MOTION RETARGETING BASED ON INTER-CHARACTER 3D GAUSSIAN MAPPING PAIRS

Anonymous authors
Paper under double-blind review

ABSTRACT

Motion retargeting aims to model the correlation between source and target characters and transfer the source motion to the target character. Conventional deep learning methods typically take the source and target meshes as a condition to directly generate target motion. These methods require large numbers of training samples, and their retargeting performance drops substantially when the two characters differ greatly in skeleton or morphology. To address these problems, we propose a motion retargeting method based on inter-character 3D Gaussian Mapping Pairs (GMPs). Our method transfers motion through the body-part primitives (BPPs) of the source and target characters. GMPs represent the mapping between their BPPs and enable source-to-target motion retargeting without requiring large numbers of training samples. We design an evolutionary framework consisting of GMP-Selection and GMP-Update to optimize GMPs. GMP-Selection uses a fitness function based on morphological and homeomorphic constraints to stabilize the search for GMPs. GMP-Update uses the same constraints to guide Gaussian adjustment, cloning, and splitting to accelerate GMP refinement. We also design a "Gaussian Binding"–"Motion Correction" strategy. It ensures motion fidelity, prevents interpenetration, and maintains continuity during motion retargeting by using tracking, collision, and smoothness constraints. Our method provides an interpretable inter-character representation for motion retargeting and supports characters with large differences. Extensive experiments on the public Mixamo dataset show that our method delivers leading performance.

1 INTRODUCTION

Motion retargeting transfers source motion to a target character for animation, games, digital humans, and virtual reality. Modeling inter-character correlation helps fit the target morphology, avoid interpenetration, and reproduce the source motion naturally and accurately.

Conventional deep learning methods encode the source and target meshes, learn their correlation, and directly use it as a condition to generate target motion. Their performance drops substantially across large skeletal or morphological differences. Figure 1(a) illustrates the resulting errors: target movement differs from the source in frames 3 and 4, and the hand penetrates the head in frame 5.

A character can be represented as a set of body-part primitives (BPPs), whose topological relationships remain stable during motion. Constructing an inter-character BPP homeomorphic mapping therefore allows source BPP motion to be transferred to corresponding target BPPs. In Figure 1(b), the target movement matches the source in frames 3 and 4, and the hand–head interpenetration in frame 5 is avoided.

We construct this mapping with Gaussian Mapping Pairs (GMPs) linking source and target BPPs represented by 3D Gaussians (Zwicker et al., 2001; Kerbl et al., 2023). Morphological constraints fit the Gaussians to their respective character volumes, while homeomorphic constraints maintain consistent local neighborhoods across paired Gaussians (Figure 2).

We design an evolutionary framework consisting of GMP-Selection and GMP-Update. GMP-Selection uses a genetic algorithm with a fitness function based on morphological and homeomorphic constraints to select GMPs for update and stabilize the search. GMP-Update uses the same constraints to guide Gaussian adjustment, cloning, and splitting to accelerate refinement. The opti-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

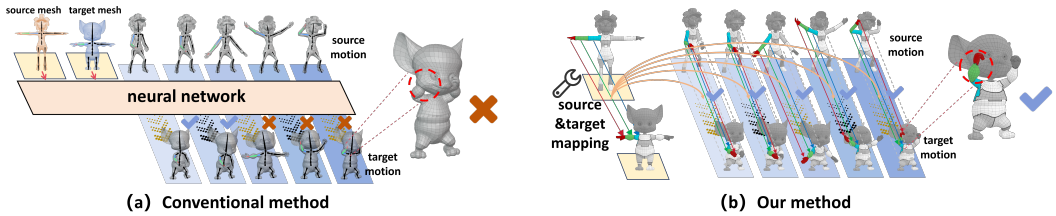

Figure 1: (a) A generative method produces motion errors and interpenetration. (b) BPP mapping preserves motion and reduces interpenetration.

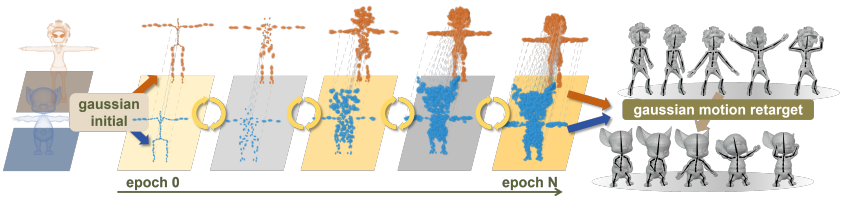

Figure 2: Overall pipeline for motion retargeting based on inter-character GMPs.

mized GMPs guide the "Gaussian Binding"–"Motion Correction" strategy: Gaussians are bound to each character's skeleton, and target skeletal joint rotations are optimized under tracking, collision, and smoothness constraints.

Our contributions are summarized as follows:

- We represent source and target BPPs with 3D Gaussians and use inter-character GMPs as a physically interpretable bridge for motion retargeting, supporting characters with large morphological differences.
- We design an evolutionary framework, which uses morphological and homeomorphic constraints to stabilize the search in GMP-Selection and guide Gaussian adjustment, cloning, and splitting in GMP-Update, thereby accelerating GMP refinement.
- We propose a "Gaussian Binding"–"Motion Correction" strategy using GMPs as the mapping reference. Optimizing target skeletal joint rotations under tracking, collision, and smoothness constraints preserves motion fidelity, reduces interpenetration, and maintains temporal continuity.

