

ReConForM : Real-time Contact-aware Motion Retargeting for more Diverse Character Morphologies

Théo Cheynel^{1,2}, Tomas Rossi¹, Baptiste Bellot-Gurlet¹, Damien Rohmer², and Marie-Paule Cani²

¹Kinetix

²LIX, École Polytechnique, CNRS, Institut Polytechnique de Paris




Figure 1: Output of our retargeting method, showcasing several of our contributions. The source pose (left scene) shows complex contacts carrying semantic information : self-contacts, foot-ground contacts, and inter-character contacts). This pose is retargeted onto very different characters from the animated movie industry, evolving on a non-flat terrain.

Abstract

Preserving semantics, in particular in terms of contacts, is a key challenge when retargeting motion between characters of different morphologies. Our solution relies on a low-dimensional embedding of the character's mesh, based on rigged key vertices that are automatically transferred from the source to the target. Motion descriptors are extracted from the trajectories of these key vertices, providing an embedding that contains combined semantic information about both shape and pose. A novel, adaptive algorithm is then used to automatically select and weight the most relevant features over time, enabling us to efficiently optimize the target motion until it conforms to these constraints, so as to preserve the semantics of the source motion. Our solution allows extensions to several novel use-cases where morphology and mesh contacts were previously overlooked, such as multi-character retargeting and motion transfer on uneven terrains. As our results show, our method is able to achieve real-time retargeting onto a wide variety of characters. Extensive experiments and comparison with state-of-the-art methods using several relevant metrics demonstrate improved results, both in terms of motion smoothness and contact accuracy.

CCS Concepts

• **Computing methodologies** → **Animation**; Motion processing;

1. Introduction

Motion transfer between animated characters, also called motion retargeting, is a crucial aspect of character animation with major applications in the fields of cinema, video games and virtual reality. The complexity of this task arises from the variety of 3D character models. While different skeletal topologies and skin meshes may make them difficult to compare, retargeting remains an ill-posed problem even when a correspondence exists between the source and target models. Differences in morphological proportions may cause self-penetrations, inaccurate contacts, and a loss of the overall sense of the pose. Finding the right compromises to mitigate these issues while maintaining the expected fidelity to the source motion is a complex task. In particular, well-thought-out metrics, capturing at least in part the semantics of motion, need to be designed. Alternatively, machine learning shows promise in enabling pose semantics to be discovered and retargeted without any explicit definition. However, the lack of paired retargeting data makes the use of supervised learning difficult.

Our work was greatly inspired by the evidence that the preservation of contacts between different body parts is of utmost importance when humans assess the quality of a retargeted motion [BOH*22]. We argue that contacts with the ground, present in almost every motion, are equally important. It would, however, be very time-consuming to manually identify the most relevant contact interactions of a specific movement, which may change over time, and then find a compromise between their influences in order to preserve the relevant mesh contacts. Instead, we introduce specific motion features to describe the relative positioning of body parts and propose an automatic, adaptive solution that dynamically selects and weights these features over time, focusing on periods around key events such as collisions and contacts. The selected weighted features are then used to compute lightweight objective functions, optimized in real time to generate the target motion. In short, our main contributions are:

- A new joint representation for the character's morphology and motion, used to compute relevant motion features;
- An adaptive proximity-based method to select and weight motion features over time;
- Extensions to multi-character retargeting and non-flat grounds, enabled by the adaptability of our framework.

The evaluation, both qualitative via a user study and quantitative thanks to various metrics, shows that our method achieves better results than the state of the art, both in terms of temporal fluidity of the target animation and semantic similarity with the source. Furthermore, adaptively selected constraints being sparse in space and time, our method (named ReConForM) achieves retargeting tasks several orders of magnitude faster than former optimization-based methods, while accommodating arbitrary morphologies and being easily extensible to new use cases.

2. Related Work

2.1. Optimization-based retargeting

A first category of methods formulate motion retargeting tasks as an optimization problem. Gleicher [Gle98] pioneered this approach with a solution requiring the manual identification of a

number of spacetime constraints. Still only considering skeletal motion, the method was improved thanks to inverse kinematics (IK) solvers [LS99], the addition of constraints to the IK framework [CK00], and of dynamic constraints to improve the respect of physical laws [TK05].


Noting that character morphology should have a significant impact on motion, recent optimization-based methods addressed motion retargeting for skinned characters. Jin et al. [JKL17] introduced the concept of Aura Mesh to retarget two characters whose meshes are interacting, and Basset et al. [BWBM20] optimized energy functions based on volume preservation and collision management. However, both methods are limited to parametric skin mesh models such as SMPL [LMR*15], used to directly pair the vertices of the source and target meshes. Ho et al. [HKT10] performed multi-character motion retargeting using an optimization method based on interacting joints and hard constraints to prevent collisions of bone-attached primitives (serving as character shape proxies). However, this method cannot retarget motion to characters of arbitrary morphologies, as the iterative motion morphing step requires the source and target characters to have identical skeleton topology.

2.2. Learning-based retargeting

Other methods leveraged machine learning techniques for motion retargeting. Shon et al. [SGHR05] modeled common latent representation of motion using Gaussian process regression, but this method required datasets of paired motion data for both characters, which are usually unavailable. To alleviate this issue, Villegas et al. [VYCL18] introduced a cycle consistency loss for an adversarial unsupervised training framework. Lim et al. [LCC19] learned to disentangle pose (joints' local coordinates) from movement (the root bone's global trajectory). By encoding motion in the shared latent space of a common primal skeleton [ALL*20; AWL*19; HZZ*23; YML23], retargeting was performed in deep feature space, thus increasing generalization to various skeletons. Yet, as these methods only considered the skeleton and not the mesh, they failed to capture self-contacts, despite them being a key component of motion.

In contrast, Villegas et al. [VCH*21] detected self-contacts and used geometry-conditioned RNN to preserve them while optimizing motion in a latent space. Despite yielding good results, this method may not be adapted to industrial use cases as the model cannot handle arbitrary skeleton topology and is not suitable for real-time applications. Besides, comparison with this method is difficult as no implementation is available online. More recently, R²ET [ZWK*23] introduced a shape-aware module that improves self-contacts as a post-processing step on the target mesh, and avoids unwanted penetrations, in real-time. Their method can handle very diverse characters, but collisions on the source mesh are not taken into account, meaning that semantic information can be lost during retargeting.

Recent works [ZCX*24] focus on motion semantics by performing a differentiable rendering of the animated character's mesh, and encoding the video through a vision-language model. However, this complex setup has only been shown to perform on simple motions



(a) Penetration with the floor (legs, belly), hands floating above ground (b) Colliding limbs (arms and legs) (c) Self-collision on the thigh, feet floating above ground (d) Hand floating above ground

Figure 2: Examples of issues found in NKN's dataset [VYCL18]. Figures 2a and 2b are called “ground-truth” although they were retargeted using Mixamo, while figures 2c and 2d show issues without having undergone any retargeting (original source characters for those motions).

that are easily transcribed into a textual description (e.g. “shrugging”, “waving”, “praying while standing up”), with little to no interaction between the limbs and minimal foot motion ; and requires the same skeleton between source and target characters.

