. \tag{13}$$

The KNN graph G_{knn} captures local Euclidean proximity:

$$G_{\text{knn}} = \big\{(v_i, v_j) \mid d(v_i, v_j) \leq \text{NN-distance}, i \neq j\big\}, \tag{14}$$

where $d(\cdot, \cdot)$ denotes Euclidean distance.

Combining G_{adj} and G_{knn}, we construct a hybrid distance matrix $D_{\text{adj}} \in \mathbb{R}^{n \times n}$:

$$D_{\text{adj}}(i, j) = \begin{cases} d(v_i, v_j), & \text{if } (v_i, v_j) \in G_{\text{adj}} \cup G_{\text{knn}}, \\ \infty, & \text{otherwise.} \end{cases} \tag{15}$$

The geodesic distance $D_{\text{g}}(i, j)$ between vertices v_i and v_j is then computed via Dijkstra's algorithm:

$$D_{\text{g}}(i, j) = \min_{P \in \mathcal{P}(v_i, v_j)} \sum_{(v_k, v_{k+1}) \in P} D_{\text{adj}}(v_k, v_{k+1}), \tag{16}$$

where $\mathcal{P}(v_i, v_j)$ is the set of all paths between v_i and v_j.

A.3 EXPERIMENT METRIC DETAILS

For structural similarity, Structural Similarity Index (SSIM) measures the similarity in shape and structure between the Morphing result and the target,we use MSE (SSIM) to measure the deviation between the actual SSIM curve and the ideal linear curve. For color consistency, ΔE is used to measure the color difference between the source object and target object, ensuring consistency in color. Finally, for edge continuity, Edge Integrity (EI) evaluates the continuity and completeness of edges during the shape morphing process, ensuring that the generated structures maintain consistent and unbroken boundaries.

SSIM for Structural Similarity

To ensure smooth and natural shape transitions during 3D morphing, we measure the Structural Similarity Index (SSIM) variation across different morphing stages. Ideally, SSIM should change linearly from the source shape A to the target shape B.

In an ideal scenario, we define the expected SSIM values at any morphing stage t as follows:

$$SSIM_{\text{ideal}}(A, G_t) = 1 - t, \tag{17}$$

$$SSIM_{\text{ideal}}(G_t, B) = t, \tag{18}$$

where G_t represents the intermediate shape at stage t. This ensures a smooth, gradual transition from A to B. For example, at specific morphing stages (At 30%): $SSIM(A, G_{30\%}) = 0.7$, $SSIM(G_{30\%}, B) = 0.3$.

To quantify how closely the actual SSIM values follow the ideal linear transition, we compute the Mean Squared Error (MSE) for each stage t:

$$\mathcal{E} = \frac{1}{N} \sum_{t \in T} \Big(\text{SSIM}_{\text{ideal}}(A, G_t) - \text{SSIM}_{\text{actual}}(A, G_t)\Big)^2 + \Big(\text{SSIM}_{\text{ideal}}(G_t, B) - \text{SSIM}_{\text{actual}}(G_t, B)\Big)^2 \tag{19}$$

A smaller error indicates that the SSIM variation is nearly linear, reflecting high-quality 3D morphing with smooth transitions and minimal distortion, whereas a larger error suggests anomalous SSIM changes, potentially indicating irregularities or distortions in the 3D morphing process.

ΔE for Color Consistenc

We evaluate color consistency using the ΔE metric in CIELAB space, calculating the average ΔE for each frame against the source, target, and adjacent frames. The CIELAB color space is chosen for its perceptual uniformity, where Euclidean distances correspond more closely to human color perception compared to RGB space.

The ΔE metric quantifies the perceptual difference between two colors and is defined as:

$$\Delta E_{ab}^* = \sqrt{(L_1^* - L_2^*)^2 + (a_1^* - a_2^*)^2 + (b_1^* - b_2^*)^2}, \tag{20}$$

where L^* represents lightness (0-100), a^* represents the green-red axis, and b^* represents the blue-yellow axis in CIELAB space.

For morphing evaluation, we compute three types of color consistency metrics: **source consistency** (ΔE_{source}) measuring deviation from the source image, **target consistency** (ΔE_{target}) evaluating progression toward the target, and **temporal consistency** (ΔE_{diff}) assessing smoothness between consecutive frames.

The final color consistency score is computed as:

$$\Delta E_{avg} = \frac{1}{3}(\bar{\Delta E}source + \bar{\Delta E}target + \bar{\Delta E}_{diff}). \tag{21}$$

Lower ΔE_{avg} values indicate better color consistency throughout the morphing sequence.

EI for Edge Continuity

Edge Integrity (EI) quantifies edge fragmentation by counting connected edge components after Canny edge detection. This metric evaluates structural quality and object boundary preservation in morphed images. EI is computed as:

$$EI = N_{Edges}(Canny(I, T_{low}, T_{high})) - 1. \tag{22}$$

where N_{Edges} represents the number of connected edge components, and the subtraction of 1 excludes the background component. Higher EI values indicate more fragmented edges, suggesting potential structural artifacts in the morphing sequence.