2 Related Work

2.1 Skeleton-Based Motion Retargeting

Conventional methods preserve end-effector, trajectory, or style constraints through spacetime optimization, hierarchical editing, online inverse kinematics, intermediate skeletons, or physics-based filtering (Gleicher, 1998; Lee & Shin, 1999; Choi & Ko, 2000; Monzani et al., 2000; Tak & Ko, 2005). Morphology-independent authoring also supports motion reuse when the rig is unknown at authoring time (Hecker et al., 2008). These methods offer explicit control but depend on manual correspondences or predefined structures.

Learning-based methods use shared motion representations. NKN combines recurrent prediction, differentiable forward kinematics, and cycle consistency (Villegas et al., 2018); PMnet separates pose from global movement (Lim et al., 2019); SAN uses skeleton-aware convolution and pooling (Aberman et al., 2020); and dilated temporal convolution provides another unsupervised formulation (Kim et al., 2020). Later approaches improve generality through topology-conditioned Transformers (Mourot et al., 2023), correspondence-free surface and skeleton features (Rekik et al., 2024), body-part attention (Hu et al., 2024), adversarial cross-morphology embeddings (Li et al., 2023), or latent-

space flow matching (Kim et al., 2025). These models do not always represent the skinned surface explicitly, so large shape differences can cause motion inconsistencies or interpenetration.

2.2 Geometry-Aware Motion Retargeting for Skinned Characters

Interaction meshes (Ho et al., 2010), Aura Mesh (Jin et al., 2018), and surface context graphs (Liu et al., 2018) preserve spatial relations through surface or volumetric proxies. Contact-aware retargeting detects self-contact and ground contact while penalizing interpenetration (Villegas et al., 2021). These proxies primarily support local preservation or collision handling, rather than a unified source–target correspondence.

Recent methods combine skeletal motion with geometry through residual corrections (R2ET; Zhang et al., 2023), masked pose modeling and collision optimization (MoMa; Martinelli et al., 2024), or semantic sensors and dense interaction fields (MeshRet; Ye et al., 2024). CAR addresses collision avoidance across different skeletal architectures (Cao & Yang, 2024), while SMRNet learns a skeleton–mesh representation (Zhang et al., 2025). ReConForM uses sparse rigged-surface features for contact-aware optimization (Cheynel et al., 2025), and lightweight spherical geometry enables online joint-level control (Guo et al., 2025). These approaches improve morphology adaptation and collision handling, but their dependence on learned representations, semantic features, or rigging relationships limits explicit, interpretable mapping across large shape differences.

2.3 Spatiotemporal Continuity

Frame-by-frame geometric correction can break trajectories or introduce jitter without temporal continuity. Physics-based retargeting promotes dynamic plausibility through simulation and reinforcement learning (Reda et al., 2023). STaR combines dense shape and limb-interpenetration constraints with a temporal Transformer and multilevel consistency losses for poses and trajectories (Yang et al., 2025). These studies motivate joint treatment of geometry, motion semantics, and temporal continuity through a unified cross-character mapping.

3 Method

3.1 Overview

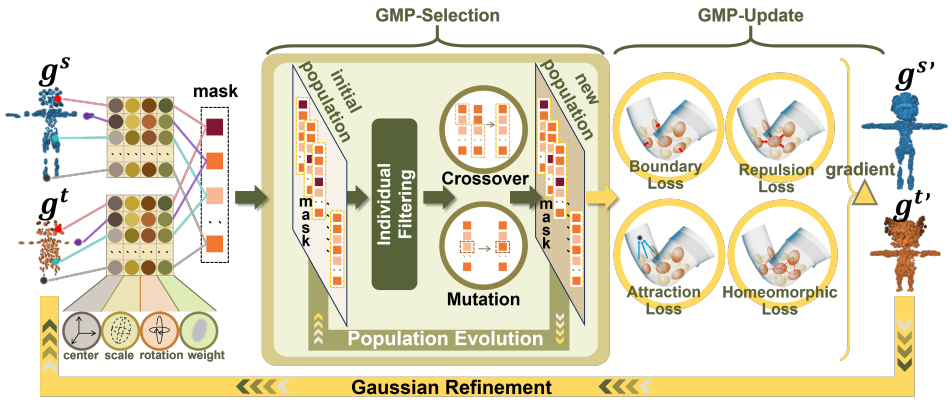

Figure 3: Evolutionary framework of GMP optimization. GMP-Selection produces a shared mask M for Gaussian sets g^s, g^t; GMP-Update refines the selected pairs into $g^{s\prime}, g^{t\prime}$, which feed the next round.

Let $c \in \{s, t\}$ denote the source or target character. We construct GMPs from their T-pose meshes, representing BPPs with anisotropic 3D Gaussians: g_i^c denotes Gaussian i in character c, $P_i = (g_i^s, g_i^t)$ is the ith GMP, and $\Phi_G = \{P_i\}_{i=1}^N$ is the set of N pairs. Each pair links T-pose primitives; the optimized correspondence is reused during motion.

We optimize Φ_G through alternating GMP-Selection and GMP-Update (Figure 3). GMP-Selection uses shared binary masks and a genetic algorithm with morphological and homeomorphic constraints to select GMPs for update and stabilize the search. GMP-Update uses the same constraints to guide Gaussian adjustment, cloning, and splitting, accelerating GMP refinement.

