Performing arts in a new era

AI-assisted creation toolbox v1

22 January 2024 - M15

Document identifier: PRMR-D5.5 -AI assisted creation toolbox v1
Version: 1
Author: Daniel Bisig (STO), Pablo Palacio (STO)
Dissemination status: PU

Funded by
the European Union

Grant Agreement n°	101061303
Project acronym	PREMIERE
Project title	Performing arts in a new era: AI and XR tools for better understanding, preservation, enjoyment and accessibility
Funding Scheme	HORIZON-CL2-2021-HERITAGE-01 (HORIZON Research and Innovation Actions)
Project Duration	01/10/2022 – 30/09/2025 (36 months)
Coordinator	Athena Research Center (ARC)
Associated Beneficiaries	• Stichting Amsterdamse Hogeschool voor de Kunsten (AHK) • Forum Danca - Associacao Cultural (FDA) • Tempesta Media SL (TMP) • Cyens - Centre of Excellence (CNS) • Kallitechniki Etaireia Argo (ARG) • Medidata.Net - Sistemas de Informacao para Autarquias SA (MED) • Fitei Festival Internacional Teatro Expresao Iberica Crl (FIT) • Instituto Stocos (STO) • Universite Jean Monnet Saint-Etienne (UJM) • Associacao dos Amigos do Coliseu Doporto (COL) • Stichting International Choreographic Arts Centre (ICK)

Project no. 101061303
PREMIERE

Performing arts in a new era: AI and XR tools for better understanding, preservation, enjoyment and accessibility

HORIZON-CL2-2021-HERITAGE-01

Start date of project: 01/10/2022

Duration: 36 months

History Chart				
Issue	**Date**	**Changed page(s)**	**Cause of change**	**Implemented by**
0.1	22/12/2023	-	Draft	Daniel Bisig (STO) Pablo Palacio (STO)
0.2	29/12/2023	ALL	Additions	Daniel Bisig (STO)
0.3	03/01/2024	ALL	Additions	Daniel Bisig (STO)
0.4	14/01/2024	ALL	Additions after review	Daniel Bisig (STO) Pablo Palacio (STO)
1.0	22/01/2024	ALL	Final changes	Daniel Bisig (STO)

Validation			
No.	**Action**	**Beneficiary**	**Date**
1	Prepared	Aggelos Gkiokas (ARC) Philippe Colantoni (UJM)	9/1/2024
2	Approved	Aggelos Gkiokas (ARC) Philippe Colantoni (UJM)	14/01/2024
3	Released	Aggelos Gkiokas (ARC) Philippe Colantoni (UJM)	22/01/2024

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

Table of Contents

Executive Summary

Deliverable 5.5 deals with the development and dissemination of the first version of an AI-assisted creation toolbox. This toolbox provides functionalities that cater to the interests and needs of creative communities in dance and theatre. The needs had been previously elucidated through questionnaires and discussions with representative members of these communities. The AI-Toolbox has also been disseminated and evaluated during several workshops in which members of these communities have participated. The current version of the AI-Toolbox offers a wide range of functionalities to capture and analy~~s~~ze movements or to generate synthetic movements, sounds, or images. These functionalities are implemented using both machine-learning based and explicit algorithms and integrated into modular tools. Not only the choice of algorithms and tools but also the method of their implementation is informed by the needs of the target communities. The central aspects of the implementation are as follows: All implementations employ programming languages that are popular among the respective communities. The implementation is modular and supports communication protocols that facilitate the exchange of data and media with other software tools. The algorithms and models are relatively simple, run in real-time on conventional consumer hardware, can easily be trained on small datasets and interacted with in live settings. The implementation supports users with different levels of technical expertise. Users with little technical expertise can run the tools as standalone applications. Users with some technical expertise but no programming skills can retrain the models and/or modify the configurations of the tools. Users with some programming skills can modify the functionality of the open-source tools or integrate the algorithms and models into their own software.

The future developments that will lead to the next version of the AI-Toolbox focus on the refinement of the already existing algorithms and tools and on the addition of new tools that are currently being adopted or developed by the PREMiERE partners. While the dissemination and evaluation of the AI-Toolbox through workshops will continue, its application for the creation of new performance works will play an increasingly important role.

