Artist-in-the-Loop AI Animation: Real-Time Control and Synthesis with AnimHost

Jonas Trottnow
jonas.trottnow@filmakademie.de
Filmakademie Baden-Württemberg, Germany

Simon Haag
simon.haag@protonmail.com
Filmakademie Baden-Württemberg, Germany

Francesco Andreussi
francesco.andreussi@gmail.com
Filmakademie Baden-Württemberg, Germany

Simon Spielmann
simon.spielmann@filmakademie.de
Filmakademie Baden-Württemberg, Germany

Volker Helzle
volker.helzle@filmakademie.de
Filmakademie Baden-Württemberg, Germany

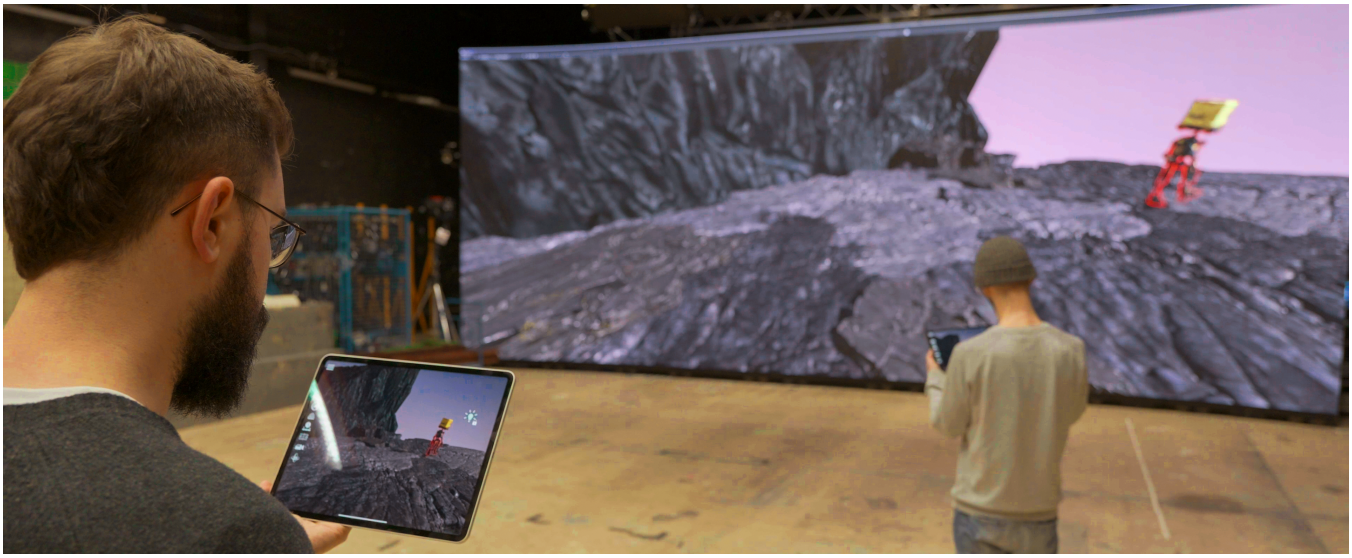

Figure 1: 'Survivor' testproduction for AnimHost - PreVis in LED Studio

ABSTRACT

AnimHost is an open-source, artist-focused application for real-time character animation driven by AI-generated motion data. It enables intuitive control of movement trajectories within familiar DCC (Digital Content Creation) tools such as Blender, supporting rapid iteration and seamless integration into standard production pipelines. Built on the TRACER[1] ecosystem, AnimHost decouples animation generation from host applications and introduces a node-based compute graph for preprocessing, inference, and retargeting. By allowing artists to train and deploy custom motion models, it maintains creative authorship and transparency over training data. AnimHost bridges the gap between research and production, offering a scalable and extensible path for integrating generative AI into real-world animation workflows.

CCS CONCEPTS

• **Applied computing** → **Media arts**; • **Computer systems organization** → **Real-time system architecture**; • **Networks** → **Application layer protocols**.