During retargeting, both characters' Gaussians are bound to their respective skeletons. GMPs guide target skeletal joint rotation optimization under tracking, collision, and smoothness constraints.

3.2 Morphological and Homeomorphic Constraints

The morphological constraint makes the BPPs fill each character's mesh interior, while the homeomorphic constraint preserves corresponding local neighborhoods across characters. Together, they evaluate GMPs during selection and guide their update (Figure 4). Gaussian g_i^c has center $x_i^c \in \mathbb{R}^3$, unit-quaternion rotation $q_i^c \in \mathbb{R}^4$, scales $s_i^c \in \mathbb{R}^3_{>0}$, and volume-contribution weight $o_i^c \in \mathbb{R}_{\geq 0}$. Its effective radius along a unit direction $u \in \mathbb{R}^3$ is $r_i^c(u)$.

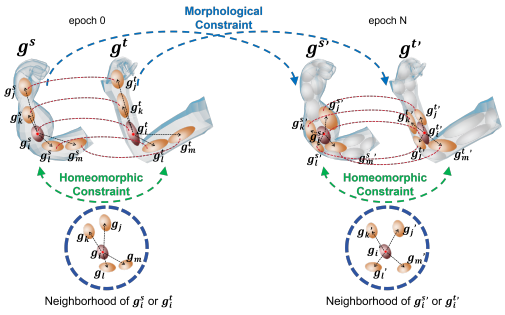

Figure 4: Morphological and homeomorphic constraints before and after GMP optimization.

3.2.1 Morphological Constraint

The morphological constraint consists of boundary, attraction, and repulsion losses. They keep Gaussians within the character volume, guide them toward insufficiently covered regions, and prevent local accumulation, respectively.

Boundary loss. Let d_i^c be the distance from x_i^c to the closest mesh triangle and n_i^c its outward unit normal. We penalize Gaussian support beyond the boundary:

$$\mathcal{L}_{\text{bnd}}^c = \frac{1}{N} \sum_{i=1}^{N} o_i^c \left[\text{ReLU}\left(r_i^c(n_i^c) - d_i^c\right)\right]^2. \tag{1}$$

Attraction loss. Let $\mathcal{P}^c$ contain points sampled in low-density interior regions, $\mathcal{N}_K^c(p)$ index the K Gaussians closest to point $p \in \mathcal{P}^c$, and u_{ip}^c be the unit direction from x_i^c to p. The loss pulls Gaussian support toward uncovered points:

$$\mathcal{L}_{\text{att}}^c = \frac{1}{K|\mathcal{P}^c|} \sum_{p \in \mathcal{P}^c} \sum_{i \in \mathcal{N}_K^c(p)} o_i^c \left[\text{ReLU}\left(\|x_i^c - p\|_2 - r_i^c(u_{ip}^c)\right)\right]^2. \tag{2}$$

Repulsion loss. For neighboring Gaussian index pairs $(i, j) \in \mathcal{E}^c$, let u_{ij}^c be the unit direction from x_i^c to x_j^c. We penalize overlap along this direction:

$$\mathcal{L}_{\text{rep}}^c = \frac{1}{|\mathcal{E}^c|} \sum_{(i,j) \in \mathcal{E}^c} o_i^c o_j^c \left[\text{ReLU}\left(r_i^c(u_{ij}^c) + r_j^c(-u_{ij}^c) - \|x_i^c - x_j^c\|_2\right)\right]^2. \tag{3}$$

The morphological loss is $\mathcal{L}_{\text{geo}}^c = \lambda_{\text{bnd}}\mathcal{L}_{\text{bnd}}^c + \lambda_{\text{att}}\mathcal{L}_{\text{att}}^c + \lambda_{\text{rep}}\mathcal{L}_{\text{rep}}^c$, where λ_{bnd}, λ_{att}, and λ_{rep} weight boundary, attraction, and repulsion, respectively.

3.2.2 Homeomorphic Constraint

The homeomorphic constraint encourages paired source and target Gaussians to maintain consistent local neighborhoods throughout GMP optimization. It consists of context consistency and orientation consistency losses, which preserve neighbor identities and align corresponding neighbor directions, respectively.

Context consistency loss. Let $\mathcal{N}_i^0$ be the fixed set of K neighboring GMP indices for pair i. In character c, $d_{i,\text{in}}^c$ is the distance to its farthest member and $d_{i,\text{out}}^c$ the distance to the nearest Gaussian outside that set, excluding i. We penalize outsiders displacing original neighbors:

$$\mathcal{L}_{\text{ctx}} = \frac{1}{N} \sum_{c\in\{s,t\}} \sum_{i=1}^{N} \left[\text{ReLU}\left(d_{i,\text{in}}^c - d_{i,\text{out}}^c\right)\right]^2. \tag{4}$$

Orientation consistency loss. For $j \in \mathcal{N}_i^0$, let $v_{ij}^c \in \mathbb{R}^3$ be the unit vector pointing from Gaussian center i to center j in character c. Comparing these vectors preserves local arrangement independently of distance:

$$\mathcal{L}_{\text{ori}} = \frac{1}{NK} \sum_{i=1}^{N} \sum_{j\in\mathcal{N}_i^0} \|v_{ij}^s - v_{ij}^t\|_2^2. \tag{5}$$

The homeomorphic loss is $\mathcal{L}_{\text{nbr}} = \lambda_{\text{ctx}}\mathcal{L}_{\text{ctx}} + \lambda_{\text{ori}}\mathcal{L}_{\text{ori}}$, where λ_{ctx} and λ_{ori} weight context consistency and orientation consistency, respectively.