Funded by
the European Union

Acronyms and abbreviations

Abbreviation	Description
AI	Artificial Intelligence
BVH	Biovision Hierarchy
ML	Machine Learning
OSC	Open Sound Control
Sclang	SuperCollider programming language
WAV	Waveform Audio File Format

1. Introduction

This deliverable describes the development of an AI-assisted creation toolbox that supports the realizsation of media rich and interactive performance works in dance and theatre. These developments encompass several objectives:

- To provide multi-modal datasets that encompass motion capture, audio, and video recordings of the movement vocabularies of the Premiere partners STO and ICK. These datasets can be used to train the ML-based components of the AI-Toolbox.
- To provide a large collection of ML-based and explicit algorithms that cover a broad range of analytic or generative applications including movement analysis, movement synthesis, sound synthesis, and image synthesis.
- To integrate these algorithms into open-source tools that run as standalone applications, that can be readily integrated into creative workflows, and that cater to artists who possess different technical skills levels.
- To disseminate and evaluate the AI-Toolbox among the targeted user groups through workshops and new performance creations.

The content of this deliverable is organised as follows. Section 2 summarises the main requirements concerning the functionality of the AI-Toolbox. These requirements have been assessed through discussions and questionnaires with the target user groups. Section 3 describes the content and scope of two datasets that contain rich multi-modal recordings of contemporary dance movement vocabularies, Section 4 provides an exhaustive overview of all the tools that have currently been implemented and made available through the AI-Toolbox. This overview is structured into several sub-sections, each of which covers a different category of tools. Section 5 provides a list of dissemination activities in which the AI-Toolbox plays a central role. Section 6 contains concluding remarks and an outlook for future developments and applications of the AI-Toolbox.

Funded by
the European Union

2. User requirements

The requirements of the functionalities of the AI-Toolbox have been assessed through discussions and questionnaires with the target user groups (artists and performers, cultural industries, students and researchers within the performing arts) (T2.1). This assessment has revealed the following interests and needs:

1. For controlling interactively generated content, users want to employ their hands or full bodies. Furthermore, interaction should not only take into account low level kinematic aspects of the users' movements but also higher-level qualitative movement principles.
2. Users are interested in employing ML for generating synthetic motions that can serve as source for choreographic experimentation and inspiration. These synthetic motions shouldn't resemble human-like motions.
3. Users are interested in employing ML-based motion synthesis to drive the behaviour of an artificial dancer that can serve as a dance partner for a human dancer. Also here, the users' main interest lies in synthetic motions that are non-human-like.
4. Users are interested in creating synthetic sounds and images in real-time both to create media-augmented performances and as multimodal feedback to increase their awareness of body movements. The synthetic media should not be realistic.

To meet these requirements, the following functionalities have been added to the AI-Toolbox.

1. The AI-Toolbox provides methods to obtain partial or full body pose information either through professional motion capture or by analysing camera images from a live video stream. The AI-Toolbox also provides algorithms to derive higher level qualitative movement principles from low-level kinematic data.
2. The AI-Toobox incorporates an autoencoder-based model for motion synthesis and provides means to navigate the model's latent space to manipulate the level of realism that the generated synthetic motions exhibit.
3. The AI-Toolbox incorporates a sequence-prediction model that generates future motions as predicted continuations from preceding motions. This model can be employed as basis for an artificial dance partner by modifying through interaction aspects of the preceding motions.
4. The AI-Toolbox provides several algorithms for sound and image synthesis that can generate output in real-time and can be interactively controlled. The sound synthesis algorithms can produce a large variety of sounds that vary in their level of realism. The image synthesis algorithms generate exclusively non-realistic images.