KEYWORDS

motion synthesis, character animation, artificial intelligence

ACM Reference Format:
Jonas Trottnow, Simon Haag, Francesco Andreussi, Simon Spielmann, and Volker Helzle. 2025. Artist-in-the-Loop AI Animation: Real-Time Control and Synthesis with AnimHost. In *SIGGRAPH Asia 2025 Technical Communications (SA Technical Communications '25), December 15–18, 2025, Hong Kong, Hong Kong*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3757376.3771385

[1] https://animationsinstitut.de/en/research/tools/tracer

SA Technical Communications '25, December 15–18, 2025, Hong Kong, Hong Kong

ACM ISBN 979-8-4007-2136-6/2025/12
https://doi.org/10.1145/3757376.3771385

1 INTRODUCTION

AI-assisted character animation has made notable advances, but practical deployment in real-world pipelines remains limited due to the lack of tools balancing automation with creative control. AnimHost addresses this gap. It is a real-time capable, open-source animation synthesis system designed with 3D animation production workflows in mind. Compatible with major DCC (Digital Content Creation) tools such as Blender and integrated into the TRACER ecosystem, AnimHost enables artists to prototype, direct, and refine motion using neural motion generators. Crucially, it prioritizes transparency and authorial control by supporting custom datasets, local model training, and modular, artist-driven workflows.

2 SYSTEM OVERVIEW

AnimHost was developed within the EU-funded MAX-R project[2] as part of the TRACER ecosystem - an open-source infrastructure for XR (extended reality) media production. TRACER supports real-time, multi-user workflows through its DataHub, which synchronizes scene and animation data across DCC tools like Blender, Unity and Unreal, as well as with a tablet-based XR production tool for simplified scene access and editing called VPET [Spielmann et al. 2018], developed by the same team. AnimHost builds on this foundation with a standalone, node-based desktop application for motion synthesis and interactive animation layout.

The architecture of the system (Figure 2) includes motion generation, data transformation, streaming, and post–processing modules, each implemented as graph nodes based on plugins. The application is cross-platform and engine-agnostic, operating independently of any rendering or animation host application. It provides bidirectional communication channels to DCCs, allowing seamless asset, motion, and annotation exchange. AnimHost was designed from the ground up with modularity and extensibility in mind. Every node in the system adheres to a clear separation of execution and data flow, allowing for live debugging, introspection, and network-wide performance tracing. TRACER has been applied and tested in diverse production scenarios, e.g. in the location-based XR Experience 'Fate of the Minotaur' [Dahn et al. 2024].

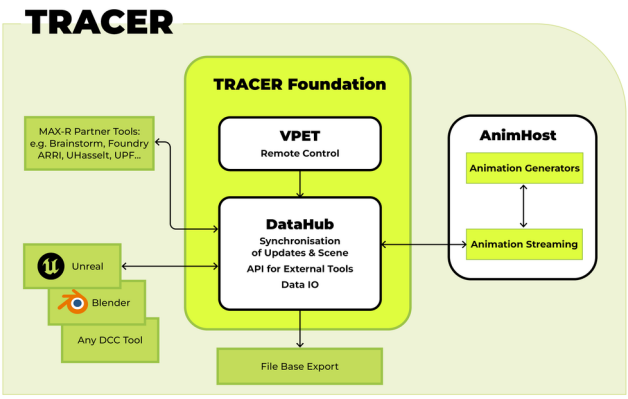

Figure 2: TRACER infrastructure overview showing data and motion stream connections across DCCs

[2]https://max-r.eu

3 NEURAL MOTION SYNTHESIS PIPELINE

3.1 Compute Graph and Node API

AnimHost adopts a compute-graph paradigm where each node encapsulates a computational function. These include:

- Skeleton preprocessing / retargeting between character rigs
- Trajectory prediction and spline-to-feature conversion
- ONNX[3] inference for autoregressive motion generation
- Post-processing filters (smoothing, alignment)
- Real-time network transmission modules

The node system is scriptable and visual (Figure 3), allowing technical directors and skilled artists to build custom data pipelines that match their animation style, solver characteristics, and production needs. A typical graph involves loading a 3D character rig, binding a skeleton, computing a target trajectory from user input, running the model and exporting results to the streaming layer.