3.3 Evolutionary Framework for GMP Optimization

GMP-Selection determines which pairs to update; GMP-Update adjusts, clones, or splits them. Both use the constraints in Section 3.2.

3.3.1 Genetic Algorithm-Based GMP-Selection

GMP-Selection evolves P shared binary masks through population initialization, individual filtering, and crossover and mutation (Figure 5).

Population initialization. Candidate $p \in \{1, \ldots, P\}$ is a mask $M^{(p)} \in \{0,1\}^N$. Entry $m_i^{(p)} = 1$ selects GMP i for update, while 0 leaves it unchanged. Nonuniform sampling initializes the masks using the mean source and target volume-contribution weight of each GMP.

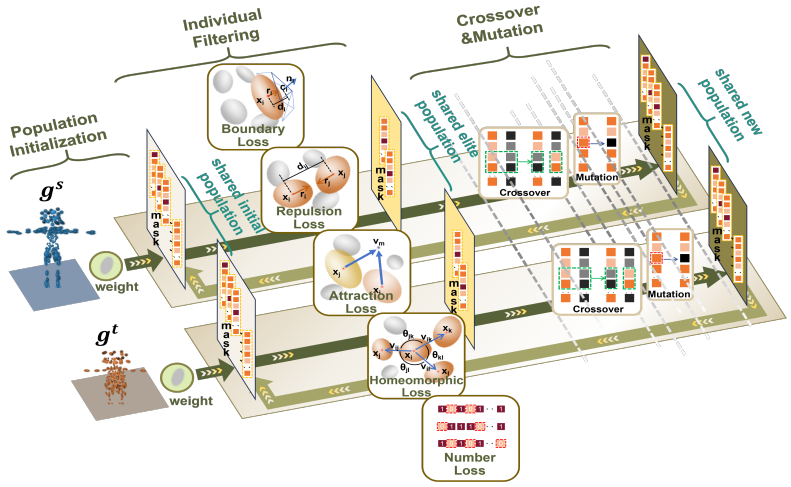

Figure 5: GMP-Selection. volume-contribution weights of g^s, g^t initialize shared masks $M^{(p)}$. Fitness $F(M)$ selects elites; crossover and mutation generate the next population.

Individual filtering. For a nonempty mask M with entries m_i, the selected count is $N(M) = \sum_{i=1}^{N} m_i$, and the selected fraction is $\mathcal{L}_{\text{num}}(M) = N(M)/N$. Let $\mathcal{L}_{\text{geo},i}^c$ and $\mathcal{L}_{\text{nbr},i}$ denote per-GMP morphological and homeomorphic loss contributions. The fitness is

$$F(M) = -\left[\frac{1}{N(M)} \sum_{i=1}^{N} m_i \left(\mathcal{L}_{\text{geo},i}^s + \mathcal{L}_{\text{geo},i}^t + \lambda_{\text{nbr}}\mathcal{L}_{\text{nbr},i}\right) + \lambda_{\text{num}}\mathcal{L}_{\text{num}}(M)\right]. \tag{6}$$

The source and target morphological losses enter equally; λ_{nbr} weights homeomorphic consistency and λ_{num} penalizes the selected fraction. The top fraction p_e of masks survives as elites.

Crossover and mutation. Crossover recombines elite masks with probability p_c, and mutation flips bits with probability p_m to promote exploration. Fitness is recomputed after GMP-Update. Population evolution continues until fitness convergence or the iteration limit, yielding the highest-fitness mask.

3.3.2 Multi-Constraint Loss-Based GMP-Update

GMP-Update computes a multi-constraint loss for the selected GMPs using the morphological and homeomorphic constraints (Figure 6). Adjustment changes Gaussian parameters, while cloning and splitting add GMPs to fill insufficiently covered regions and fit finer volume details.

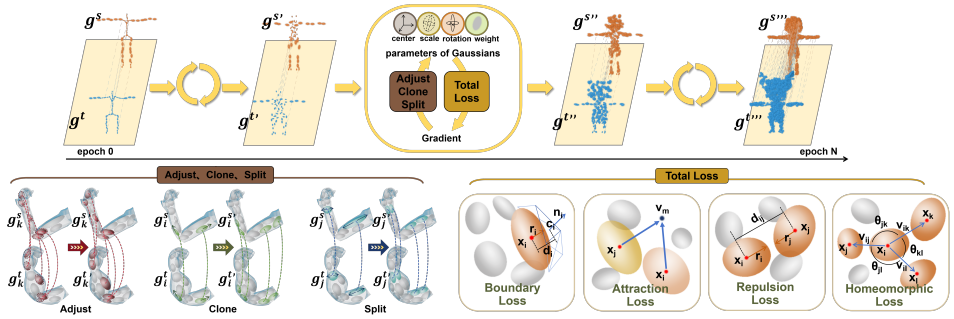

Figure 6: GMP-Update. The multi-constraint gradient $\nabla_{\theta_i^c}\mathcal{L}_{\text{total}}$ guides adjustment, cloning, and splitting of g^s, g^t into $g^{s\prime}, g^{t\prime}$, then $g^{s\prime\prime}, g^{t\prime\prime}$ in subsequent rounds.