Lastly, some recent works [RWY*23; ZGY*23] used a combination of deep reinforcement learning and physically-based simulation to allow for a physically accurate retargeting, taking mesh contacts into account. The main issue is that of generalization: unlike previously mentioned methods, the model has to be retrained for each new target character, making it impossible to scale up to the diversity of virtual characters used in most 3D computer graphics applications.

2.3. Mesh correspondence

Some methods [ZBP20; WWF*20; LYS*22] aim to disentangle the pose and shape of a mesh to transfer poses to a target with a different shape. Because we have skeletal animations and a common standard pose amongst our characters, this disentanglement is trivial. However, we still need a way to compare shapes by computing a correspondence between the source and target meshes (see [VZHC11] and [Sah20] for surveys on this topic). A recent family of solutions relied on optimal transport. Solomon et al. [SPKS16] used a Gromov-Wasserstein distance with an entropy term to find a dense mapping between 3D meshes. Mandad et al. [MCK*17] increased efficiency by using a coarse-to-fine approach relying on diffusion geometry, while Schmidt et al. [SPK23] used an automatic triangulation to approximate the source and target surfaces. In this work, we take inspiration from these techniques to compute a correspondence between the source and target models.

2.4. Retargeting datasets

Recent works in the field of motion retargeting [LCC19; ZWK*23; ALL*20] made use of the dataset introduced in NKN [VYCL18], with 2400 motion sequences sampled from Mixamo [Ado24]. We visually assessed its quality, and expose two major shortcomings. First, the variety of characters was obtained using Mixamo's in-house retargeting algorithm, despite many of these retargeted animations showing degraded motion semantics (where the intention

behind a movement is completely lost), severe self-collisions, and ground collisions. Examples of such issues are shown in Figures 2a and 2b. Second, some of the original motions from Mixamo show severe flaws on their original character (before any retargeting), as illustrated by Figures 2c and 2d. Therefore, we claim that using this dataset as ground truth, with metrics such as the MSE of joint positions, is not an accurate way to estimate the quality of motion retargeting methods. It may even have impacted the quality of previous work: while it was used for unsupervised training by Villegas et al. [VYCL18], some other methods used this dataset as evaluation data, training data [LCC19; ZWK*23] and even as ground-truth for supervised learning [ALL*20].

In contrast, we use a few quality metrics such as self-penetration, foot sliding, and jerk, to assess the quality of the animations we generate, while referring to human feedback, through a user-study, to assess the overall semantic-preserving quality of our method.

3. Shape and motion descriptors

Let us consider a *source* and a *target* humanoid characters, both with a skin-mesh rigged to a skeleton. Neither the skeletons or the meshes need to have the same topologies. The skeletons may have different numbers of joints so long as a *bone mapping*, i.e. a list of paired bones on both characters, is predefined; and the meshes can exhibit different numbers of vertices and large shape differences. In addition, an input, kinematic animation to be transferred is provided for the source character.

To compute relevant, semantic-based motion transfer from source to target, some common representation for the characters' shapes and a set of motion descriptors for the input animation are required. Tracking the distance between relevant pairs of points was already used as a simple and robust way to encode the semantic of motion in the fields of robotics [HVV*20] and motion retargeting [HKT10; ZGY*23]. Using points located on the skin mesh instead of skeletal joints captures finer semantic information and helps preventing collisions [KNK*02]. Taking inspiration from these works, we propose a light morphological representation based on key-vertices, described in section 3.1, and then use it to build a set of time-varying descriptors encoding the successive poses in the input animation.

3.1. Sparse shape encoding and correspondence

The ReConForM method uses a generic humanoid template mesh (based on the SMPL model [LMR*15]), on which we pre-selected N specific vertices, called *key-vertices* ($N = 41$ unless mentioned otherwise, see Figure 3). The latter are chosen as to provide (i) a sparse yet comprehensive coverage of the character's surface, and (ii) a good sampling of typical areas prone to contact, such as hands and feet. While our method was developed for humanoid characters, its principle could easily be extended to different categories of characters, just by changing this template model.

To provide a sparse one-to-one correspondence between the source and target model, we automatically transfer the template's keypoints to their two skin meshes, as follows. Taking inspiration from optimal transport between arbitrary meshes [SPKS16], we view the task as an optimization problem. We input both the template and destination (i.e., source or target) meshes in their T-pose. Using the skinning weights, we first split the mesh into its various limbs (arms, legs, torso, head, feet and hands). For each limb, we consider the position of all N^l vertices, and normalize them to have zero mean and unit variance. For each limb l , we convert the template and destination point-clouds into two distributions :

$$\mathcal{T} = \frac{1}{N_t^l} \sum_{i=1}^{N_t^l} \delta_{v_{i,t}^l} \quad \text{and} \quad \mathcal{D} = \frac{1}{N_d^l} \sum_{j=1}^{N_d^l} \delta_{v_{j,d}^l}$$

where $v_{i,t}^l$ and $v_{j,d}^l$ are, respectively, the normalized position of the i -th and j -th vertices of limb l for the template and destination meshes, and δ is the Dirac distribution.

The constraints of our setup (namely, that the meshes are in T-pose and normalized) allow us to use a simple euclidian distance as criterion, instead of relying on the Gromov-Wasserstein distance used in Solomon et al. [SPKS16]. Our goal is to find an optimal transport plan, which is a matrix $P^l \in \mathbb{R}_+^{N_t^l \times N_d^l}$ defining a coupling between the two distributions $\mathcal{T}$ and $\mathcal{D}$.

In order to account for an uneven distribution of vertices, we weight each vertex by the inverse of the local vertex density :

$$w_i^l = \sum_{f \in \mathcal{F}_i} \frac{A(f)}{n(f)}$$

where $\mathcal{F}_i$ is the set of faces containing vertex i , $A(f)$ the area of face f , and $n(f)$ its number of vertices.

We solve the following optimization problem:

$$\begin{aligned} P^{l*} &= \underset{P^l}{\operatorname{argmin}} \langle P^l, C^l \rangle = \underset{P^l}{\operatorname{argmin}} \sum_{i=1}^{N_t^l} \sum_{j=1}^{N_d^l} P_{i,j}^l C_{i,j}^l \\ \text{s.t. } \mathbb{1} \cdot P^l &= (w_{0,d}^l; \dots; w_{N_d^l,d}^l) \\ P^l \cdot \mathbb{1}^T &= (w_{0,t}^l; \dots; w_{N_t^l,t}^l)^T \end{aligned}$$

where $C_{i,j}^l = \|v_{i,t}^l - v_{j,d}^l\|^2$, and $\mathbb{1}$ is a vector filled with ones.

Similarly to Feydy et al. [FSV*19], we approximate this optimal transport plan by solving the entropy regularized version of the problem, which is faster to compute using Sinkhorn's algorithm

[Cut13]. Thus, for each key-vertex on the template, we are thus able to get a corresponding vertex on the destination mesh, which effectively gives us the position of the key-vertices for both the source and target characters.