3.2 Training from Custom Motion Capture Data

To ensure traceable data sources for training data and ethical use of AI, AnimHost provides preprocessing and training routines based on user-supplied datasets. We recorded more than 57,000 frames of motion using an optical marker-based system [Animationsinstitut 2025a], following the "dance card" methodology [Zadziuk 2016]. Animations were segmented into pose sequences, joint velocities, and phase descriptors. Using DeepPhase [Starke et al. 2022], we trained a periodic autoencoder and an autoregressive prediction network. Models were exported to ONNX.

Recent advances in motion generation span from PFNN [Holden and Komura 2017], Robust Motion In-betweening [Harvey et al. 2021], Real-time Controllable Motion Transition [Park et al. 2022], and RSMT [Kim et al. 2023], to diffusion-based frameworks [Lee et al. 2024]. As an open platform, AnimHost can integrate such methods as plugins, making their deployment into production pipelines straightforward. AnimHost itself focuses on artist-centric integration, real-time bidirectional DCC coupling, and open-source reproducibility within the TRACER ecosystem.

Training is performed externally using PyTorch and the reference implementation from Starke et al [Starke et al. 2022]. The resulting network accepts spline-based trajectories as input and produces full-body locomotion at 30fps. A preprocessing graph embedded in AnimHost handles phase calculation and motion cleanup, including end-effector locking and joint trajectory regularization. A training monitor visualizes loss curves, phase stability and frame-wise reconstruction error to assist users in evaluating model quality.

3.3 Live Generation and User Controls

Users interact with the animation generator by defining input splines in Blender or VPET. The spline is converted into a trajectory tensor and fed to the generator node. Resulting motion is streamed in real-time to Blender, where it is applied to a retargeted skeleton. Users can adjust mix-weights, path-following aggressiveness, turning rates, stride length modulation, and ground projection strength, offering granular control over the AI output.

Latency for live animation generation is under 100ms. During production scenarios, full shot sequences can be generated and

[3]https://onnxruntime.ai

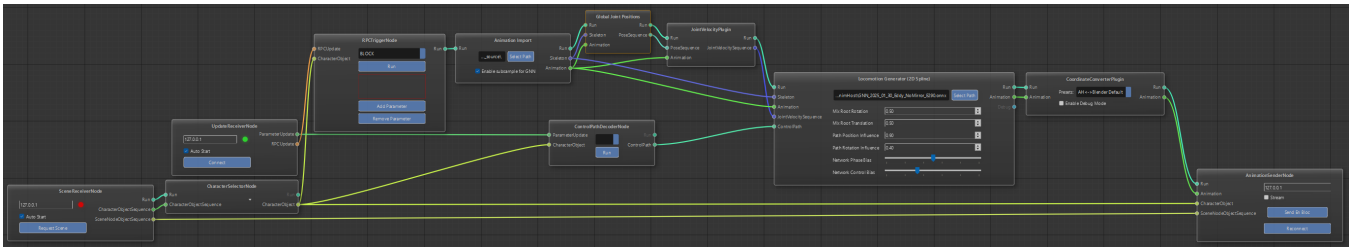

Figure 3: AnimHost provides a node graph interface for simplified user interaction with parameters and data-flow

cached, allowing AnimHost to function as a keyframe provider for offline refinement.

4 WORKFLOWS AND USE CASES

4.1 Blender Add-on and Rig Architecture

Blender serves as the primary interface for spline authoring and animation refinement. The custom add-on allows users to:

- Define walk paths using standard curve objects with adjustable resolution and timing markers
- Assign look-at targets, emotion tags and custom gait metadata to curve segments
- Bake generated motion onto an animation rig or control rig
- Use FK/IK blending and constraints for fine-tuning

The rig architecture includes a dual-layer setup: a motion rig that receives the AI-generated keyframes, and a control rig that overlays animator-defined modifications. During baking, constraints and drivers are activated to mirror bone transforms; afterward, artists can switch to manual control and override parts of the motion.

This setup allows for reuse of industry-standard tools including the Graph Editor, Dope Sheet, NLA Editor and Pose Libraries. Rig constraints (e.g. foot roll, inverse pole vector) continue to work seamlessly with AI data input, and can be layered using the standard animation stack.

4.2 VPET Integration for Exploration of On-Set Direction in Immersive Environments

VPET, TRACER's tablet-based XR interface, transforms character animation direction into an interactive spatial experience. Designed for LED studio volumes and XR stages, VPET allows directors or artists to:

- Define walk paths by pointing on target locations
- Play back motion in real-time as it updates in connected clients like Blender
- Adjust timing, and edit spline curvature with touch gestures

The interaction possibilities and features are continuously extended and currently in an early testing phase.