Adjustment. The parameter tuple $\theta_i^c = (x_i^c, q_i^c, s_i^c, o_i^c) \in \mathbb{R}^{11}$ collects center, rotation, scales, and weight. The construction objective is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{geo}}^s + \mathcal{L}_{\text{geo}}^t + \lambda_{\text{nbr}}\mathcal{L}_{\text{nbr}}, \tag{7}$$

with the same λ_{nbr} as in fitness. Gradient descent updates both Gaussians in each active pair with learning rate η; inactive pairs remain fixed.

Cloning and splitting. Cloning and splitting add GMPs according to center gradients and scales. Let G_i be the mean center-gradient magnitude of the source and target Gaussians, and S_i the mean of their largest axis scales. Let ϵ_{grad} and ϵ_{scale} denote the gradient and scale thresholds. When $G_i > \epsilon_{\text{grad}}$ and $S_i \le \epsilon_{\text{scale}}$, the GMP is cloned, with each new Gaussian inheriting its respective parent's parameters. When $G_i > \epsilon_{\text{grad}}$ and $S_i > \epsilon_{\text{scale}}$, it is split along its principal axes into two smaller GMPs. Both operations are synchronized across characters to preserve the mapping relation.

3.4 GMP-Based Motion Retargeting

Given N GMPs Φ_G^* optimized in T-pose and source rotations $Q^s \in \mathbb{R}^{T\times J_s\times 4}$ over T frames, we use the "Gaussian Binding"–"Motion Correction" strategy. Frame τ indexes the source and target motions, with $c \in \{s, t\}$ identifying the character.

Gaussian Binding. We bind each character's Gaussians to its own skeleton using linear blend skinning (LBS). For J_c joints, $Q^c(\tau) \in \mathbb{R}^{J_c\times 4}$ stacks unit quaternions $Q_j^c(\tau)$. Binding weights $w_{ij}^c \ge 0$, with $\sum_j w_{ij}^c = 1$, are interpolated from nearby mesh vertices. Let $\text{FK}_j(Q^c(\tau); K^c)$ be joint j's global affine transform from forward kinematics on skeleton K^c, and $B_j^{c,0}$ its T-pose transform. FK and Gaussian LBS give

$$\begin{aligned}\Delta B_j^c(\tau) &= \text{FK}_j(Q^c(\tau); K^c) \circ (B_j^{c,0})^{-1},\\ x_i^c(\tau) &= \sum_{j=1}^{J_c} w_{ij}^c \Delta B_j^c(\tau)[x_i^{c,0}], \quad u_i^c(\tau) = \frac{x_i^c(\tau) - x_i^{c,0}}{\sigma_{\ell_i}^c}.\end{aligned} \tag{8}$$

Here, $x_i^{c,0} \in \mathbb{R}^3$ is the optimized T-pose center, ℓ_i is the body-part label, and $\sigma_{\ell_i}^c > 0$ is its characteristic scale. Source rotations are given; target rotations are optimized. Normalized displacements $u_i^c(\tau) \in \mathbb{R}^3$ are comparable across body proportions; their dependence on Q^c is implicit below.

Motion Correction. Using the bound Gaussians, we optimize the target skeletal joint rotation sequence $Q^t \in \mathbb{R}^{T\times J_t\times 4}$. GMPs identify corresponding source and target displacements for the following constraints.

The tracking constraint compares normalized displacements of paired Gaussians to preserve motion fidelity:

$$\mathcal{E}_{\text{track}}(Q^t; Q^s) = \frac{1}{TN}\sum_{\tau=1}^{T}\sum_{i=1}^{N}\left\|u_i^t(\tau) - u_i^s(\tau)\right\|_2^2. \tag{9}$$

The collision constraint uses the same signed separation as the repulsion loss in Eq. 3: center distance minus the sum of facing effective radii. Let $\mathcal{C}^t$ contain non-adjacent target BPP pairs and $\delta_{ij}^t(\tau)$ denote this clearance on the bound target Gaussians. Negative clearance indicates interpenetration:

$$\mathcal{E}_{\text{col}}(Q^t) = \frac{1}{T|\mathcal{C}^t|}\sum_{\tau=1}^{T}\sum_{(i,j)\in\mathcal{C}^t}\left[\text{ReLU}\left(-\delta_{ij}^t(\tau)\right)\right]^2. \tag{10}$$

The smoothness constraint uses $R_j^t(\tau) = \text{Rot}(Q_j^t(\tau))$ in $\text{SO}(3)$, where Rot converts a unit quaternion to a 3×3 local rotation matrix. Defining $\Delta^2 R_j^t(\tau) = R_j^t(\tau+1) - 2R_j^t(\tau) + R_j^t(\tau-1)$, we penalize second temporal differences with the Frobenius norm $\|\cdot\|_{\text{F}}$:

$$\mathcal{E}_{\text{smo}}(Q^t) = \frac{1}{(T-2)J_t}\sum_{\tau=2}^{T-1}\sum_{j=1}^{J_t}\left\|\Delta^2 R_j^t(\tau)\right\|_{\text{F}}^2. \tag{11}$$

The weighted objective jointly optimizes all frames over unit quaternions:

$$Q^{t,*} = \arg\min_{Q^t}\ \left[\lambda_{\text{track}}\mathcal{E}_{\text{track}}(Q^t; Q^s) + \lambda_{\text{col}}\mathcal{E}_{\text{col}}(Q^t) + \lambda_{\text{smo}}\mathcal{E}_{\text{smo}}(Q^t)\right]. \tag{12}$$

The output $Q^{t,*}$ contains the optimized target local joint quaternions. Target mesh vertices can then be obtained by LBS. Further explanations are provided in the separate appendix.