3. Datasets

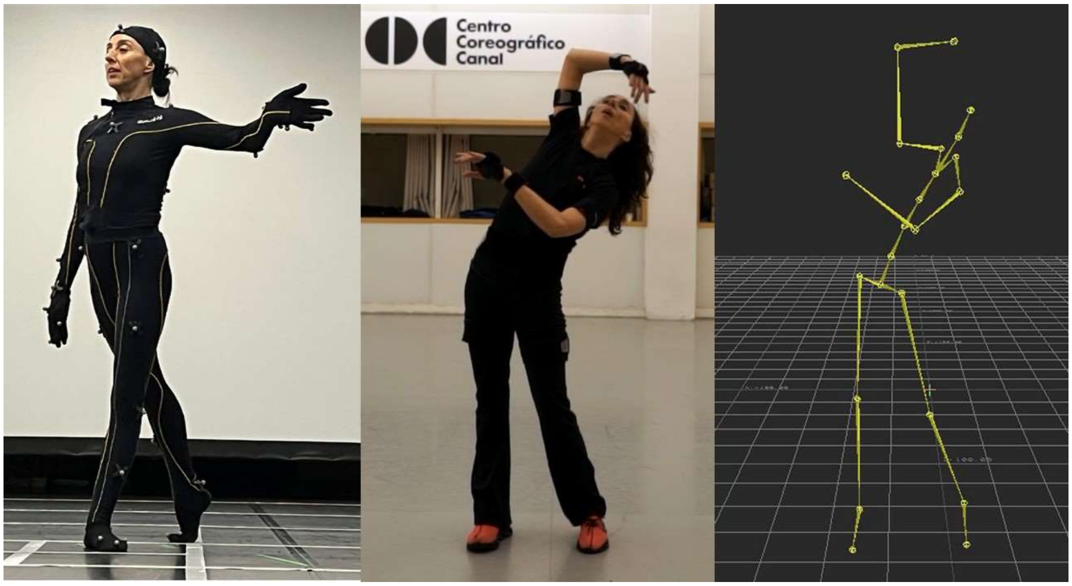

Figure 1: Motion capture sessions to record different movement qualities of Muriel Romero. From left to right: setup using a Qualisys optical motion capture system, setup using a XSens IMU-based motion capture system, a skeleton representation based on a motion capture recording.

Two datasets serve as resources for motion, sound, and image material. These resources have been extensively used to train the ML-based tools that form part of the AI-Toolbox. In addition, the datasets are also used to document, archive and disseminate some of the idiosyncratic dance techniques that are employed by the two dance companies (STO and ICK) both of which are part of the Premiere consortium (T3.5).

The first dataset has been created in 2021 during a Marie Curie Fellowship (H2020-MSCA-IF-2018 - Grant agreement no: 840465) in which STO played the role of a secondment partner. This dataset contains motion capture, audio, and video recordings of Muriel Romero performing different movement qualities that she has developed and works with in her choreographies. The recordings have been conducted using a Qualisys motion capture system that was kindly provided by the dance company Cie Gilles Jobin. The dataset is available on Zenodo[1].

A second and new dataset is currently in the process of being created. This dataset comprises solo and duet performances of idiosyncratic movement principles developed and employed by STO and ICK. This dataset contains motion capture, video, and audio recordings that meet the recommendations for capturing and storing performances (T2.3). The recordings have been conducted using an XSens motion capture system. This system has been chosen by PREMiERE consortium as standard for obtaining ground truth motion data.

[1] https://zenodo.org/records/7034917

4. Tools

The AI-Toolbox provides a large set of tools that employ both ML-based and explicit algorithms to analyse or generate movement, sound, and images. The ML-based and explicit algorithms play complementary roles. Since ML-based algorithms learn from data, they excel at capturing nuanced and idiosyncratic aspects of performance recordings if these recordings are used as data to train the algorithms. Explicit algorithms do not learn but instead consist of a fixed set of instructions that have been manually programmed. Explicit algorithmis typically provide more predictable results, and their output is often of higher quality (in particular in the case of sound synthesis).

The AI-Toolbox follows a modular approach in that each tool can function in isolation or combined with other tools (that are or are not part of the AI-Toolbox). This interoperability, which is based on the OSC protocol, facilitates the integration of the functionalities provided by the AI-Toolbox into the software infrastructure that is frequently employed in the fields of creative coding and computer music. The OSC protocol has been designed for real-time control of sound and media and offers great flexibility with regards to content and functionality. It also forms the basis for combining the AI-Toolbox with the Virtual 3D Theatre (T5.3). The OSC protocol is employed in a bidirectional manner to send data generated by the AI-Toolbox to other applications or to receive control commands to interact with and alter the configuration of the AI-Toolbox.

