A Generative Motion Rig for Artist-Driven Motion Authoring

Jakob Buhmann
DisneyResearch|Studios
Zürich, Switzerland
jakob.buhmann@disneyresearch.com

Dhruv Agrawal
ETH Zürich
Zürich, Switzerland
DisneyResearch|Studios
Zürich, Switzerland
dhruv.agrawal@inf.ethz.ch

Dominik Borer
DisneyResearch|Studios
Zürich, Switzerland
dominik.borer@disneyresearch.com

Luca Vögeli
DisneyResearch|Studios
Zürich, Switzerland
voegeli.luca@gmail.com

Robert W. Sumner
DisneyResearch|Studios
Zürich, Switzerland
ETH Zürich
Zürich, Switzerland
sumner@disneyresearch.com

Martin Guay
DisneyResearch|Studios
Zürich, Switzerland
martin.guay@disneyresearch.com

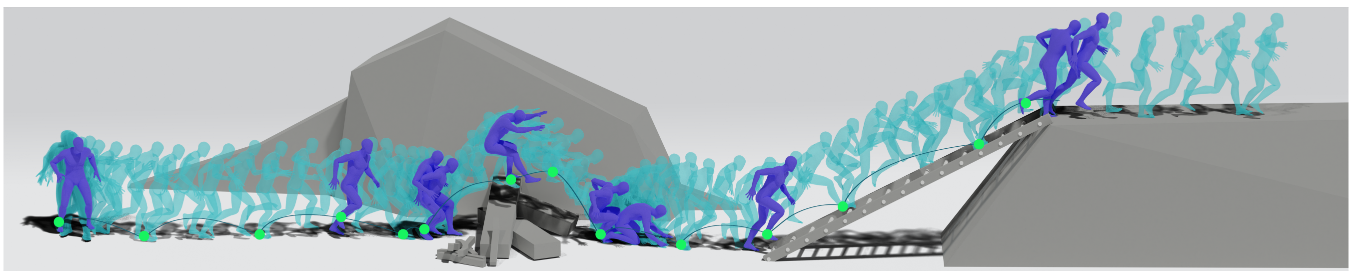

Figure 1: An animation sequence (cyan) created with our Generative Motion Rig. Artists can use sparse keyframes (blue), click-and-drag on neural motion curves (green), and sample different motions, enabling a new generative motion authoring paradigm.

Abstract

The recent success of generative modeling has led to entirely new capabilities to author 3D motion. By manipulating only a few sparse handles and poses, it is now possible to generate entire motion sequences, at scale and with many variations. To bridge the gap between model research and real-world integration, we developed a Generative Motion Rig as a Blender plugin, built atop a general motion model. Our rig supports a new "generative keyframing" workflow, where artists author movements by manipulating sparse poses, handles, window lengths, and noise sampling. We show in our accompanying video how artists use these controls to rapidly make a short animation. We also show how the same rig can support *generative* motion editing, allowing one to edit and extend mocap clips. Finally, we share insights and future challenges to help close the gap between generative and traditional animation workflows.

CCS Concepts

• **Computing methodologies** → *Artificial intelligence*; **Animation**.

ACM ISBN 979-8-4007-2541-8/2026/07
https://doi.org/10.1145/3799818.3812088

Keywords

Generative Motion Models, Control Rig, Character Animation

ACM Reference Format:
Jakob Buhmann, Dhruv Agrawal, Dominik Borer, Luca Vögeli, Robert W. Sumner, and Martin Guay. 2026. A Generative Motion Rig for Artist-Driven Motion Authoring. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH Talks '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3799818.3812088

1 Generative Motion Rig

In this work, we describe the design of our Generative Motion Rig (GMR), including the new workflows it enables, artist feedback, and insights for bringing this novel authoring paradigm into animation software.

Our GMR parameterizes motion with a single generative model and provides familiar handles via a posing rig (traditional or learned [Oreshkin et al. 2022]), together with sparse handles onto Neural Motion Curves (NMCs) [Agrawal et al. 2024]. This enables a new authoring experience, where complex motion can be sampled from a few constraints and manipulated by intuitive click-and-drag operations. The user additionally controls the generation length and can sample motion variations.

Implementation: An overview of our system can be found in Fig. 2. To integrate GMR into software such as Maya or Blender,

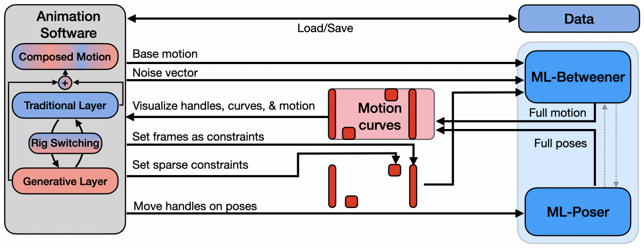

Figure 2: Overview of the client (grey) server (blue) setup. The UI (pink) reflects the constraints and the generated motion that is visualized for the user and that the user can influence.

