Generative Choreographies: The Performance Dramaturgy of the Machine

Esbern T. Kaspersen[1], Dawid Górny[2], Cumhur Erkut[1][a] and George Palamas[1][b]
[1]Dept. Architecture, Design, and Media Technology, Aalborg University Copenhagen, DK-2450, Denmark
[2]Makropol.dk, Denmark

Keywords: Generative Art, Virtual Agents Animation, Dramaturgy, Neural Networks, Human Computer Interaction, Motion Expressiveness, New Media.

Abstract: This paper presents an approach for a full body interactive environment in which performers manipulate virtual actors in order to augment a live performance. The aim of this research is to explore the role of generative animation to serve an interactive performance, as a dramaturgical approach in new media. The proposed system consists of three machine learning modules encoding a human's movement into generative dance, performed by an avatar in a virtual world. First, we provide a detailed description of the technical aspects of the system. Afterwards, we discuss the critical aspects summarized on the basis of dance practice and new media technologies. In the process of this discussion, we emphasize the ability of the system to conform with a movement style and communicate choreographic semiotics, affording artists with new ways of engagement with their audiences.

1 INTRODUCTION

New advances in machine learning have empowered computational designers with advanced tools for data-driven design, exploiting the critical importance to form meaningful representations from raw sensor data (Crnkovic-Friis and Crnkovic-Friis, 2016). Among design practices, generative choreography can show promising results in producing movement phrases that exhibit motion consistency, realistic appearance and aesthetics. These phrases might then be utilized by virtual performers, sharing the same stage with humans, such as soft agents, robots or other mechanical performers (Schedel and Rootberg, 2009).

In a research creation context, we examine three independent neural network architectures trained on raw motion data, captured from a human performer, which then used to generate original dance sequences. Except of the primary objective for collaborative human-machine choreography, we believe that the proposed system would be a useful tool for artistic exploration.

[a] https://orcid.org/0000-0003-0750-1919
[b] https://orcid.org/0000-0001-6520-4221

1.1 Dramaturgy and New Media

Full body interactive environments have increasingly become part of the dramaturgy of live performance events (Seo and Bergeron, 2017). On the other hand, autonomous virtual agents, a new area of research, can provide an attractive abstraction that is driven from human motion data. These virtual agents, or robots, should be able to communicate ideas, symbolisms and metaphors.

New media definition is changing as per requirements and technology advancements. Defining what is new media is not easy (Lister et al., 2008). It is what makes their artifacts, practices, and arrangements different from those of other technological systems but also the social dimensions; the exchange of ideas is of primary importance to new media (Socha and Eber-Schmid, 2014). By considering this we can form an opinion whether performers can share the same stage with virtual actors that possess a degree of autonomy (Kakoudaki, 2014). These actors may act without the intervention of a human but the human might be able to influence the behavior of the agent.

However, most proposed systems suggest one to one mappings between a set of human actions and a direct visual interpretation of the expressiveness of the body. The main approach usually involves esti-

Kaspersen, E., Górny, D., Erkut, C. and Palamas, G.
Generative Choreographies: The Performance Dramaturgy of the Machine.
DOI: 10.5220/0008990403190326
In *Proceedings of the 15th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2020) - Volume 1: GRAPP*, pages 319-326
ISBN: 978-989-758-402-2; ISSN: 2184-4321
Copyright © 2022 by SCITEPRESS – Science and Technology Publications, Lda. All rights reserved

mating a set of body state, such as pose or motion, that link to a visual element. In this approach, a team consisting of software engineers and artists work together to build multi-modal interactions and immersive visual effects for stage, such as real-time projection mappings (Mokhov et al., 2018). But the dialogue between the dancer and the 'visualization' is predetermined, automatically executed, and the actor has little input to explore.

For these reasons we seek to develop an experimental framework where we can describe, conceptualize, design and direct an interactive performance. Observing the developments in new media, such as generative art, and the duality of autonomous virtual actors and human performers sharing the same stage, we can see a re-position of the performer's presence as a creative entity. Within this emerging role, the performer's motivation can be drastically changed, influencing the artistic development and the audience perception.

2 BACKGROUND

The use of generative technology provides new ways to explore and express the artistic intent. Some supporters of generative systems consider that the art is not anymore in the achievement of the formal shape of the work but in the design of a system that explores all possible permutations of a creative solution modulated by the quality of the dialogue.

2.1 The New Role of the Performer

The use of highly sophisticated video and audio content, along with powerful projectors and machines, such as robots, reflects a subtle re-positioning of bodily presence, rather than signaling the absence of the human bodily presence (Eckersall et al., 2017).

Moreover, new media have given rise to multiple forms of distributed co-presence (Webb et al., 2016), between performers and audiences, across a range of performative spaces, both real and mediated (Feng, 2019).