The AI-Toolbox focuses on the provision of algorithms that operate in real-time. In the performing arts, real-time capabilities are a pre-requisite to establish tight feedback loops between gesture-based interaction and media generation. For sound control, the latency between gesture and acoustic result should ideally not exceed 10 ms. For other types of synthetic media such as images or motions, the latency can be larger (< 200 ms) In the context of the Premiere project, a focus on real-time capabilities extends the usability of the AI-Toolbox beyond use case 4 and makes it also attractive for use cases 2 and 3.

Most of the ML-based tools that form part of the AI-Toolbox employ relatively small and simple models. This choice of models has been made for several reasons: to run the models in real-time interactive settings on standard consumer PC hardware, to train the models with small datasets that can be created by the artists themselves, and to facilitate the modification of the models by users who possess programming skills but are not experts in ML.

The AI-Toolbox caters to target users with a wide range of technical skill levels. For users who do not possess any technical skills, the AI-Toolbox provides tools that can be run as standalone applications and that are pre-configured to generate interesting synthetic motions, sounds, and visuals "out of the box". The AI-Toolbox is also attractive for users who possess regular skills in working with computers. These users can change the configuration of the tools either by modifying control parameters or by training the ML-based tools with new datasets. The AI-Toolbox is also attractive for users who possess some programming skills. These users benefit from the fact that the tools are open source and implemented in programming languages that are popular among the respective communities (Python for tools involving ML and Sclang for tools involving explicit algorithms for sound synthesis). Accordingly, programming proficient users can easily modify the source code of the tools and therefore thoroughly modify their functionality.

The following subsections describe each of the tools that currently form part of the AI-Toolbox in some detail.

4.1. Tools for working with motion capture recordings.

Motion Capture Player

Figure 2: A screenshot of the motion capture player that forms part of the AI-Toolbox.

The AI-Toolbox provides a tool for reading and playing motion capture recordings that are stored in the BVH format. BVH is a common but fairly old standard for storing skeleton and motion data in a human-readable format. The excerpt of the recording and the rate at which it is played

4.2. Tools for obtaining and analysing motion capture data

The AI-Toolbox provides several tools for obtaining and analysing motion capture data.

XSens to OSC Converter

For use in combination with the XSens motion capture system, the converter tool receives motion capture data obtained through the native XSens protocol and forwards this data using the OSC protocol.

Pose Estimation from Video

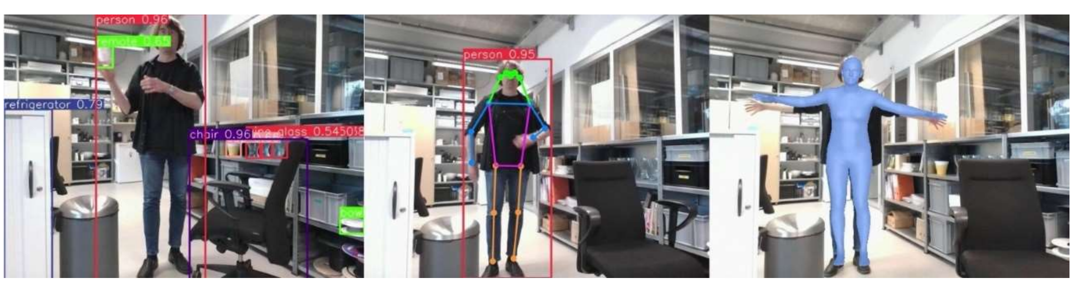

Figure 3: Screenshots of several ML-based computer vision models that have been integrated into the AI-Toolbox. From left to right: YoloV8 object detection, YoloV8 2D key point detection, ROMP 3D key point and shape detection.

Funded by
the European Union

To facilitate affordable motion capture using low-end consumer cameras, the AI-Toolbox integrates three existing ML-based computer vision models. These models infer from a live stream of monocular camera images the location and type of objects (YOLOv8 Detect), the 2D positions of body key points (YOLOv8 Pose [Terven et al. 2023]), or 3D positions of body key points and shapes (ROMP [Sun et al. 2021]). The corresponding tools wrap the functionality of these computer vision models and forwards the models' outputs using the OSC protocol. It is planned to further extend this functionality of the AI-Toolbox by integrating those computer vision models that have been selected as part of T4.2.