Examples of key-vertex transfer is shown in Figure 3, and results on additional characters are shown in the Supp. Mat, along with results for our extended template with $N = 96$ keypoints. Note that the output of this automatic transfer can be overridden, or further refined if needed via interactive authoring. This enables the user to fine-tune key-vertex positions on the source, target, or even on the template shape, to account for very large morphological differences, or to highlight specific areas where careful sampling needs to be considered.

3.2. Motion encoding via time-varying pose descriptors

Using the key-vertices defined before, time-varying pose descriptors can be computed on the source animation to extract the semantic elements, in order to transfer them to the target character. Yet,



Figure 3: Chosen location of key-vertices on the template mesh (top-left), and results of the key vertices transfer to several character from Mixamo [Ado24]. Key vertices are shown with corresponding colored spheres to show the automatic transfer to the different meshes.

not all descriptors convey the same amount of semantic information: it is necessary to adaptively focus on the pairs of key-vertices interacting with each other or with the ground.

Given an animated character, we call p_i (or $p_{i,t}$ explicitly when necessary) the 3D position of the $i^{th} \in \llbracket 1, N \rrbracket$ key-vertex at time t , and n_i its normal to the mesh surface. We compute three time-varying descriptors $\mathcal{M}_{\text{dist}}$, $\mathcal{M}_{\text{dir}}$, and $\mathcal{M}_{\text{pen}}$, each stored in a matrix-like structure representing the pairwise relationship between all possible pairs of key-vertices $(i, j) \in \llbracket 1, N \rrbracket^2$ of a character.

These descriptors respectively encode:

- The distances between all pairs of key-vertices,

$$\mathcal{M}_{\text{dist}} \in \mathbb{R}^{N \times N}, \text{ where } \mathcal{M}_{\text{dist}}(i, j) = \|p_j - p_i\|;$$

- The vector, and therefore the offset direction, between pairs of key-vertices,

$$\mathcal{M}_{\text{dir}} \in \mathbb{R}^{N \times N \times 3}, \text{ where } \mathcal{M}_{\text{dir}}(i, j) = p_j - p_i;$$

- The signed distance between two key vertices i and j , computed along the mesh normal at vertex i , which locally serves as a measure of penetration,

$$\mathcal{M}_{\text{pen}} \in \mathbb{R}^{N \times N}, \text{ where } \mathcal{M}_{\text{pen}}(i, j) = n_i \cdot (p_j - p_i).$$

Two additional descriptors, $\mathcal{M}_{\text{height}}$ and $\mathcal{M}_{\text{sliding}}$, are used to represent the relationship between individual key-vertices and the environment, considered to be an horizontal ground floor at this stage. Calling $\vec{up}$ the unit vertical direction, these two descriptors are stored in vector-like containers encoding:

- The height of each key-vertex with respect to the ground,

$$\mathcal{M}_{\text{height}} \in \mathbb{R}^{N \times 1}, \text{ where } \mathcal{M}_{\text{height}}(i) = \vec{up} \cdot p_i;$$

- The horizontal velocity of each key-vertex,

$$\mathcal{M}_{\text{sliding}} \in \mathbb{R}^{N \times 2}, \text{ where } \mathcal{M}_{\text{sliding}}(i) = \frac{H(p_{i,t+\Delta t}) - H(p_{i,t})}{\Delta t}$$

$$\text{and } H(p_i) = p_i - (\vec{up} \cdot p_i) \vec{up}.$$

4. Adaptive weighting of constraints

The motion descriptors defined above are computed at each time on the source animation and used as constraints to be matched on the target animation. Trying to maintain all these constraints at once might however be unfeasible when the source and target morphologies are different, as they would conflict with one another. Instead, our insight is that, depending on the input animation, only a few of these descriptors are significant at each time step to the perceived semantics of motion. We therefore introduce a new, dynamic weighting formulation to select only the relevant ones over time and weight them accordingly.

4.1. Weighting coefficients

Our descriptor weighting method exploits the notion of spatial proximity between pairwise key-vertices (indicating an interaction between limbs), and of height relative to the ground (indicating a likely occurrence of floor contact). We therefore make use of two

weighting characteristics, also stored in matrix/vector-like structures with values in $[0, 1]$, obtained via a clamped normalized representation of the distances between, and heights of key-vertices:

$$\mathcal{W}_{\text{interaction}}(i, j) = \text{clamp} \left(1 - \frac{\mathcal{M}_{\text{dist}}(i, j) - d_{\min}}{d_{\max} - d_{\min}} \right)$$

$$\mathcal{W}_{\text{floor}}(i) = \text{clamp} \left(1 - \frac{\mathcal{M}_{\text{height}}(i) - h_{\min}}{h_{\max} - h_{\min}} \right)$$

where $\text{clamp}(x) = \min(1, \max(0, x))$ and $d_{\min}$, $d_{\max}$, and $h_{\min}$, $h_{\max}$ are per-character constant thresholds characterizing typical distances for which nearby limb interaction or ground contact should be considered. For a character of height h_c , we typically set $d_{\min} = h_{\min} = 5\% h_c$, and $d_{\max} = h_{\max} = 15\% h_c$.

These weights provide a measure of the importance of each constraint, at each animation frame. More precisely, $\mathcal{W}_{\text{interaction}}(i, j)$ (resp. $\mathcal{W}_{\text{floor}}(i)$) will be close to 1 when a constraint between two interacting limbs (resp. between a limb and the floor) must be taken into account, and be 0 for unimportant relationships to be discarded. Furthermore, even though a pair of key-vertices may interact at a specific animation frame, the associated weight decreases back to 0 as soon as they spread apart. The resulting sparsity of constraints, in space and time, is a key-advantage for the efficiency of our solution.

4.2. Sparse objective to be optimized

We formulate the retargeting task as an optimization problem over the skeletal poses variables defined by the local joint rotation $\mathbf{q}$ and root joint position in world space p , which are related to the key-vertex positions p_i and normal n_i through linear-blend skinning.

The optimization starts with a rough initialization of $\mathbf{q}$ and p through a naive retargeting, obtained by copying the rotation of the source bones onto their equivalent target bones using the *bone mapping* mentioned in Section 3.

We then optimize $(\mathbf{q}, p)$ over the entire animation at once to minimize the combination of the following three losses :

$$\min_{\mathbf{q}, p} w_{\text{reg}} \mathcal{L}_{\text{reg}} + w_{\text{smooth}} \mathcal{L}_{\text{smooth}} + w_{\text{sem}} \mathcal{L}_{\text{sem}} \quad (1)$$

The regularization loss $\mathcal{L}_{\text{reg}} = \sum_{i \in \llbracket 1, N \rrbracket} \sum_t \|p_{i,t} - p_{i,t}^{\text{init}}\|^2$ is a simple Mean Squared Error (MSE) comparing the key-vertices' positions to that of the naive retargeting $p_{i,t}^{\text{init}}$, ensuring that the solution stays in the neighborhood of the initialization.