The integration of new media, especially computationally generated visuals, with a live performance is nonetheless a common practice (Grba, 2017) merging real and virtual worlds into a single experience. A number of works draw on a range of technologies exploiting emergent technologies such as motion tracking.

2.2 Human Motion Tracking

Human motion tracking can be achieved with a wide range of technologies that utilize optical, magnetic, mechanical and inertial sensors. The tracking accuracy depends on the sensors, as well as post-processing, the intended application, dramaturgical purposes and environmental parameters.

2.2.1 Motion Features

A multitude of features related to human movement expressiveness have been used to drive interactive performance environments (Alemi and Pasquier, 2019). The quantities that sensors capture are summarized in (Alemi and Pasquier, 2019) within these categories:

1. Joint positions and rotations: motion capture systems.
2. Joint acceleration and orientation: accelerometer and gyroscope.
3. Biometric features: electromyography, electroencephalography, breath, heart rate, and galvanic skin response.
4. Location of the body: Radio Frequency ID (RFID), Global Positioning System (GPS), and Mobile Networks.

As we will see in Sec. 3.1, our motion capture system makes use of the first two categories within an integrated hardware/software package.

2.3 Generative Animation

Generative art, as an artistic approach, utilize an autonomous system controlled by a set of predefined elements with different parameters balancing between unpredictability and order. Thus the generative system produces artworks by formalizing the uncontrollability of the creative process (Grba, 2017), (Dorin et al., 2012). According to (Galanter 2003) "Generative art refers to any art practice where the artist uses a system, .., which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art".

Generative animation can encompass a range of stochastic methods for motion synthesis, modulated by the motion of a human performer. Consequently, generative animation can be seen as a way of exploring a space of creative solutions spanned by a set of choreographic rules.

2.4 Human Computer Interaction

The research on movement generation usually follows one or more of following three themes (Alemi and Pasquier, 2019): (a) achieving believability, (b) controlling and manipulating the characteristics of the generated movements, and (c) supporting real-time and continuous generation. Believability is one of the fundamental notions in virtual agent animation and human-computer interaction (HCI); even non-movement-expert audience can notice the smallest details that make movement look unnatural (Alemi and Pasquier, 2019). Believability is usually achieved by (a) *physical validity*, and (b) *expressivity*. Next, for a meaningful dialog between the performer and the audience through the virtual agents, the agent movements should be controllable. The control should allow the natural *motor variability*. Humans never exactly repeat the same movement even when they try to do so. Consequently, the virtual agents that replicate the same execution will be perceived as more mechanical than natural. The most important factor for an interactive system is the real-time computational constraints (c), which we visit in the next section.

3 DESIGN AND SYSTEM IMPLEMENTATION

Our implementation is a multidisciplinary framework that uses full body interaction to create believable, that is physically-valid and highly expressive visuals. This is influenced by the meta-instrument concept, although this research does not focus on gestural interaction between performers and virtual actors. Before presenting the system implementation as a whole, we first introduce movement acquisition and feature tracking.

3.1 Feature Tracking

To capture the movements of the dancer, a Rokoko Smartsuit Pro[1] was used. This suit uses 19 inertial measurement units (IMUs) to sense the orientation of the body parts they are placed on. The sensors are placed in a wearable, which has sensor slots at various joints, as well as along the back of the suit (hip joints). Each sensor outputs a quaternion rotation vector and a position vector. The sensors have a sampling rate of 30Hz.

[1]https://www.rokoko.com/en/products/smartsuit-pro

3.2 Data Acquisition and Data Flow

The data from the suit was streamed, thourhg Rokoko Studio, to a Unity 3D[2] scene, using the Rokoko Unity plugin, where it was used to animate a digital avatar. The data was put to a machine learning model and the obtained results were streamed back to the host computer and applied to another avatar in the unity scene. This streaming between computers was performed with Node-red[3] using open sound control (OSC) connections.

The Unity scene was also rendered on the host PC, and sent to a HTC Vive VR head-mounted display (HMD), which the dancer was wearing. The scene was also sent to two other computers, one rendering the viewpoint of the guest and sending it to their HTC Vive VR HMD, and another one rendering the scene from a third person perspective. The last renderer was routed to a large-screen LED TV for the audience.

Before being used for training, a set of joints were chosen because they were considered to be vital to the expression of the dance, such as the knee, elbow and hip joints. Other joints were discarded, such as most vertebrae joints, since these could be approximated by knowing the quaternions of the hip, and neck joints. This selection was done to eliminate redundant information, as many of the vertebrae joints would often have similar rotations, and thus provide an easy local minimum for the neural networks to find. The approximation of the rotations of the discarded joints rotations was done using inverse kinematics when the predicted joint quaternions were applied to the model in Unity. The data used for training the machine learning model was obtained from a performance by a dancer wearing the suit and was written to a file instead of streaming. This data flow is illustrated on Fig. 1 and includes both the positions and quaternions for each sensor. The details of the data processing and filtering roughly correspond to the general workflow in motion retargeting, see e.g., (Villegas et al., 2018). For more details on the quaternion-based learning, see (Pavllo et al., 2019).