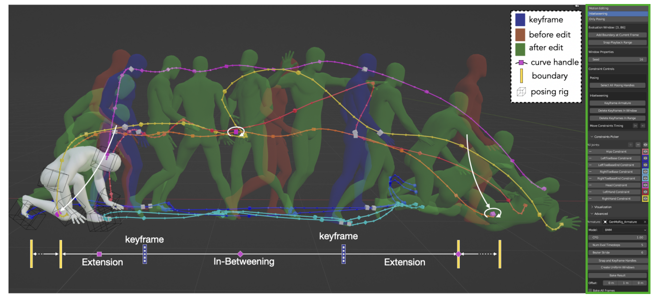

Figure 3: Visualization of our GMR in Blender with NMCs, posing handles (cubes), keyframes (blue), and sparse constraints (pink). This sequence depicts motion extension and the effect of editing with only three sparse constraints (before in red, after in green). The user-specified evaluation boundaries are indicated in yellow on the timeline.

we use a client-server paradigm. This decouples the models running on a dedicated GPU server from the host software with its often restrictive development environment, enabling flexible client swapping to other animation platforms. When the user modifies data influencing motion, such as keyframes, the changes are sent to the server and the updated motion is returned to the authoring software.

On the server, we use ML-Poser, a neural IK-model [Oreshkin et al. 2022], to complete full-body poses from sparse joint constraints. For the motion model, the ML-Betweener module uses the Implicit Bézier Motion Model (IBMM) [Vögeli et al. 2025], which leverages a Bézier representation that ensures temporal sparsity and smoothness, aligning with traditional animation workflows. While we showcase results with IBMM, our framework is compatible with other generative motion engines [Cohan et al. 2024].

In the authoring software, we expose posing handles via an armature and the NMCs for visualization and direct manipulation of the generated motion. Each control point on the curves is a keyable object editable in the 3D viewport or via the graph editor. To manage visual density, users can toggle curve visibility, customize coloring, and display time indices for each control handle. In addition, the system provides access to relevant data including keyframes, spatial constraints, temporal evaluation boundaries, and the initial noise at each timestep. For storing the motion, we opted for a layered approach such that poses and motion can be transferred between the "generative layer" and the "traditional layer," or mixed together for motion editing. Lastly, the layers offer the user the flexibility of blending or switching between them to compose the final motion.

2 Generative Authoring Capabilities

In the following, we highlight some of the capabilities that our GMR offers for authoring motion.

Direct Control. Combining IBMM and NMC enables directly manipulating the generated motion (Fig. 3). A user can provide a combination of sparse joint-wise constraints or full keyframes for conditioning the generated motion. By translating sparse constraints on the *feet*, the foot pattern can be easily changed. Or, by rotating a single sparse *hip* constraint, one can change walking motion into complex spinning or falling dynamics.

Noise Manipulation. GMR allows the user to resample the seed noise, either for the full motion interval or more locally for a single timestep. Additionally, resampling the positional part of the motion vector leads to different behavior than resampling the orientational part. While the former can change the overall outline drastically, the latter maintains the shape and changes the pose at that timestep.

Temporal Control. Animators can slide constraints—both sparse handles and full-body keyframes—along the timeline to exert continuous control over motion dynamics. Moving a constraint earlier in the timeline forces the model to hit that target sooner, effectively speeding up the preceding motion. Thus, a walk can be easily transformed into a sprint.

Motion Extension. Besides in-betweening, GMR also supports motion extension, where the model predicts motion outside the bounds of the given constraints. As shown in our video, this can be used to continuously extend a clip forward in time by generating new motion and shaping it with sparse constraint handles, while maintaining freedom over the final pose.

Motion Editing. In addition to generating new motion, GMR also enables editing existing animations by optionally sending a base motion to be preserved by the ML-Betweener on the server, as shown in Fig. 2. On the server side, the diffusion process is guided towards the base motion using a custom inpainting strategy. GMR gives users precise control over which parts of the motion to maintain, and where to generate new motion conforming to the given constraints. This generative motion editing concept can be used for stitching, extending, and directly manipulating mocap—providing a natural bridge to traditional animation techniques.

Animation Layering and Rig Switching. For seamless integration with professional workflows, artists must be able to concurrently utilize traditional rigs alongside our GMR within a single animation sequence. This hybrid approach allows artists to iterate on the timing and the layout using the GMR, subsequently transitioning to more precise traditional rigs for further polish, such as to stylize the motion and pushing it beyond physical realism. To facilitate this, we build a rig-switching logic between the FK-controls of a traditional rig and our GMR, as shown in Fig. 2. By building upon existing switches between IK- and FK- controls, our system allows the user to easily transfer poses between all three control types.

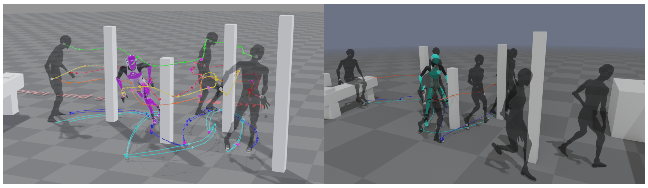

(a) Non-artist **(b) Artist**

Figure 4: Final constraints added by the two participants for animating a zig-zag motion. The black silhouettes show full body keyframes and the pink cubes show sparse constraints.