The smoothness loss $\mathcal{L}_{\text{smooth}} = \sum_{i \in \llbracket 1, N \rrbracket} \sum_t \|\ddot{p}_{i,t}\|$ ensures that the change in acceleration, or jerk, is as small as possible.

The semantic loss $\mathcal{L}_{\text{sem}}$ is defined as the weighted sum of the following terms (where $\circ$ denotes the element-wise product):

- A term ensuring that the distance between pairs of key-vertices is preserved:

$$\mathcal{L}_{\text{dist}} = \sum_t \left\| \mathcal{W}_{\text{interaction}}^{\text{src} \rightarrow \text{targ}} \circ (\mathcal{M}_{\text{dist}}^{\text{source}} - \mathcal{M}_{\text{dist}}^{\text{target}}) \right\|^2$$

- A term ensuring that the direction between pairs of key-vertices




Figure 4: Top: graphical representation of our method (left: key-vertices transfer presented in Section 3.1; middle: motion retargeting process presented in Sections 3.2 through 4.2). Top-right shows two poses being retargeted onto different characters, with key-vertices shown in red. Bottom shows a complex pose showcasing collisions with the floor as well as self-collisions (the leftmost character is the source character, all the others are target characters from Mixamo).

is preserved (with $S_{\cosine}$ being the element-wise cosine similarity):

$$\mathcal{L}_{\text{dir}} = \sum_t \left\| \mathcal{W}_{\text{interaction}}^{\text{src} \rightarrow \text{targ}} \circ S_{\cosine} \left(\mathcal{M}_{\text{dir}}^{\text{source}}, \mathcal{M}_{\text{dir}}^{\text{target}} \right) \right\|^2$$

- A term matching the amount of penetration of pairs of interacting key vertices, between source and target:

$$\mathcal{L}_{\text{pen}} = \sum_t \left\| \mathcal{W}_{\text{interaction}}^{\text{src} \rightarrow \text{targ}} \circ \left(\mathcal{M}_{\text{pen}}^{\text{source}} - \mathcal{M}_{\text{pen}}^{\text{target}} \right) \right\|^2$$

- A term penalizing key-vertices penetrating the floor, and encouraging similar heights on source and target for key-vertices close to the ground:

$$\mathcal{L}_{\text{height}} = \sum_t \left\| \min \left(0, \mathcal{M}_{\text{height}}^{\text{target}} \right) \right\|^2 + \sum_t \left\| \mathcal{W}_{\text{floor}}^{\text{src} \rightarrow \text{targ}} \circ \left(\mathcal{M}_{\text{height}}^{\text{source}} - \mathcal{M}_{\text{height}}^{\text{target}} \right) \right\|^2$$

- A term preventing key-vertices that are in contact with the ground from sliding across consecutive frames:

$$\mathcal{L}_{\text{sliding}} = \sum_t \left\| \mathcal{W}_{\text{floor}}^{\text{src} \rightarrow \text{targ}} \circ \left(\mathcal{M}_{\text{sliding}}^{\text{source}} - \mathcal{M}_{\text{sliding}}^{\text{target}} \right) \right\|^2$$

The matrices $\mathcal{W}_{\text{interaction}}^{\text{src} \rightarrow \text{targ}}$ and $\mathcal{W}_{\text{floor}}^{\text{src} \rightarrow \text{targ}}$ are a combination of the

weighting matrices of the source and target characters:

$$\begin{aligned} \mathcal{W}_{\text{interaction}}^{\text{src} \rightarrow \text{targ}} &= \mathcal{W}_{\text{interaction}}^{\text{source}} + \alpha \text{Sg} \left(\mathcal{W}_{\text{interaction}}^{\text{target}} \right), \\ \mathcal{W}_{\text{floor}}^{\text{src} \rightarrow \text{targ}} &= \mathcal{W}_{\text{floor}}^{\text{source}} + \alpha \text{Sg} \left(\mathcal{W}_{\text{floor}}^{\text{target}} \right), \end{aligned} \quad (2)$$

where the mixing parameter α increases linearly from 0 to 1 during the optimization process. Both $\mathcal{W}_{\text{interaction}}^{\text{target}}$ and $\mathcal{W}_{\text{floor}}^{\text{target}}$ are weighting matrices computed from the target's motion descriptors at the current optimization step. This evolution from source to target weighting helps to guide the numerical optimization toward a correct solution during the first steps, and prevent false positive interactions (i.e., ending up with contacts that were not present in the source motion) for later steps. $\text{Sg}()$ is the stop-gradient operation, through which gradient flow is prevented during the optimization to reduce computational time and selectively prevent the optimization process from attempting to minimize unrelated parameters.

The terms of $\mathcal{L}_{\text{sem}}$ were weighted manually to give each loss a roughly equal contribution at the start of the optimization. The values of the hyperparameters, which yield satisfying results regardless of the animation and characters, are provided in supplementary material.

5. Results

5.1. Implementation

We implemented our ReConForM method using the Pytorch framework to leverage the differentiability of the losses relative to

the variables $(\mathbf{q}, p)$. We optimized the target poses using Adam [KB14].

5.2. Evaluation data

We tested our method using animations and characters from Mixamo [Ado24]. The evaluation sets provided by previous works [ZWK*23; VYCL18; ALL*20] lack quality even on the source character. Instead, we manually selected 45 animations from Mixamo [Ado24] with accurate contacts and asked a 3D animation expert to annotate their estimated difficulty to be retargeted based on the amount of floor contacts and self-collisions present in them: 24 of them were quantified as *Hard*, 21 as *Easy*.

5.3. Quantitative results

5.3.1. Inference time

Our optimization objectives being very sparse, our method is able to consistently run in real-time (for motions longer than 3 seconds) on a laptop with a Nvidia RTX 3060 GPU, with an asymptotic speed of 67 frames per second. This makes it suitable for use in video games or motion capture pre-visualization. Since all frames are processed at once in a batched fashion, our method reaches higher framerates for longer animations. A three-second batch (around 75 frames) gives a good compromise between speed and delay for use cases where an online, on-the-fly retargeting is necessary, such as motion capture visualization. Figure 5 shows the inference speed for animations of various durations.




Figure 5: Speed and framerate of our retargeting method

5.3.2. Metrics

We evaluated, and report in Table 1, the performance of ReConForM through several metrics, comparing it to the naive baseline (“Copy Rotations” – see Supp. Mat. for implementation details), to SAN [ALL*20], as well as R²ET [ZWK*23] with and without the addition of a 1€ filter [CRV12] to reduce its noise. We also provide comparison with industrial retargeting algorithms of Mixamo [Ado24], Unreal Engine 5 [Epi20], and Autodesk Maya/MotionBuilder [Aut25]. Overall, our approach provides the best compromise to avoid self-penetration and foot skating, while preserving a smooth animation.

We evaluated the motion smoothness using the average and maximum magnitude of the jerk of joint positions throughout each animation sequence. Thanks to the integration of our smoothness loss

within the optimization, our approach achieves remarkably low jerk (same order of magnitude as the source), leading to very smooth motion.