4 Experiments

4.1 Experimental Setup

Dataset. We use 2,752 non-overlapping clips from 18 Mixamo characters (Adobe Systems Inc., 2020) with diverse skeletons, mesh topologies, body shapes, and motions. Each source motion is retargeted to multiple targets; results are averaged over valid frames, clips, and character pairs. Meshes are made watertight, root position, orientation, coordinates, and scale are normalized; body parts are partitioned by joint name.

Implementation details. Each character is divided into 22 body parts, with $K = 8$ neighbors per Gaussian. GMP construction uses 500 warm-up and at most 19,500 further iterations, with early stopping. Our PyTorch implementation runs on an NVIDIA RTX 5090. GMP-Selection uses $P = 32$ masks, crossover probability $p_c = 0.9$, and mutation probability $p_m = 0.1$. The weights $(\lambda_{\text{bnd}}, \lambda_{\text{att}}, \lambda_{\text{rep}}, \lambda_{\text{ctx}}, \lambda_{\text{ori}})$ are $(2, 5, 1, 2, 3)$, with learning rate $\eta = 0.003$. Target skeletal joint rotations are optimized with $(\lambda_{\text{track}}, \lambda_{\text{col}}, \lambda_{\text{smo}}) = (4, 2, 1)$ and learning rate 0.03.

Evaluation metrics. We measure character-height-normalized mean squared joint error (MSE), root-aligned local joint error (MSE_{lc}), the percentage of limb vertices penetrating non-adjacent body parts (Pen), and mean squared distance between target samples whose source counterparts contact (Con). Following STaR (Yang et al., 2025), Curv measures the change in average squared limb-joint acceleration relative to reference motion. Lower is better for all metrics.

4.2 Comparison with Baselines

We compare with Copy, NKN (Villegas et al., 2018), SAN (Aberman et al., 2020), R2ET (Zhang et al., 2023), MeshRet (Ye et al., 2024), and STaR (Yang et al., 2025), covering direct kinematic copying, skeleton-based retargeting, and geometry-aware retargeting.

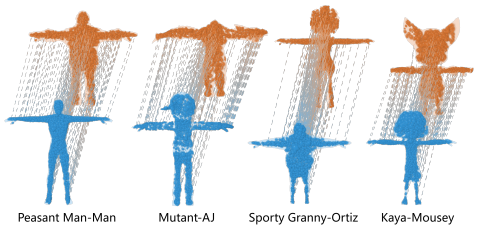

Figure 7: Optimized GMPs for four character pairs. Each line connects the source and target Gaussians in one GMP.

Table 1: Quantitative comparison with baselines.

Method	MSE ↓	MSE_{lc} ↓	Pen (%) ↓	Con ↓
Source	–	–	3.641	–
Copy	**0.068**	**0.045**	5.013	0.959
NKN	0.092	0.063	6.284	1.767
SAN	0.069	**0.045**	4.674	1.974
R2ET	0.085	0.047	4.369	1.271
MeshRet	0.277	0.069	5.247	0.902
STaR	0.221	0.058	4.915	0.863
Ours	0.197	0.061	**4.113**	**0.835**

Qualitative results. Figure 7 shows examples of optimized GMPs. The correspondences adapt to large body-proportion differences while preserving the homeomorphic BPP mapping. Figure 8 separately compares the retargeted motions: baselines cause hand–body interpenetration or shift the original contact, whereas our constraints better preserve contact and surface separation.



Figure 8: Qualitative comparison of retargeted motions. Our method better preserves contact and reduces hand–body interpenetration across different target morphologies.

Quantitative results. Table 1 shows that Copy has the lowest MSE but increases Pen over the source motion. Ours achieves the lowest Pen and Con among retargeting methods, improving the respective strongest baselines by 5.9% and 3.2%. Its MSE is also lower than those of MeshRet and STaR.

4.3 Ablation Studies

GMP-Selection. GMP-Selection uses constraint-based fitness to select shared masks, reducing redundant updates and lowering the final loss for all four character pairs (Figure 9) while reducing construction time from 9,520 s to 7,532 s, or 20.9% (Table 2).

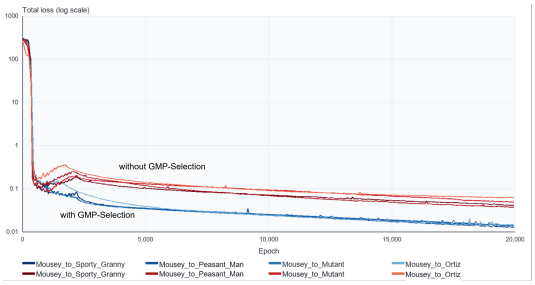

Figure 9: GMP loss with and without GMP-Selection on four character pairs.