Motion Analysis

The AI-Toolbox provides a tool for analysing motion capture data. This tool integrates a set of explicit algorithms to derive from raw kinematic measurements higher level quantitative or qualitative movement features. The quantitative features include temporal derivatives of joint positions and rotations (velocity, acceleration, jerk) and the spatial extension of body parts (bounding sphere, bounding box). The qualitative features include the four Laban Effort Factors (Time Effort, Flow Effort, Weight Effort, Space Effort). The Effort Factors form part of the Laban Movement Analysis (LMA) system that has been introduced by Rudolf Laban [Laban 1975]. The Effort Factors describe aspects that relate to the dynamics, energy, and inner intention of movement, all of which contribute to the expressivity of movement [Bartenieff and Lewis 2013]. The computational analysis of the Laban Effort Factors is based on their mathematical description by Larboulette and Gibet [Laboulette and Gibet 2015]. This tool forwards the analysed movement features using the OSC protocol. For future versions of this tool, it is planned to incorporate an ML-based model that can be trained to classify idiosyncratic movement qualities (T3.2).

Motion Clustering

The AI-Toolbox provides a tool to group motion excerpts according to their movement features. The tool employs a simple unsupervised clustering method (K-Means). Once the clustering has been completed, the motions in each cluster can either be played back or exported as BVH files. When playing, the tool sends pose information using the OSC protocol. While this tool is primarily meant to be used to study choreographic movement vocabularies, it can also be employed for generative purposes by collaging the grouped motion excerpts into a new movement. For future versions of this tool, it is planned to provide motion clustering across modalities (e.g. clustering of motions based on image and audio features and vice versa, clustering of audio or images based on movement features).

4.3. Tools for movement synthesis

The AI-Toolbox currently provides three different ML-based tools for generating synthetic movements.

Movement Autoencoder

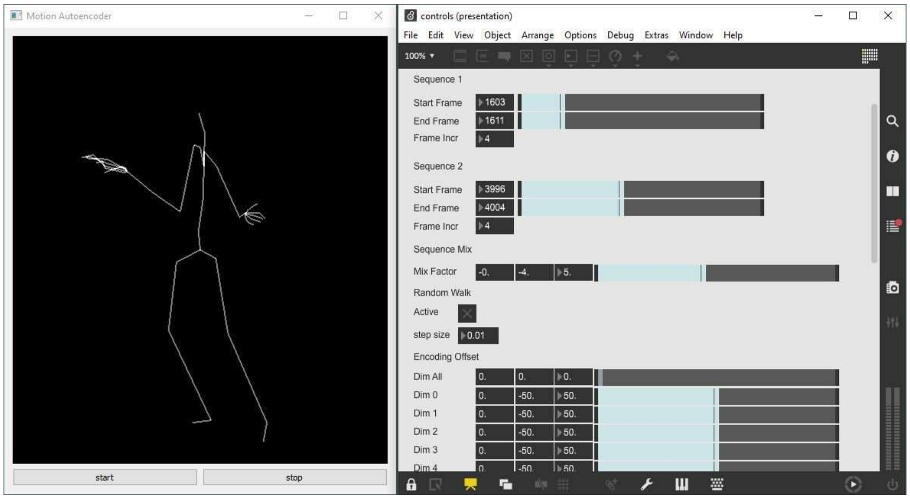

Figure 4: Screenshot of an adversarial autoencoder for movement synthesis. From left to right: the movement generated the model shown as simple line drawing of a skeleton, a graphical user interface to control the behaviour of the model.

This tool implements an adversarial autoencoder for encoding and decoding short motion experts and subsequently concatenating these excerpts into longer motion sequences. This autoencoder is based on an earlier implementation of a choreographic tool named Granular Dance [Bisig 2021]. The tool transforms existing movements into novel movements by interactively navigating the model's latent space. The tools provide different methods for latent space navigation which include direct manipulation of encodings, random walk-in latent space, latent trajectory following, and latent trajectory inter- and extrapolation. The tool either exports the generated movements as BVH files or plays them in real-time. When playing, the tool sends pose information using the OSC protocol. Future plans include an implementation of the autoencoder using skeleton aware graph convolutional neural networks [Aberman et al. 2020].