3.3 Training an Agent

Three approaches for modeling this data were experimented with, 1) mixture density networks (MDN) (Bishop, 1994), 2) auto-encoders (Liou et al., 2014), and 3) long short-term memory (LSTM) networks (Hochreiter and Schmidhuber, 1997).

[2]https://unity.com/

[3]https://nodered.org/

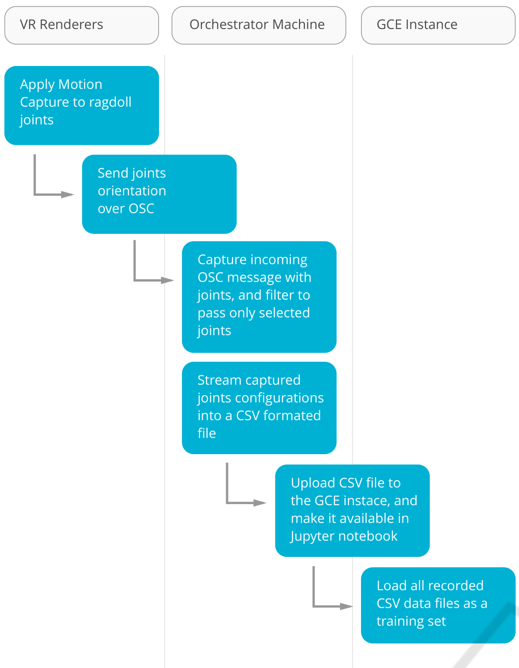

Figure 1: Training data flow. The dancer's movements were captured by the suit, smoothed and the resulted quaternions saved into a CSV file.

3.3.1 Mixed Density Network

A MDN was applied to this problem with promising results (Alemi and Pasquier, 2019) and thus was the first technique that was focused on. A MDN is a network that does not produce a single definite output like a normal neural network, but rather the mean and standard deviation (STD) of a Gaussian distribution (figure 2). From this distribution, a value can then be drawn and used as an output. This is useful if one input has to map to different outputs, since this is a case normal neural networks cannot handle well. However, in the case of an MDN this input can map to a Gaussian distribution that covers all output possibilities, as well as how often those possibilities were encountered during training.

All of these possibilities can sometimes not be covered by a single distribution, without that distribution having too high of a standard deviation to be useful. In this case multiple distributions should be used. This is done by making the MDN output more pairs of mean and STD, along with a third parameter, the mixture coefficient, which is a weighting of that distribution in the overall distribution.

By using this technique it would be possible to teach the AI the distribution of possible dance moves based on a given pose of the dancer. In this case, each quaternion distribution describes a space of choreographical solutions. These distributions could then be combined to a single multi-variate Gaussian distribution, for each quaternion, from which a new quaternion could be sampled.

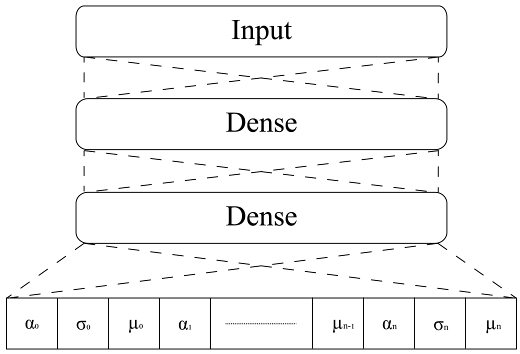

Figure 2: The architecture of the MDN network used. The α, σ, μ represent mixing coefficient, standard deviation and mean respectively for each of the n inputs.

In theory this approach would generate all possible movements from a given pose, however in practice this was not always desirable. If the STD of the distribution became large, the method would sample from a too wide range of quaternions, generating wildly different movements very quickly. This happened often during the training of the network, and caused the generated movements to seem random.

This problem was caused by the network not having the right amount of distributions available to describe the data accurately, and thus it instead generated larger distributions that covered the data too much. A possible solution to this might be an exhaustive fine-tuning the amount of means and STDs outputted for each quaternion, along with the hyperparameters used to train the network. Instead another model was considered, the Auto-encoder, as it was known to be easily trained and is also a powerful network structure for generative tasks (Shu et al., 2018).