Table 2: Effect of GMP-Selection on construction time.

Selection strategy	Build (s) ↓
Without GMP-Selection	9,520
With GMP-Selection	**7,532**

Table 3: Ablation of the three motion constraints.

Objective	MSE ↓	Pen (%) ↓	Con ↓	Curv ↓
Tracking	**0.172**	4.835	0.902	1.038
+ Collision	0.204	**4.068**	0.863	0.887
+ Collision + Smoothness	0.197	4.113	**0.835**	**0.864**

Table 4: Accuracy–efficiency trade-off across GMP counts.

GMP count	Build (s) ↓	FPS ↑	MSE ↓	Pen (%) ↓	Con ↓
100–500	**5,324**	**17.8**	**0.191**	4.432	0.898
500–2,000	7,532	13.7	0.197	**4.113**	**0.835**
2,000–5,000	13,440	6.3	0.195	4.130	0.841

Tracking, collision, and smoothness constraints. Tracking alone gives the lowest MSE (Table 3). Adding collision reduces Pen and Con by 15.9% and 4.3%, respectively, with increased MSE. Adding smoothness improves MSE and yields the lowest Con and Curv, while slightly raising Pen. The constraints balance fidelity, geometry, and continuity.

Number of GMPs. With 100–500 GMPs, construction is fastest, FPS highest, and MSE lowest (Table 4). Using 500–2,000 yields the lowest Pen and Con at 13.7 FPS; 2,000–5,000 lowers MSE but increases cost, Pen, and Con. We favor 500–2,000 GMPs for this geometry–efficiency trade-off.

4.4 User Study

We recruited 100 participants. Each rated five randomly ordered, anonymized clips on motion semantics preservation, geometric plausibility, and overall quality. Our method received higher ratings than the baselines (see the separate appendix).

5 Conclusion

This work addresses the motion retargeting task by representing source and target BPPs with 3D Gaussians and optimizing inter-character GMPs through an evolutionary framework. The designed morphological and homeomorphic constraints stabilize the search in GMP-Selection and guide adjustment, cloning, and splitting in GMP-Update. The "Gaussian Binding"–"Motion Correction" strategy optimizes target skeletal joint rotations under tracking, collision, and smoothness constraints. Experiments on Mixamo show the lowest Pen and Con among the compared retargeting methods, with ablations supporting the constraint design and chosen GMP count range.

Limitations and future work. The stability and computational cost of GMP optimization leave room for improvement. Future work will improve construction stability and runtime efficiency and reduce the effect of source-motion noise on retargeting.

AI USE STATEMENT

AI tools were used to assist with English translation and language polishing, as well as the retrieval and discovery of relevant literature. The authors take full responsibility for the final text, citations, and scientific claims.

REPRODUCIBILITY STATEMENT

To facilitate reproducibility, Section 3 describes GMP construction and the motion retargeting objectives. Section 4 reports the dataset, preprocessing, evaluation metrics, optimization settings, and hardware used in our experiments. The separate supplementary material provides a notation and formula guide, additional implementation details, and the user-study setup.

REFERENCES

Kfir Aberman, Peizhuo Li, Dani Lischinski, Olga Sorkine-Hornung, Daniel Cohen-Or, and Baoquan Chen. Skeleton-aware networks for deep motion retargeting. *ACM Transactions on Graphics*, 39 (4):1–14, 2020. doi: 10.1145/3386569.3392462.

Adobe Systems Inc. Mixamo. `https://www.mixamo.com/`, 2020. Accessed: 2026-07-27.

Yu Cao and Minghui Yang. CAR: Collision-avoidance retargeting for varied skeletal architectures. In *Proceedings of the ACM SIGGRAPH Asia Technical Communications*, pp. 1–4. ACM, 2024. doi: 10.1145/3681758.3698007.

Théo Cheynel, Thomas Rossi, Baptiste Bellot-Gurlet, Damien Rohmer, and Marie-Paule Cani. ReConForM: Real-time contact-aware motion retargeting for more diverse character morphologies. *Computer Graphics Forum*, 44(2):e70028, 2025. doi: 10.1111/cgf.70028.

Kwang-Jin Choi and Hyeong-Seok Ko. On-line motion retargetting. *The Journal of Visualization and Computer Animation*, 11(5):223–235, 2000.

Michael Gleicher. Retargetting motion to new characters. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 33–42. ACM, 1998. doi: 10.1145/280814.280820.

Tianze Guo, Zhedong Chen, Yi Jiang, Linjun Wu, Xilei Wei, Lang Xu, Yeshuang Lin, He Wang, and Xiaogang Jin. Ultrafast and controllable online motion retargeting for game scenarios. *ACM Transactions on Graphics*, 44(6):1–11, 2025. doi: 10.1145/3763351.

Chris Hecker, Bernd Raabe, Ryan W. Enslow, John DeWeese, Jordan Maynard, and Kees van Prooijen. Real-time motion retargeting to highly varied user-created morphologies. *ACM Transactions on Graphics*, 27(3):1–11, 2008. doi: 10.1145/1360612.1360626.

Edmond S. L. Ho, Taku Komura, and Chiew-Lan Tai. Spatial relationship preserving character motion adaptation. *ACM Transactions on Graphics*, 29(4):1–8, 2010. doi: 10.1145/1778765.1778770.