Quantifying the self-penetration by computing the collision volume between the limbs, normalized by the total volume of the character, we found that our method surpasses the current state-of-the-art while remaining temporally smooth. On the other hand, R²ET, being solely dedicated to reducing self-collisions, shows higher temporal noise. Using a post-process filter [CRV12] helps reduce the noise, but spoils, in return, the amount of self-penetration.

Similarly, we evaluated the penetration with the floor by computing the relative proportion of the volume of the character that is below the ground at each frame. Our method shows a 70% reduction in floor penetration compared to the state-of-the-art.

We finally evaluated the quality of the feet-to-ground contacts along two quantitative measures. First, we call a foot to be *grounded* when its distance to the floor is between -1% and 1% of the character’s height. Second, we consider that the foot is *locked* on its position if its horizontal speed is lower than 0.1% of the character’s height per second. We computed the F1 and ROC AUC scores to quantify how much the semantic information of feet-to-ground contacts is kept by the retargeting method for these two metrics, with the source motion serving as ground truth. In both cases, our scores indicates that our approach generates more accurate feet-to-ground contact than the academic state-of-the-art, much closer to the quality of IK-based methods such as Mixamo and MotionBuilder. These methods indeed allow for a very precise placement of the feet, however they tend to do so without consideration for other semantic aspects of motion, leading for example to higher self-collisions and more frequent loss of semantic contacts.

5.3.3. Ablation study

We replaced the vertex transfer process described in section 3.1 by key-vertices placed manually by a human expert, to evaluate its impact on the overall retargeting method. We also evaluate with different amounts of key-vertices. These results (bottom lines of Table 1) confirm the robustness of our method with respect to the definition and placement of key-vertices. Additional experiments demonstrating the usefulness of each loss can be found in the supplementary material.

5.4. Qualitative results

5.4.1. Visual output

Figure 1 and Figure 4 show the results of our retargeting on several poses, and Figure 13 provides a visual comparison to R²ET. Additional results are displayed in the Supp. Mat.

5.4.2. User study

We designed a preferential blind user study to evaluate the quality of ReConForM compared to the state-of-the-art R²ET. We gathered a panel of 133 online respondents including both experts and non-experts. Each participant evaluated twelve randomly selected animations displayed as videos containing the source animation, the result of our retargeting, and the one obtained using

	Jerk (m.s^{-3}) ↓		Self-penetration (%) ↓		Floor penetration (%) ↓		Grounded feet		Foot sliding	
	mean	max	mean	max	mean	max	F1 ↑	ROC AUC ↑	F1 ↑	ROC AUC ↑
Source	254.9	3980.2	$1.68 \cdot 10^{-8}$	$1.31 \cdot 10^{-7}$	$1.54 \cdot 10^{-3}$	$9.08 \cdot 10^{-3}$	–	–	–	–
Mixamo	281.2	6495.44	$5.65 \cdot 10^{-8}$	$2.84 \cdot 10^{-7}$	$6.56 \cdot 10^{-3}$	$3.31 \cdot 10^{-2}$	0.944	0.908	0.923	0.900
Unreal Engine 5	261.2	4020.3	$9.50 \cdot 10^{-8}$	$4.74 \cdot 10^{-7}$	$1.76 \cdot 10^{-2}$	$7.14 \cdot 10^{-2}$	0.573	0.515	0.283	0.541
Maya + MotionBuilder	275.1	4364.8	$6.43 \cdot 10^{-8}$	$3.33 \cdot 10^{-7}$	$7.10 \cdot 10^{-3}$	$3.45 \cdot 10^{-2}$	0.945	0.922	0.928	0.927
Copy Rotations	258.6	3854.9	$1.05 \cdot 10^{-7}$	$4.82 \cdot 10^{-7}$	$9.01 \cdot 10^{-3}$	$3.92 \cdot 10^{-2}$	0.826	0.787	0.670	0.701
SAN	293.9	3803.1	$6.78 \cdot 10^{-8}$	$3.91 \cdot 10^{-7}$	$7.18 \cdot 10^{-3}$	$2.26 \cdot 10^{-2}$	0.302	0.508	0.139	0.510
R ² ET	1540.1	17249.9	$5.75 \cdot 10^{-8}$	$3.52 \cdot 10^{-7}$	$9.85 \cdot 10^{-3}$	$4.41 \cdot 10^{-2}$	0.828	0.790	0.505	0.641
R ² ET + 1€ filter	398.5	4628.6	$5.96 \cdot 10^{-8}$	$3.27 \cdot 10^{-7}$	$9.82 \cdot 10^{-3}$	$5.53 \cdot 10^{-2}$	0.773	0.708	0.379	0.586
Ours ($N = 41$)	212.8	2978.8	$4.65 \cdot 10^{-8}$	$3.60 \cdot 10^{-7}$	$3.22 \cdot 10^{-3}$	$1.72 \cdot 10^{-2}$	0.925	0.901	0.835	0.804
Ours ($N = 96$)	238.8	3358.4	$3.62 \cdot 10^{-8}$	$3.55 \cdot 10^{-7}$	$2.76 \cdot 10^{-3}$	$1.46 \cdot 10^{-2}$	0.908	0.888	0.820	0.795
Ours (w/o k-v transfer)	200.1	2560.5	$4.09 \cdot 10^{-8}$	$3.34 \cdot 10^{-7}$	$3.99 \cdot 10^{-3}$	$2.02 \cdot 10^{-2}$	0.926	0.900	0.836	0.805

Table 1: Metrics on our evaluation set. *Green* : best, *blue* : second best. Bottom two lines show the results of the ablation study mentioned in section 5.3.3.

R²ET [ZWK*23] (in random order) on one of eight different characters taken from Mixamo [Ado24]. Participants were asked to choose a preferred output along two questions: first with respect to *fidelity* with the source, and second, with respect to *pleasantness* of the motion.

Figure 6 summarizes the reported preferences of the users over all the animations (Overall) covering a total of 1596 answers, as well as on both *Easy* and *Hard* subsets. We observe a clear preference for our results, with our outputs obtaining the majority of

the votes on 41 out of 45 animations. Moreover, we can note that the average score on animations from the *Hard* subset exceeds that on the *Easy* subset. This makes us hypothesize that the benefits of using our approach increase with the complexity of the retargeting. A χ^2 test confirms this claim with $p < 0.01$.

Figure 7 shows another interesting result: the more respondents are familiar with the task of retargeting, the more they perceive our outputs to be best. A χ^2 test indicates a statistically significant distribution shift between the answers of experts and neophytes, with $p < 0.01$. All study details are in the supplementary material.




Figure 6: User preference across all answers, split by difficulty




Figure 7: User preference across all answers, split by expertise level

5.5. Additional characters

To demonstrate the robustness of our method on various characters, we provide additional results on production-ready characters provided by Blender Studio (CC-BY license). Despite their extreme proportions, and their different skeleton structure (fewer spine bones, different hierarchy, presence of additional “roll bones”...), our key-vertex transfer method provides a remarkably robust output. Several retargeted poses are shown in Figure 12. Whenever a good target pose is reachable, our optimization yields good results, and when the characters’ morphology prevents it from mimicking a complex pose, the outputs stay visually pleasing, without any catastrophic failure.