This workflow lowers the barrier for non-technical creative staff to engage in motion layout. Directors can block entire scenes without touching a mouse or setting a keyframe. Because VPET operates on the same TRACER DataHub as Blender and AnimHost, changes made in one interface appear instantly in all others. This supports rapid prototyping, collaborative iterations, and allows decisions about pacing, blocking and choreography to happen in a collaborative immersive environment.

Through the Blender integration this animation pipeline also supports versioning. Multiple takes can be saved, tagged and compared before a final shot is committed to animation refinement.

4.3 Improved Foot Contact Stability

Early evaluations revealed visible foot-sliding artifacts with the sample implementation of DeepPhase [Starke et al. 2022] inference, primarily due to small temporal inconsistencies in autoregressive velocity prediction. To mitigate these effects we plan to introduce a lightweight contact-stabilization layer that operates as a post-inference pass. During motion generation, foot joints shall be monitored for low-velocity phases. When the horizontal velocity of a foot joint falls below a configurable threshold, the joint position is temporarily locked in world space using a damped least-squares inverse-kinematics (IK) correction. To avoid temporal discontinuities, the lock is smoothly blended in and out using a window over 3–5 frames. This is a clear example of the benefits of the user configurable, modular architecture of AnimHost. It can just be added as plugin to AnimHost and added to the GUI node graph.

5 EVALUATION

AnimHost was evaluated in professional production environments at the Filmakademie Baden-Württemberg through the animated short film *Survivor* and the diploma project *KITSUNE*. Both productions provided realistic testbeds to assess usability, workflow impact, and creative potential of AI-assisted character animation.

5.1 Production Tests

Within the EU-funded MAX-R project, AnimHost and the TRACER infrastructure were integrated into Blender via the "TRACER for Blender" add-on. The system enabled real-time motion generation for previsualization and blocking, connecting directly to Unreal Engine for final rendering. 12 shots—approximately one third of the film (around 800 frames)—were produced using AnimHost-generated animation. The typical workflow consisted of spline-based path creation in Blender, immediate motion preview through AnimHost, baking to a control rig, and refinement before lighting and rendering in Unreal. This setup allowed for rapid experimentation with multiple takes and versions of each shot.

Animators reported a one-third reduction in blocking time and appreciated the ability to preview AI-generated motion, describing the workflow as "immediate, transparent, and creatively liberating."

While the results were primarily used for previsualization and layout, the generated animations were rated as sufficiently stable for use as a foundation for final performance refinement. Minor foot-sliding and jittering at higher movement speeds were identified as the main areas for improvement but could be fixed with of the shelf solutions within production tools like Blender.

In the 2025 diploma production *KITSUNE*, AnimHost was employed for fast blocking and locomotion prototyping within a hybrid film/game pipeline. Artists confirmed its effectiveness for quickly generating walk cycles and directional transitions, particularly for background and secondary characters.

Figure 4: Shot from 3D short film 'Survivor' used for evaluation of the proposed tool-set and workflow

5.2 Summary

Across both productions, AnimHost demonstrated that AI-assisted motion generation can be integrated seamlessly into professional animation workflows. It accelerated layout and previsualization stages while preserving artistic control, establishing itself as a reliable tool for creative iteration. The case studies also revealed valuable directions for future development: improved contact handling, a richer motion dataset and continued refinement of the user interface to support both technical and non-technical artists.

6 DISCUSSION AND FUTURE WORK

The evaluations confirmed that AnimHost effectively accelerates early animation phases while maintaining transparency and authorial control. They also exposed typical limitations of AI-based character animation in production, such as limited motion diversity, occasional foot-sliding, and the need for stylistic adaptability.

To address these findings, ongoing development focuses on three areas:

- a lightweight velocity-based inverse-kinematics correction module for improved ground contact stability
- expansion of training datasets with additional actions including climbing, jumping, and stylized gestures
- usability enhancements such as adaptive UI layouts and color-coded node grouping

Further research explores scene-aware navigation, multi-character coordination, and motion in-betweening from sparse pose constraints. Integration with VPET aims to enable multi-character and emotion-driven control within LED and XR stages, allowing directors to prototype and direct performances in real-time.