Lei Hu, Zihao Zhang, Chongyang Zhong, Boyuan Jiang, and Shihong Xia. Pose-aware attention network for flexible motion retargeting by body part. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):4792–4808, 2024. doi: 10.1109/TVCG.2023.3277918.

Taeil Jin, Meekyoung Kim, and Sung-Hee Lee. Aura mesh: Motion retargeting to preserve the spatial relationships between skinned characters. *Computer Graphics Forum*, 37(2):311–320, 2018. doi: 10.1111/cgf.13363.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. doi: 10.1145/3592433.

SangBin Kim, Inbum Park, Seongsu Kwon, and JungHyun Han. Motion retargetting based on dilated convolutions and skeleton-specific loss functions. *Computer Graphics Forum*, 39(2):497–507, 2020. doi: 10.1111/cgf.13947.

Wontaek Kim, Tianyu Li, and Sehoon Ha. MoReFlow: Motion retargeting learning through unsupervised flow matching, 2025.

Jehee Lee and Sung Yong Shin. A hierarchical approach to interactive motion editing for human-like figures. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 39–48. ACM, 1999. doi: 10.1145/311535.311539.

Tianyu Li, Jungdam Won, Alexander Clegg, Jeonghwan Kim, Akshara Rai, and Sehoon Ha. ACE: Adversarial correspondence embedding for cross-morphology motion retargeting from human to nonhuman characters. In *Proceedings of the ACM SIGGRAPH Asia Conference*, pp. 1–11. ACM, 2023. doi: 10.1145/3610548.3618255.

Jongin Lim, Hyung Jin Chang, and Jin Young Choi. PMnet: Learning of disentangled pose and movement for unsupervised motion retargeting. In *Proceedings of the British Machine Vision Conference*, 2019.

Zhiguang Liu, Antonio Mucherino, Ludovic Hoyet, and Franck Multon. Surface-based motion retargeting by preserving spatial relationship. In *Proceedings of the ACM SIGGRAPH Conference on Motion, Interaction and Games*, pp. 1–11. ACM, 2018. doi: 10.1145/3274247.3274507.

Giulia Martinelli, Nicola Garau, Niccolò Bisagno, and Nicola Conci. MoMa: Skinned motion retargeting using masked pose modeling. *Computer Vision and Image Understanding*, 249:104141, 2024. doi: 10.1016/j.cviu.2024.104141.

Jean-Sébastien Monzani, Paolo Baerlocher, Ronan Boulic, and Daniel Thalmann. Using an intermediate skeleton and inverse kinematics for motion retargeting. *Computer Graphics Forum*, 19 (3):11–19, 2000. doi: 10.1111/1467-8659.00393.

Lucas Mourot, Ludovic Hoyet, François Le Clerc, and Pierre Hellier. HuMoT: Human motion representation using topology-agnostic transformers for character animation retargeting, 2023.

Daniele Reda, Jungdam Won, Yuting Ye, Michiel van de Panne, and Alexander Winkler. Physics-based motion retargeting from sparse inputs. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–19, 2023. doi: 10.1145/3606928.

Rim Rekik, Mathieu Marsot, Anne-Hélène Olivier, Jean-Sébastien Franco, and Stefanie Wuhrer. Correspondence-free online human motion retargeting. In *Proceedings of the International Conference on 3D Vision*, pp. 707–716, 2024. doi: 10.1109/3DV62453.2024.00032.

Seyoon Tak and Hyeong-Seok Ko. A physically-based motion retargeting filter. *ACM Transactions on Graphics*, 24(1):98–117, 2005. doi: 10.1145/1037957.1037963.

Ruben Villegas, Jimei Yang, Duygu Ceylan, and Honglak Lee. Neural kinematic networks for unsupervised motion retargetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8639–8648, 2018. doi: 10.1109/CVPR.2018.00901.

Ruben Villegas, Duygu Ceylan, Aaron Hertzmann, Jimei Yang, and Jun Saito. Contact-aware retargeting of skinned motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9720–9729, 2021. doi: 10.1109/ICCV48922.2021.00958.

Xiaohang Yang, Qing Wang, Jiahao Yang, Gregory Slabaugh, and Shanxin Yuan. STaR: Seamless spatial-temporal aware motion retargeting with penetration and consistency constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12947–12955, 2025. doi: 10.1109/ICCV51701.2025.01203.

Zijie Ye, Jia-Wei Liu, Jia Jia, Shikun Sun, and Mike Zheng Shou. Skinned motion retargeting with dense geometric interaction perception. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-4000.

Jia-Qi Zhang, Miao Wang, Fu-Cheng Zhang, and Fang-Lue Zhang. Skinned motion retargeting with preservation of body part relationships. *IEEE Transactions on Visualization and Computer Graphics*, 31(9):4923–4936, 2025. doi: 10.1109/TVCG.2024.3423426.

Jiaxu Zhang, Junwu Weng, Di Kang, Fang Zhao, Shaoli Huang, Xuefei Zhe, Linchao Bao, Ying Shan, Jue Wang, and Zhigang Tu. Skinned motion retargeting with residual perception of motion semantics & geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13864–13872, 2023. doi: 10.1109/CVPR52729.2023.01332.

Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. EWA volume splatting. In *Proceedings of IEEE Visualization*, pp. 29–36, 2001. doi: 10.1109/VISUAL.2001.964492.