5.6. Failure cases

5.6.1. Key-vertices transfer

We identified several failure cases which might cause suboptimal results:

- **Lack of geometry:** low-poly character with few vertices can cause poor results because of a low number of vertices to choose from.
- **Asymmetry:** while usually insignificant, some extreme asymmetries in the mesh can “pull” the key-vertices of some limbs toward one side, causing an asymmetry in the key-vertices that might not always be preferred.

- **Accessories / Props:** character modeled with props (i.e. bags, hats, large haircuts...) can cause local errors on the placement of key vertices (directly on or close to the prop).
- **Overlapping meshes:** when meshes overlap, the transfer might place key-vertices on a “hidden” part of the mesh, which could prevent the retargeting process from considering the outer part of the mesh.

A more complete description of the failure cases, along with examples can be found in the Supplementary Material. We also propose several straightforward solutions to alleviate these specific issues.

5.6.2. Motion Retargeting

We also identified some failure cases of the retargeting process:

- **Complex poses:** poses with several points of contact (such as yoga poses) are harder to retarget, and finding a compromise between all semantic aspects of the pose can prove challenging when limb proportions are very different.
- **Unrealistic motion trajectories:** preventing floor collisions can cause a modification of the trajectory of the character, without enforcing the laws of dynamics, might create motions that are not physically accurate, although this effect is only visible on some cartoon characters with extreme proportions.
- **Fine-grained contacts on undersampled parts of the mesh:** using few keypoints constrains us to a simplified representation of the mesh, which might fail to capture contacts on some parts of the mesh.

A detailed analysis of those failure cases can be found in the Supplementary Material, with examples of each type of failure case.

6. Extensions

Thanks to the modular nature of our method, the optimization terms can be easily extended to account for specific and challenging scenarios previously unaddressed in the literature, as well as providing some degree of authoring and interactive control.

6.1. Multi-character retargeting

Handling multi-character interaction can be achieved by extending the pose descriptors to account for all interactions between the characters. These interactions may encompass not only physical contact but also some long-distance interactions such as gaze direction. In particular, handling eye contact is key in modeling plausible human interaction.

To this end, we first expand $\mathcal{W}_{\text{interaction}}$ and $\mathcal{W}_{\text{height}}$ to account for all possible distances between pairs of key-vertices of all the characters. The respective dimensions of these matrices become, respectively, $NN_{\text{char}} \times NN_{\text{char}}$ and $NN_{\text{char}} \times 1$, where N_{char} designates the number of characters. Second, we propose a modified $\mathcal{W}_{\text{interaction}}$ descriptor formulation accounting for an additional gaze direction term. Let us call $\mathcal{I}^{\text{eye}} = (i_c^{\text{eye}})_{c \in [1, N_{\text{char}}]}$ the set of indices of the key-vertices associated with the eyes of a character c , and typically positioned at the center of the character's face. We further call gaze_c the averaged gaze direction of character c at a

given time, and define the following sparse matrix:

$$\mathcal{M}_{\text{gaze}} \in \mathbb{R}^{N_{\text{char}} N \times 3}, \text{ where } \mathcal{M}_{\text{gaze}}(i) = \begin{cases} \text{gaze}_c & \text{for } i \in \mathcal{I}^{\text{eye}} \\ 0 & \text{otherwise.} \end{cases}$$

We finally define the extended $\mathcal{W}_{\text{interaction}}$ as :

$$\mathcal{W}_{\text{interaction}} += \text{clamp} \left(\frac{S_{\text{cosine}}(\mathcal{M}_{\text{dir}}, \mathcal{M}_{\text{gaze}}) - \cos(a_{\min})}{\cos(a_{\min}) - \cos(a_{\max})} \right)$$

The thresholds $a_{\min}$ and $a_{\max}$ are characteristic angles, set to 2° and 5° . Figure 8 illustrates our results on multi-character interactions and shows the separate impacts of adding the inter-character physical contact and the gaze-based extension. Note how the optimization automatically changes the relative position of the characters, as well as their poses, to account for their height difference.

We also experiment with multi-character retargeting with strictly more than two characters. We design and retarget a five-character motion inspired from a traditional Breton dance, shown in Figure 9. We found that the multi-character retargeting only runs 19% slower than five independent retargeting processes, due to the sparse nature of the interactions in the method. The speed and memory footprints could likely be optimized further using specific sparse computations libraries.

6.2. Non-flat grounds

A second extension of our approach is to generalize the retargeting to non-flat grounds. This can be straightforwardly implemented by adapting the height evaluation of the key-vertices to become coordinates-dependent. Considering a terrain defined as a differentiable height-field $y = f(x, z)$ and key-vertex position coordinates as $p_i = (x_i, y_i, z_i)$, we can replace $p_i \leftarrow p_i - f(x_i, z_i)$ in the computation of $\mathcal{M}_{\text{height}}$ and $\mathcal{M}_{\text{sliding}}$ descriptors. Illustrations of results obtained using this approach are provided in Figure 10.

6.3. Identifying and solving conflicting constraints

An advantage of our formulation is its ability to automatically detect and identify conflicting semantic objectives during the optimization process. This identification can be coupled with an interactive adjustment of the weights assigned to the associated losses, thus offering an efficient authoring tool to handle complex retargeting scenarios. Let us consider the loss component as a scalar function $\mathcal{L}_k(\mathbf{q}_b)$, where k is the type of loss (among dist, dir, pen, height, and sliding), and b is a body part. The cosine similarity

$$S_{\text{cosine}}(\nabla \mathcal{L}_{k_1}(\mathbf{q}_b), \nabla \mathcal{L}_{k_2}(\mathbf{q}_b)) \quad (3)$$

is an indicator of the conflict between the loss k_1 and k_2 for the body part b . More precisely, a negative value close to -1 indicates strong conflicting goals between two losses. Figure 11 shows an example such use case where a conflict is identified between $\mathcal{L}_{\text{pen}}$ and $\mathcal{L}_{\text{dist}}$ on the torso, as the character is unable to reach its feet because of its large belly and short limbs. The user can then use an interactive cursor to adapt a weight between the two losses, so as to find the most pertinent output for their use case. In one case (Figure 11c), collision is neglected to account for more accurate contacts; in the other (Figure 11d), the output is free of collisions, but contacts are lost.

7. Conclusion and future work

We presented ReConForM, a novel motion retargeting method for characters of arbitrary shape and skeleton structure, focusing on self-collisions and ground contacts. Our solution efficiently encodes shape and motion by creating time-varying pose descriptors based on the trajectory of selected key vertices of the character's meshes. Thanks to a proximity-based criterion, we designed a way to select and weight the descriptors that carry the most importance in each frame. By optimizing the target poses to conform as well as possible to the source's motion descriptions, we are able to retarget motion in real time. Compared to state-of-the-art methods, our algorithm achieves more accurate contacts with the ground, smoother motion, and effectively avoids self-collisions.