Beyond film production, AnimHost shows potential for performing arts and live events, where AI-driven characters can react interactively to performers or audience. Its open, modular architecture provides a long-term foundation bridging AI research and production practice, empowering artists to co-create motion in real-time while retaining full ownership of both data and creative intent.

7 CONCLUSION

AnimHost offers a comprehensive toolchain for AI-based character animation tailored to production needs. Its integration with DCC tools like Blender and VPET enables seamless transitions between automation and artistic control. By leveraging user-owned datasets and real-time workflows, it empowers artists to co-create with AI while preserving authorship and adaptability. The open-source nature and extensibility of AnimHost ensure continued innovation across academic, educational, and industrial contexts. Further information, models, and code are available via the public GitHub repository [Animationsinstitut 2025b] and the project showcase site [Animationsinstitut 2025a].

8 ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon Europe Research and Innovation Programme under Grant Agreement No 101070072. The authors thank Leszek Plichta, Eduard Schaefer, Alexandru Tufis-Schwartz and Thomas Krüger.

REFERENCES

Animationsinstitut. 2025a. *AnimHost Motion Capture Dataset 2024 & 2025.* Retrieved August 11, 2025 from https://animationsinstitut.de/en/research/projects/max-r#anchor3

Animationsinstitut. 2025b. *Github repository R&D department of Animationsinstitut, Filmakademie Baden-Wuerttemberg.* Retrieved August 11, 2025 from https://github.com/FilmakademieRnd

Andreas Dahn, Leszek Plichta, Simon Spielmann, Eduard Schäfer, and Justus Blönnigen. 2024. Fate of the Minotaur: A scalable location based VR experience. In *ACM SIGGRAPH 2024 Immersive Pavilion* (Denver, CO, USA) *(SIGGRAPH '24)*. Association for Computing Machinery, New York, NY, USA, Article 4, 2 pages. https://doi.org/10.1145/3641521.3664406

Felix Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. 2021. Robust Motion In-betweening. In *ACM Transactions on Graphics (SIGGRAPH 2021)*, Vol. 40. 76:1–76:12. https://doi.org/10.1145/3450626.3459804

Daniel Holden and Taku Komura. 2017. Phase-Functioned Neural Networks for Character Control. In *ACM Transactions on Graphics (SIGGRAPH 2017)*, Vol. 36. 42:1–42:13. https://doi.org/10.1145/3072959.3073663

Jungdam Kim, Hanbyul Park, and Jehee Lee. 2023. RSMT: Real-time Stylized Motion Transition for Characters. In *ACM Transactions on Graphics (SIGGRAPH 2023)*, Vol. 42. 91:1–91:13. https://doi.org/10.1145/3592122

Sangmin Lee, Youngjoon Kim, Youngjoon Yoon, and Jehee Lee. 2024. Taming Diffusion Probabilistic Models for Character Control. In *ACM Transactions on Graphics (SIGGRAPH 2024)*, Vol. 43. 135:1–135:15. https://doi.org/10.1145/3658160

Hanbyul Park, Jungdam Won, and Jehee Lee. 2022. Real-time Controllable Motion Transition for Characters. In *ACM Transactions on Graphics (SIGGRAPH 2022)*, Vol. 41. 104:1–104:12. https://doi.org/10.1145/3528223.3530091

Simon Spielmann, Volker Helzle, Andreas Schuster, Jonas Trottnow, Kai Götz, and Patricia Rohr. 2018. VPET: virtual production editing tools. In *ACM SIGGRAPH 2018 Emerging Technologies (SIGGRAPH '18)*. Association for Computing Machinery, Article 22, 2 pages. https://doi.org/10.1145/3214907.3233760

Sebastian Starke, Ian Mason, and Taku Komura. 2022. DeepPhase: periodic autoencoders for learning motion phase manifolds. *ACM Trans. Graph.* 41, 4, Article 136 (jul 2022), 13 pages. https://doi.org/10.1145/3528223.3530178

Kristjan Zadziuk. 2016. *Motion Matching, The Future of Games Animation... Today, GDC.* Retrieved August 11, 2025 from https://youtu.be/KSTn3ePDt50