When switching from a traditional rig to GMR, we add the current deform pose as a full body constraint. Hence, the motion generated by GMR conforms to the specified pose from the traditional rig. Additionally, the user can use the *synchronize* mode where manipulating the traditional rig immediately adds a full body constraint to the GMR. Thus, artists can pose with traditional rig armature but view the updated generative motion on GMR's armature for the current frame and on NMCs for the remaining frames.

Note that our rig generates a whole motion sequence as opposed to the sparse per-frame information of traditional rigs. Therefore, when switching from GMR to the traditional rig, transferring the whole motion would add constraints on every frame, rendering the traditional rig uneditable. Instead, we choose to transfer only the current pose, as this allows the artist to pick out the key poses that they want to transfer to their traditional rig for further polishing.

This choice also highlights a limitation of current motion generation models, where regenerating with a previously predicted pose as a condition causes the generated motion to change. This can cause a *pop* in the generated motion in *synchronize* mode, as the transferred poses are added as new full-body keyframes.

Lastly, as GMR is an additional rig alongside a traditional rig, the user can use standard animation tools to blend between each rig's influence across different parts of the sequence.

3 Workflows

This section details the different workflows discovered during various tests and their compatibility with traditional workflows.

Freestyle Test. For this test, we tasked a professional artist to create any animation of their choice using our GMR. As shown in Fig. 1 and the accompanying video, the artist created a long chase between two characters. For this example, the artist combined keyframes, sparse manipulation of the NMCs, and sampling motion and key poses from it in a "generative keyframing" workflow. Compared to cubic interpolation, our motion model results in highly non-linear and natural movement. Note that this complex 22 second clip was authored in less than two days, including the time required to learn the tool and storyboard the shot. The artist was particularly happy with how fast it was to author the full motion, especially in dynamic parts of the scene. However, they also noticed challenges with polishing and finalizing. Additionally, the concept of sampling different movements while having the window length influence the motion is a novel and unfamiliar concept.

Guided Test. In another test, we laid out a parkour scene with obstacles and animation prompts. The test was conducted with two users, one professional artist and one of the authors with no animation experience. Both participants were given 1.5 hours to complete the course with a 45 seconds animation.

While given identical instructions, we noticed varying workflows between the two participants. The author, a non-artist, preferred to visualize all NMCs and used a mix of full-body and sparse point constraints to lay out the motion. In contrast, the professional artist used ML-poser to lay out full-body keyframes and only used sparse constraints for later cleanup. The final constraints added by both participants are visualized in Fig. 4.

Given the flexibility of our rig, the two participants were able to build their own workflows and successfully complete the task. We include the final animations in our supporting video. While there are differences in the resulting motions, we attribute them to the large gap in the participants' experience and not their choice of workflow. It is worth noting that the non-artist would not have been able to create a similar animation without our generative rig.

Motion Editing Test. We demonstrate the compatibility of GMR with off-the-shelf motion capture systems such as Synth2Track [Buhmann et al. 2025] by editing a captured jumping sequence using the motion editing functionality of our GMR. As shown in our video, we can edit this temporally dense motion using the same controls as for generating new motion. The final animation changes the distance of the jump and adds a backward step.

4 Conclusion and Future Work

While we showcase a new system for generative motion authoring, there are still many opportunities for future exploration. First, generative models are inherently bound by their training data, which can limit artists when trying to achieve "out-of-distribution" movements, such as non-physical dynamics and more stylized animations. While a layered approach is a first step that allows generative and traditional animation to coexist in the same software, finding the optimal way to blend these disparate sources without losing physical or stylistic consistency remains an exciting area for future exploration. Additionally, current state-of-the-art generative models require improved stability when adding new constraints for a more seamless user experience. Lastly, scaling to more complex rigs and different character morphologies remains an open challenge.

References

Dhruv Agrawal, Jakob Buhmann, Dominik Borer, Robert W Sumner, and Martin Guay. 2024. SKEL-Betweener: a Neural Motion Rig for Interactive Motion Authoring. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–11.

Jakob Buhmann, Douglas L. Moore, Dominik Borer, and Martin Guay. 2025. Synth2Track Editor for Efficient Match-Animation. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH Talks '25)*. Association for Computing Machinery, New York, NY, USA.

Setareh Cohan, Guy Tevet, Daniele Reda, Xue Bin Peng, and Michiel van de Panne. 2024. Flexible Motion In-betweening with Diffusion Models. *arXiv preprint arXiv:2405.11126* (2024).

Boris N. Oreshkin, Florent Bocquelet, Felix G. Harvey, Bay Raitt, and Dominic Laflamme. 2022. ProtoRes: Proto-Residual Network for Pose Authoring via Learned Inverse Kinematics. In *International Conference on Learning Representations*.

Luca Vögeli, Dhruv Agrawal, Martin Guay, Dominik Borer, Robert W Sumner, and Jakob Buhmann. 2025. Implicit Bézier Motion Model for Precise Spatial and Temporal Control. In *Proceedings of the 2025 18th ACM SIGGRAPH Conference on Motion, Interaction, and Games*. 1–10.