Our method's simplicity and flexibility allow easy adaptation to new use cases. We presented three extensions: accounting for inter-character contacts and gaze alignment to achieve multi-character retargeting; adapting the height of the ground for retargeting motion on non-flat terrains; and identifying conflicting constraints, allowing for interactive feedback and authoring, as prioritizing one of the constraints is often more desirable than getting high errors for them all.

7.1. Limitations and future work

Our method, while achieving real-time results, has a few limitations, which we would like to address in future work. Firstly, it

does not account for how a character's morphology (e.g., weight) affects its motion style (e.g., inertia). Addressing this may request additional modifications beyond pure pose-based analysis, paving the way to include dynamic speed-based or acceleration-based descriptors. Secondly, while our method can handle different skeleton topologies and large geometrical variations from source to target, it cannot transfer motions between drastically different articulated structures that lack a common template, e.g. a target character with additional limbs. To handle this case, additional work would be required to extend the method to non-bijective mapping between key-vertices. In addition, our method would highly benefit from an adaptive set of key-vertices, generated in collision areas of the source motion, and automatically transferred to the target character. This could enable the use of an even sparser set, but of more precise of motion descriptors over time. Lastly, combining our semantic losses with deep learning methods could accelerate inference, while reinforcement learning methods could be used to ensure physically plausible motions. The latter may be particularly useful when retargeting to non-flat grounds, where equilib-




Figure 8: Results of the multi-character extension on a high-five pose. Arrows denote the gaze direction.




Figure 9: Results of the multi-character extension on a five-character motion.




Figure 10: Results of the non-flat terrain extension.




Figure 11: Example of conflicting losses and conflict resolution

rium might be lost. Finally, since assessing the quality of retargeting strongly benefits from human feedback, human-in-the-loop reinforcement learning might give new insights on how to move beyond manually-designed motion descriptors.

References

- [Ado24] ADOBE. *Mixamo*. <https://www.mixamo.com/>. Accessed: 2023-09-12. 2024 **3**, **4**, **7**, **8**.
- [ALL*20] ABERMAN, KFIR, LI, PEIZHUO, LISCHINSKI, DANI, et al. “Skeleton-Aware Networks for Deep Motion Retargeting”. *ACM Transactions on Graphics (TOG), Proc. SIGGRAPH* 39.4 (2020), 62 **2**, **3**, **7**.
- [Aut25] AUTODESK, INC. *Maya (MotionBuilder)*. Version 2025.2. 2025. URL: <https://www.autodesk.com/fr/support/technical/product/motionbuilder> **7**.
- [AWL*19] ABERMAN, KFIR, WU, RUNDI, LISCHINSKI, DANI, et al. “Learning character-agnostic motion for motion retargeting in 2D”. *ACM Transactions on Graphics* 38.4 (July 2019), 1–14. ISSN: 1557-7368. DOI: [10.1145/3306346.3322999](https://doi.org/10.1145/3306346.3322999). URL: <http://dx.doi.org/10.1145/3306346.3322999> **2**.
- [BOH*22] BASSET, JEAN, OUANNAS, BADR, HOYET, LUDOVIC, et al. “Impact of Self-Contacts on Perceived Pose Equivalences”. *Motion, Interaction and Games (MIG)*. 2022. ISBN: 9781450398886. DOI: [10.1145/3561975.3562946](https://doi.org/10.1145/3561975.3562946). URL: <https://doi.org/10.1145/3561975.3562946> **2**.
- [BWB*20] BASSET, JEAN, WUHRER, STEFANIE, BOYER, EDMOND, and MULTON, FRANCK. “Contact preserving shape transfer: Retargeting motion from one shape to another”. *Computers & Graphics* 89 (2020), 11–23. ISSN: 0097-8493. DOI: <https://doi.org/10.1016/j.cag.2020.04.002>. URL: <https://www.sciencedirect.com/science/article/pii/S00978493203004062>.
- [CK00] CHOI, KWANG-JIN and KO, HYEONG-SEOK. “Online motion retargeting”. *The Journal of Visualization and Computer Animation* 11.5 (2000), 223–235 **2**.
- [CRV12] CASIEZ, GÉRY, ROUSSEL, NICOLAS, and VOGEL, DANIEL. “1€ filter: a simple speed-based low-pass filter for noisy input in interactive systems”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2012, 2527–2530 **7**.
- [Cut13] CUTURI, MARCO. “Sinkhorn Distances: Lightspeed Computation of Optimal Transport”. *Advances in Neural Information Processing Systems*. Ed. by BURGESS, C.J., BOTTU, L., WELLING, M., et al. Vol. 26. Curran Associates, Inc., 2013. URL: https://proceedings.neurips.cc/paper_files/paper/2013/file/af21d0c97db2e27e13572cbf59eb343d-Paper.pdf **4**.
- [Epi20] EPIC GAMES. *Unreal Engine*. Version 4.22.1. 2020. URL: <https://www.unrealengine.com> **7**.
- [FSV*19] FEYDY, JEAN, SÉJOURNÉ, THIBAUT, VIALARD, FRANÇOIS-XAVIER, et al. “Interpolating between Optimal Transport and MMD using Sinkhorn Divergences”. *The 22nd International Conference on Artificial Intelligence and Statistics*. 2019, 2681–2690 **4**.
- [Gle98] GLEICHER, MICHAEL. “Retargeting motion to new characters”. *Proc. ACM SIGGRAPH*. 1998, 33–42 **2**.
- [HKT10] HO, EDMOND, KOMURA, TAKU, and TAI, CHIEW-LAN. “Spatial Relationship Preserving Character Motion Adaptation”. *ACM Trans. Graph.* 29 (July 2010). DOI: [10.1145/1833351.1778770](https://doi.org/10.1145/1833351.1778770) **2**, **3**.
- [Hvy*20] HANDA, ANKUR, VAN WYK, KARL, YANG, WEI, et al. “Dex-Pilot: Vision-Based Teleoperation of Dexterous Robotic Hand-Arm System”. *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 2020, 9164–9170. DOI: [10.1109/ICRA40945.2020.9197124](https://doi.org/10.1109/ICRA40945.2020.9197124) **3**.
- [HZZ*23] HU, LEI, ZHANG, ZHAO, ZHONG, CHONGYANG, et al. “Pose-Aware Attention Network for Flexible Motion Retargeting by Body Part”. *IEEE Transactions on Visualization and Computer Graphics* (2023), 1–17. ISSN: 2160-9306. DOI: [10.1109/tvcg.2023.3277918](https://doi.org/10.1109/tvcg.2023.3277918). URL: <http://dx.doi.org/10.1109/TVCG.2023.3277918> **2**.
- [JKL17] JIN, TAEIL, KIM, MEEKYUNG, and LEE, SUNG-HEE. “Motion Retargeting to Preserve Spatial Relationship between Skinned Characters”. *Symposium on Computer Animation (SCA)*. 2017. ISBN: 9781450350914. DOI: [10.1145/3099564.3106647](https://doi.org/10.1145/3099564.3106647). URL: <https://doi.org/10.1145/3099564.3106647> **2**.
- [KB14] KINGMA, DIEDERIK P. and BA, JIMMY. “Adam: A Method for Stochastic Optimization”. *CoRR* abs/1412.6980 (2014). URL: <https://api.semanticscholar.org/CorpusID:66281067>.
- [KNK*02] KUFFNER, J., NISHIWAKI, K., KAGAMI, S., et al. “Self-collision detection and prevention for humanoid robots”. *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*. Vol. 3. 2002, 2265–2270 vol.3. DOI: [10.1109/ROBOT.2002.1013569](https://doi.org/10.1109/ROBOT.2002.1013569) **3**.
- [LCC19] LIM, JONGIN, CHANG, HYUNG JIN, and CHOI, JIN YOUNG. “PMnet: Learning of Disentangled Pose and Movement for Unsupervised Motion Retargeting”. *British Machine Vision Conference (BMVC)*. 2019 **2**, **3**.
- [LMR*15] LOPER, MATTHEW, MAHMOOD, NAUREEN, ROMERO, JAVIER, et al. “SMPL: A Skinned Multi-Person Linear Model”. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 34.6 (2015), 248:1–248:16 **2**, **4**.
- [LS99] LEE, JEHEE and SHIN, SUNG YONG. “A hierarchical approach to interactive motion editing for human-like figures”. *Proc. ACM SIGGRAPH*. 1999, 39–48 **2**.
- [LYS*22] LIAO, ZHOUYINGCHENG, YANG, JIMEI, SAITO, JUN, et al. “Skeleton-free Pose Transfer for Stylized 3D Characters”. *European Conference on Computer Vision (ECCV)*. Springer. Oct. 2022 **3**.
- [MCK*17] MANDAD, MANISH, COHEN-STEINER, DAVID, KOBELT, LEIF, et al. “Variance-Minimizing Transport Plans for Inter-surface Mapping”. *ACM Transactions on Graphics, Proc. SIGGRAPH* 36 (2017), 14. DOI: [10.1145/3072959.3073671](https://doi.org/10.1145/3072959.3073671). URL: <https://inria.hal.science/hal-01519006> **3**.
- [RWY*23] REDA, DANIELE, WON, JUNG DAM, YE, YUTING, et al. “Physics-based Motion Retargeting from Sparse Inputs”. *SCA, Proceedings of the ACM on Computer Graphics and Interactive Techniques* 6 (Aug. 2023), 1–19. DOI: [10.1145/3606928](https://doi.org/10.1145/3606928) **3**.