Movement Continuation

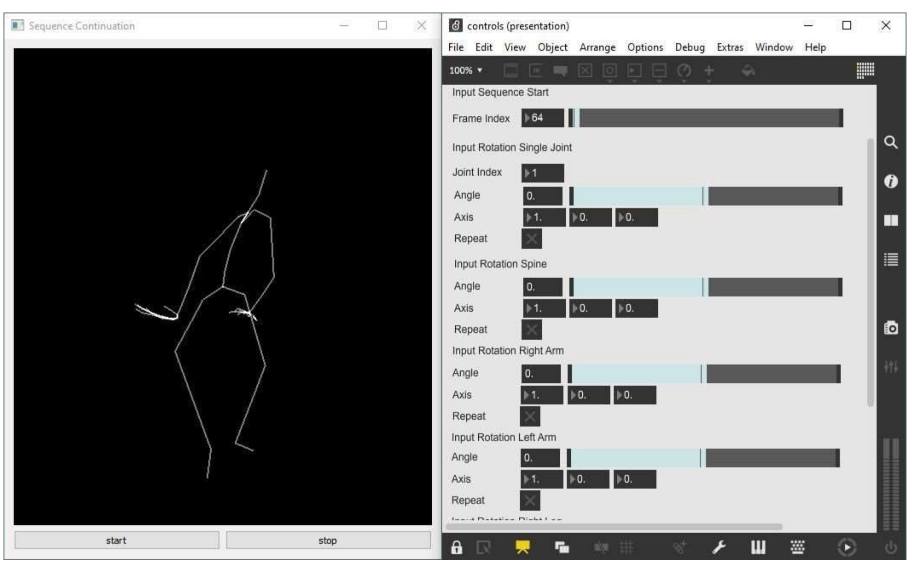

Figure 5: Screenshot of a sequence continuation model for movement synthesis. From left to right: the movement generated the model shown as simple line drawing of a skeleton, a graphical user interface to control the behaviour of the model.

This tool implements a recurrent neural network that predicts future poses based on a sequence of previous poses. This model is trained using a non-teacher forcing approach to improve its long-term prediction capabilities. In non-teacher forcing, the model learns to predict the next element in a sequence using its own previous and likely erroneous predictions as input. In the more conventional teacher forcing scenarios, the model is always provided with the correct input when making its predictons. Non-teacher forcing allows the model to predict longer sequences since it has learned to make predictions from more or less correct input data.The~~is~~ model can be interacted with by replacing some of the joint rotations in the sequence of input poses. If the rotations are replaced by joint rotations obtained through live motion capture of a dancer, the model can play the role of an artificial dancer that a human dancer can improvise with. The predicted movements can either be exported as BVH files or played in real-time. When playing, the tool sends pose information using the OSC protocol. Future plans include a replacement of this model with a Transformer-based implementation. The attention mechanism employed by Transformers have proven to perform better on long sequence prediction tasks than recurrent neural networks.

It is planned to add several additional ML-based models for movement synthesis to the AI-Toolbox. This includes a model that employs reinforcement learning to create synthetic movements for characters with non-humanoid morphologies. Reinforcement learning is a machine learning approach in which an agent learns through trial and error how to interact with an environment in a manner that maximises the agent’s long term reward [Sutton and Barto, 1998]. This model will be based on prior work on a system named “Expressive Aliens” [Bisig 2022]. Furthermore, a transformer-based model that translates music into motion and a diffusion-based model that are both currently being developed as part of T3.2 will also be added to the toolbox.

4.4. Tools for sound synthesis

The AI-Toolbox provides several different sound synthesis techniques that have been adapted for interactive movement sonification.

<u>Non-standard sound synthesis</u>

Non-standard synthesis models: two first order dynamic stochastic synthesis models in which amplitude and time values that conform a waveform (Gendy) are defined by probability functions. In this sound synthesis approach the waveform is articulated by a variable number of breakpoints (polygonyzed), each of them controlled by two random walks, one for the time values (abscissae) and another for the amplitude values (ordinates). Each new waveform is computed applying stochastic variatioins to the previous oneThe parameters of these functions can be interactively controlled by movement principles. An extended version of this approach concatenates several waveforms (Gendys) into a single one, each of them assigned to body parts.

<u>Physical Modeling Synthesis</u>

Physical models: So far three models have been implemented. One based on Feedback Delay Networks simulates bowed string sounds, ranging from extremely realistic sounds to complex abstract spectra. A second model simulates a membrane (waveguide mass physical model) and produces percussive sounds. A third model simulates the air flow in a tube (stochastic waveguide synthesis) and creates sounds ranging from multiple wind instruments, to breathing and wind sounds.