A.4 USER STUDY

To evaluate the 3D morphing quality from a human perspective, we conducted a user study with 54 participants. Each participant viewed 17 questions on multiple pairs of objects, randomly selected from our method and three baseline techniques, to evaluate texture and geometric shape comparisons, as well as the ablation results of our mesh-guided strategy and geometric distortion loss. They were asked to select the best set of results based on the following criteria: structural similarity, color consistency, edge continuity, and overall quality. The questionnaire used in our user study, designed to evaluate the quality and effectiveness of the morphing results, is shown in the Figure below:

User Study for Morphing
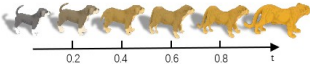
We have generated morphing results at 20%, 40%, 60%, and 80% transformation levels. Please select the morphing method that best meets the given criteria:
✓ Structural Similarity
Does the method preserve the structure of the intermediate frames (20%, 40%, 60%, 80%) well?
✓ Texture Consistency
Do the intermediate frames exhibit smooth and natural color transitions? Are there any abrupt or unnatural color jumps?
✓ Edge Continuity
Are the edges in the intermediate frames smooth and continuous? Are there any breaks, artifacts, or unnatural edge distortions?
✓ Overall Score
Considering "Structural Similarity", "Texture Consistency", and "Edge Continuity", please give an overall score for the morphing results of this method.

Please select the method that you think best achieves the transformation from "dog" to "lion" based on the following different results.
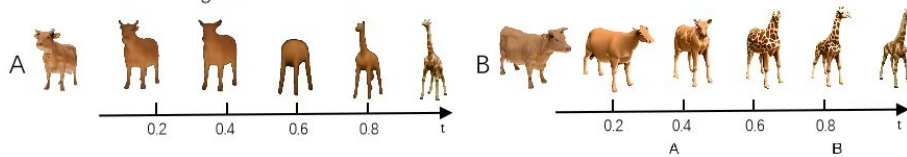
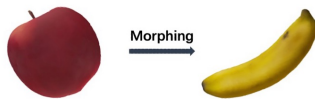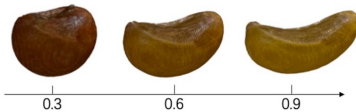Please select the method that you think best achieves the transformation from "dog" to "lion" based on the following different results.

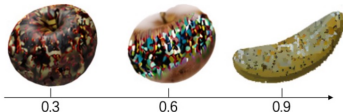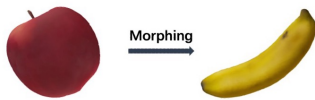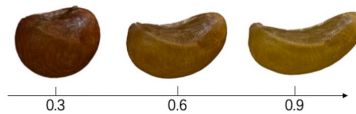|	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

Please select the method that you think best achieves the transformation from "dog" to "lion" based on the following different results.

	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

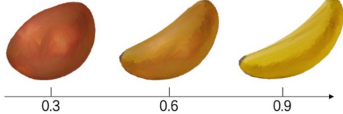Please select the method that you think best achieves the transformation from "cow" to "giraffe" based on the following different results.

A

0.2   0.4   0.6   0.8   t

B

0.2   0.4   0.6   0.8   t
A   B

	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

Please select the method that you think best achieves the transformation from "cow" to "giraffe" based on the following different results.

A

0.2   0.4   0.6   0.8   t

B

0.2   0.4   0.6   0.8   t
A   B

	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

Morphing

Please select the method that you think best achieves the transformation from "apple" to "banana" based on the following different results.
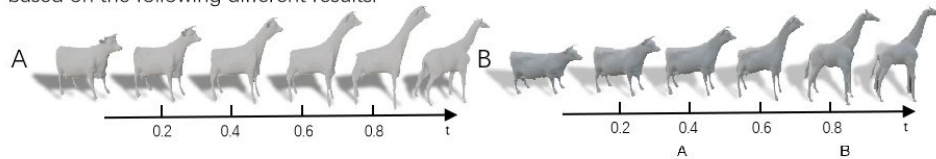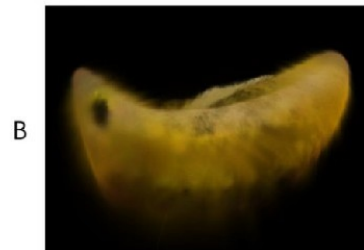
A

0.3   0.6   0.9

B

0.3   0.6   0.9

	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

Morphing

Please select the method that you think best achieves the transformation from "apple" to "banana" based on the following different results.

A

0.3   0.6   0.9

B

0.3   0.6   0.9

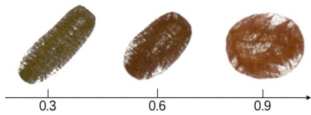|	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

Please select the method that you think best achieves the transformation from "cow" to "giraffe" based on the following different results.

	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

The two images below show the 90% transformation of an apple into a banana using different methods. Which one do you think has better results?

The reason you think this image has better results is...

- ☐ Structural stability
- ☐ Edge continuity
- ☐ Smooth color

Please select the method that you think best achieves the transformation from "zucchini" to "orange" based on the following different results.

A
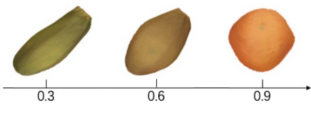
0.3　　　0.6　　　0.9

B

0.3　　　0.6　　　0.9

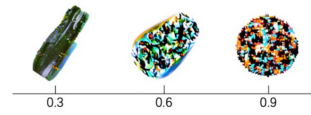|	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○

Please select the method that you think best achieves the transformation from "zucchini" to "orange" based on the following different results.

A

0.3　　　0.6　　　0.9

B

0.3　　　0.6　　　0.9

	A	B
The method with better Structural Similarity during the morphing process is	○	○
The method with better Texture Consistency during the morphing process is	○	○
The method with better Edge Continuity during the morphing process is	○	○
The method with better Overall Score is	○	○