Figure 12: Results of retargeting on characters from Blender Studio. Source poses on the left, followed by the results on several characters (from left to right : Phil, Sprite, Rex, Rain, Gabby)

- [Sah20] SAHILIOĞLU, YUSUF. “Recent Advances in Shape Correspondence”. *Vis. Comput.* 36.8 (Aug. 2020), 1705–1721. ISSN: 0178-2789. DOI: [10.1007/s00371-019-01760-0](https://doi.org/10.1007/s00371-019-01760-0). URL: <https://doi.org/10.1007/s00371-019-01760-0>.
- [SGHR05] SHON, AARON, GROCHOW, KEITH, HERTZMANN, AARON, and RAO, RAJESH P. “Learning shared latent structure for image synthesis and robotic imitation”. *Advances in neural information processing systems* 18 (2005) 2.
- [SPK23] SCHMIDT, PATRICK, PIEPER, DÖRTE, and KOBELT, LEIF. “Surface Maps via Adaptive Triangulations”. *Computer Graphics Forum* 42.2 (2023) 3.
- [SPKS16] SOLOMON, JUSTIN, PEYRÉ, GABRIEL, KIM, VLADIMIR G., and SRA, SUVRIT. “Entropic Metric Alignment for Correspondence Problems”. *ACM Transactions on Graphics* 35.4 (July 2016). ISSN: 0730-0301. DOI: [10.1145/2897824.2925903](https://doi.org/10.1145/2897824.2925903). URL: <https://doi.org/10.1145/2897824.2925903>.
- [TK05] TAK, SEYOON and KO, HYEONG-SEOK. “A physically-based motion retargeting filter”. *ACM Transactions on Graphics* 24.1 (2005), 98–117 2.
- [VCH*21] VILLEGAS, RUBEN, CEYLAN, DUYGU, HERTZMANN, AARON, et al. “Contact-aware retargeting of skinned motion”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, 9720–9729 2.
- [VYCL18] VILLEGAS, RUBEN, YANG, JIMEI, CEYLAN, DUYGU, and LEE, HONGLAK. “Neural Kinematic Networks for Unsupervised Motion Retargeting”. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2018 2, 3, 7.
- [VZHC11] VAN KAICK, OLIVER, ZHANG, HAO, HAMARNEH, GHASAN, and COHEN-OR, DANIEL. “A survey on shape correspondence”. *Computer graphics forum*. Vol. 30. 6. Wiley Online Library. 2011, 1681–1707 3.
- [WWF*20] WANG, JIASHUN, WEN, CHAO, FU, YANWEI, et al. “Neural Pose Transfer by Spatially Adaptive Instance Normalization”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, 5831–5839 3.
- [YML23] YAN, YASHUAI, MASCARO, ESTEVE VALLS, and LEE, DONGHEUI. “ImitationNet: Unsupervised Human-to-Robot Motion Retargeting via Shared Latent Space”. *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*. 2023, 1–8. DOI: [10.1109/Humanoids57100.2023.10375150](https://doi.org/10.1109/Humanoids57100.2023.10375150) 2.
- [ZBP20] ZHOU, KEYANG, BHATNAGAR, BHARAT LAL, and PONS-MOLL, GERARD. “Unsupervised Shape and Pose Disentanglement for 3D Meshes”. *European Conference on Computer Vision (ECCV)*. Aug. 2020 3.
- [ZCX*24] ZHANG, HAODONG, CHEN, ZHIKE, XU, HAOCHENG, et al. “Semantics-aware Motion Retargeting with Vision-Language Models”. (2024) 2.



Figure 13: Results of our method on several animations taken from our validation dataset. Animation references from top to bottom : Phone Call, Shoulder Rubbing, Pike Walk, Angry, Gangnam Style. From left to right : Source, Copy Rotations, R^2ET [ZWK*23], Ours.

- [ZGY*23] ZHANG, YUNBO, GOPINATH, DEEPAK, YE, YUTING, et al. “Simulation and Retargeting of Complex Multi-Character Interactions”. *ACM SIGGRAPH 2023 Conference Proceedings*. 2023. ISBN: 9798400701597. DOI: [10 . 1145 / 3588432 . 3591491](https://doi.org/10.1145/3588432.3591491). URL: <https://doi.org/10.1145/3588432.3591491> 3.
- [ZWK*23] ZHANG, JIAXU, WENG, JUNWU, KANG, DI, et al. “Skinned Motion Retargeting With Residual Perception of Motion Semantics & Geometry”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2023, 13864–13872 [2](#), [3](#), [7](#), [8](#), [13](#).