Additive/Subtractive Synthesis

Additive/Subtractive: A series of additive and subtractive models combined with stochastic processes adapted to create very realistic vocal sounds in multiple tessituras, percussion sounds and string sounds (plucked and bowed).

The AI-Toolbox also provides three ML-based tools for generating synthetic audio.

Audio Autoencoder

This tool implements an adversarial autoencoder for encoding and decoding raw audio waveforms. The tool is partially based on the Rave model architecture [Caillon and Esling, 2021] but replaces among others variational training by adversarial training. This tool can be used to transform existing waveforms into novel waveforms by blending their respective encodings.

Audio Deep Dream

This tool applies the Deep Dream principle for sound synthesis. The tool can be used to generate audio either in real-time or export audio as WAV files. The tool provides two different model architectures that have been pre-trained on three different audio datasets [Bisig and Wegner 2023]. Currently, work is conducted to train different models in a self-supervised manner by employing either audio reconstruction or contrastive loss functions. Once completed, these models will complement the currently available models.

Motion to Sound Translation

This tool employs a simple recurrent neural network to translate motion into audio. Audio is processed as MEL spectrograms that possess the same temporal resolution as the motion capture data (50 fps). After the model has generated synthetic MEL spectrograms, these spectrograms are converted into audio waveforms using the Griffim Lim algorithm. This model successfully recreates audio when working with motion it has been trained and also generates new audio of reasonable quality when working with new motions. At the moment, this tool doesn’t generate audio in real-time due to the computationally demanding conversion of spectrograms into waveforms. Future improvements of this model will address this issue either by employing an ML-based spectrogram conversion or by replacing the spectrogram representation with encodings generated by an audio encoder. Additional improvements will deal with the replacement of the model with a cross-modal transformer.

Future improvements of all the sound synthesis tools deal with a 3D binaural rendering of the generated audio and its steaming into the 3D Virtual Theatre (T5.3).

4.5. Tools for image synthesis

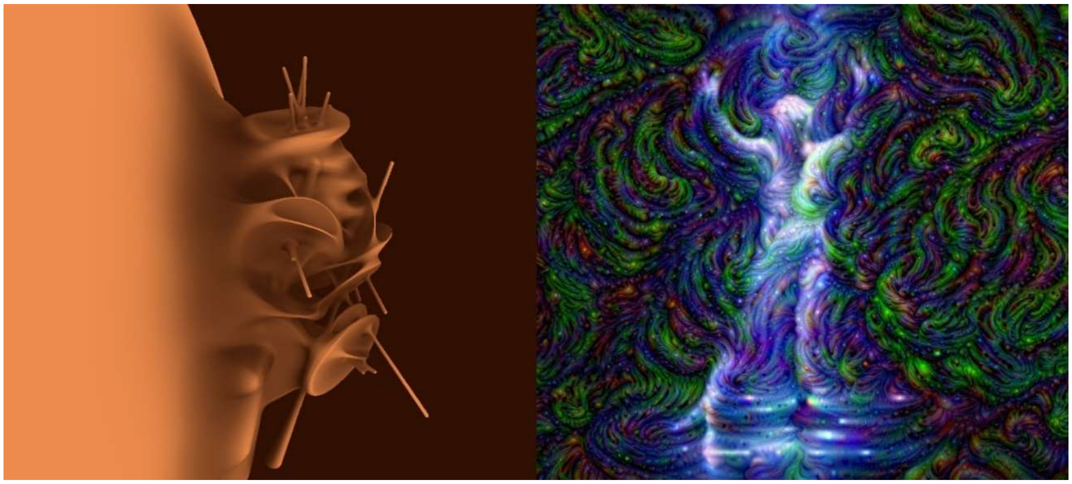

Figure 6: Two examples of synthetic images that have been generated using the AI-Toolbox. From left to right: A ray-marching-based visualisation of motion capture data, Deep Dream applied to a still image of a video recording of a dancer.

The AI-toolbox currently provides three tools for image synthesis.

Raymarching

This tool employs the raymarching rendering technique [Tomczak 2012] to render synthetic or captured motion capture data of a single dancer as abstract images. This tool can be configured to create a wide variety of visual results that range from humanoid skeletons consisting of platonic solids to amorphous surfaces. It is planned to adapt this tool to combine conventional mesh rendering techniques with raymarching.

Image Autoencoder

This tool implements an adversarial autoencoder for encoding and decoding images. This tool can be used to transform existing images into novel images and animations by interactively navigating the model's latent space. The tool currently doesn't operate in real-time. Future developments of this tool aim to improve its performance.

Image Deep Dream

This tool implements the Deep Dream technique to generate synthetic images by conducting gradient ascent on an input image to maximise the activity of individual layers and feature maps in a pre-trained model [Mordvintsev et al. 2015]. The tool currently offers two types of pre-trained models. One type includes image classification models that have been trained in a supervised manner on standard image datasets. The other type consists of the encoder part of image autoencoders that have been trained in a self-supervised manner on an image reconstruction task or using contrastive learning of visual representations. The Deep Dream technique can be used in combination with any of these models as long as the models have learned to discriminate between image features. Future developments of this tool aim to improve the quality of the generated images that are obtained from models trained in a self-supervised manner.

Additional improvements of all the image synthesis tools deal with the creation of stereoscopic images that can be streamed in real-time into the 3D Virtual Theatre (T5.3).

5. Dissemination and evaluation

Figure 7: Photos of workshop situations during which the AI-Toolbox as presented. From left to right: Stocos workshop in Gipuzkoako Dantzagunea, Choreographic Coding Lab in Cologne.

The AI-toolbox has been disseminated and evaluated among the main target user communities during several workshops.

The workshops that have been organised so far, the user communities that were targeted, and the functionalities of the AI-toolbox that were presented are as follows:

Gipuzkoako Dantzagunea (February 24 – 26, 2023)
Target Groups: Dancers, Choreographers, Musicians
Activities: Performing with interactive sonification models, Experimentation with machine learning models for movement synthesis.

Conservatorio María Ávila (March 6 & 7, 2023)
Target Groups: Choreographers
Activities: Performing with Interactive sonification models, experimentation with machine learning models for movement synthesis.

Choreographic Coding Lab Cologne (September 6 – 10, 2023)
Target Groups: Creative Coders, Media Artists, Experts in Dance & Technology
Tools provided during workshop: Motion Analysis, Motion Synthesis (Autoencoder, Sequence Continuation, Motion Clustering). Demo realisations: Duet performance between dancer and Autoencoder: mapping mocap on latent encoding. Facial tracking-based motion synthesis using autoencoder, mapping facial feature on latent encoding.

IDLab Workshop Amsterdam (September 21, 2023)
Target Groups: Dancers, Creative Coders, Computer Musicians
visualisation of live motion capture using raymarching, experimentation with machine learning models for movement synthesis, Live performance with movement sonification systems and robotic lights.

Two additional workshops are planned in the near future. One workshop will take place on February 1, 5, 8, and 12 at the Conservatorio María Ávila and targets mainly choreographers. A second workshop will take place from March 20 to 22 at La Casa Encendida in Madrid and targets dancers, creative coders, and computer musicians.

The AI-Toolbox will also play an important role in two upcoming dance productions. One dance production is realised as collaboration between STO and AHK and will be premiered in

January in Amsterdam. A second dance production with the title Incubatio will be realised by STO and premiered in November in Navarra.

6. Conclusion

The current version of the AI-Toolbox provides a collection of algorithms and tools that cover a wide range of functionalities for analysing movement and generating synthetic movements, sounds, and images. The need for these functionalities has been elucidated through surveys among the main target user categories of the AI-Toolbox. While this first version of the AI-Toolbox is still preliminary, it has already been received very well and adopted by the participants during several workshops.

In the upcoming months, it is planned to refine the algorithms and tools that are part of the AI-Toolbox. In parallel to these activities, a strong focus will be placed on the integration of the tools that are currently being adopted or developed by the Premiere partners. These are computer-vision based models for performance tracking (T4.2), generative models for movement synthesis (T3.2), and protocols and pipelines for exchanging media and control data with the 3D Virtual Theatre (T5.3).

The dissemination and evaluation of the AI-Toolbox through workshops will continue. But in addition to that, the application of the AI-Toolbox for the creation of new performance works will play an increasingly prominent role from now on.