3.3.2 Auto-encoder Network

Auto-encoder is a self-supervised network topology for representation learning (Liou et al., 2014). It can be thought as a generative model that can generate outputs sharing common structural elements, such as correlations, with the trained data. Auto-encoders can automatically exploit these relationships without explicitly defining them (Liou et al., 2014). An auto-encoder is a neural network that has a "hourglass" structure, meaning it compresses its input and then extracts the output from that compressed representation 3. These structures are usually used in an unsupervised manner, where the output of the network is trained to be the same as the input, i.e. if a 10 is

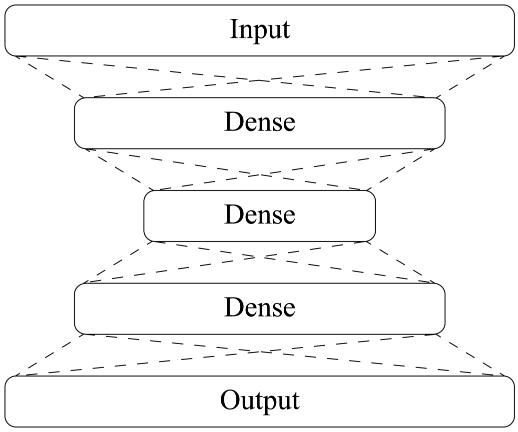

Figure 3: The hourglass-like architecture of the Auto-encoder. The output is trained to be the same as the input.

inputted a 10 should be outputted. This has the effect, that a very abstract representation of the input is found in the middle of the network (the most narrow point). From this representation, the output can be reconstructed. In this case, this means that every pose of the dancer is represented by a point in this abstract space, and thus if a translation in this space occurs, a new pose will be generated by the network. The idea was therefore to input the real-time movements of the dancer into this network, obtain the abstract representation, translate that representation by some set amount and then generate a new movement which would be applied to the avatar.

Autoencoders are easy to train as reconstructing the input is a simple task if the abstract space is big enough. For these experiments we used a 5 dense layer structure, 3 layers for the encoder and 2 for the decoder. The input to this model was a 1-D array of 15 quaternions with 4 values each, totalling 60 values. The middle layer contained 20 neurons. This short topology was chosen as deeper topologies would be more prone to overfitting. However, this approach over-fitted too easily and thus produced wrong results for inputs it had not encountered before. These errors were also unstable, meaning a small variation in input, resulted in a big variation in output. The movements the network produced over time, were very jittery and noisy, as the small movements of the actor between frames resulted in larger unpredictable movements of the avatar.

3.3.3 Long Short-term Memory Network

Long short-term memory networks(Hochreiter and Schmidhuber, 1997) were the answer to the instability problem of the auto-encoders. The LSTM network consisted of three LSTM layers, of which the first two output a time-series while the third only output a single vector. Rather than processing a single pose of the dancer at a time, this network processed the last 5 poses and tried to predict what the next pose would be. This resulted in the network learning about temporal coherence, and made it more stable than the auto-encoder.

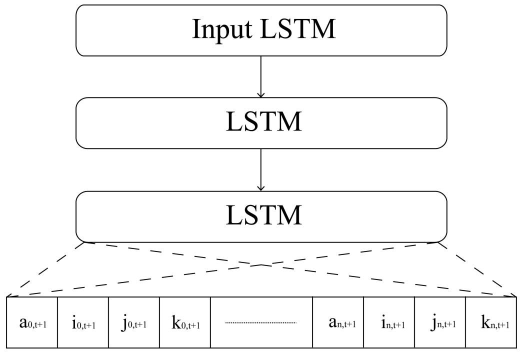

Figure 4: The 3 layer LSTM structure trained on a time series of t (5) $\times$ n (15) quaternions. The output is the predicted real coefficients of the n (15) quaternions at time t+1.

The LSTM layers can learn temporal relationships in the data due to the way they process the data. As opposed to traditional dense layers that process all data at once, LSTM layers are recurrent layers that iterate over time-series and can find their temporal relations. For each iteration the current data-point is processed by a dense layer, together with the output from the previous iteration. The result of this iteration is then stored to be used for the next iteration. This result can also be used as an output, if it is desired that the layer should return the full time-series. If only a single output is desired, only the output of the last iteration will be returned.

However, the vanishing/exploding gradient problem(Hochreiter, 1998) can occur if no other processing is done between iterations. This is the problem that LSTM layers solve, as they use forget-gates to retain or discard information between iterations, thus only important information is retained.

The LSTM network was trained to predict the movements of the dancer, rather than replicate them. This is a very important aspect to generative design because the generated movements were in fact interpretations rather than reproductions. The LSTM was also more stable during training, less prone to overfitting, as predicting the next move is a much more difficult task than simply reproducing a pose. Because of these reasons, the LSTM network was used instead of the auto-encoder and the MDN.

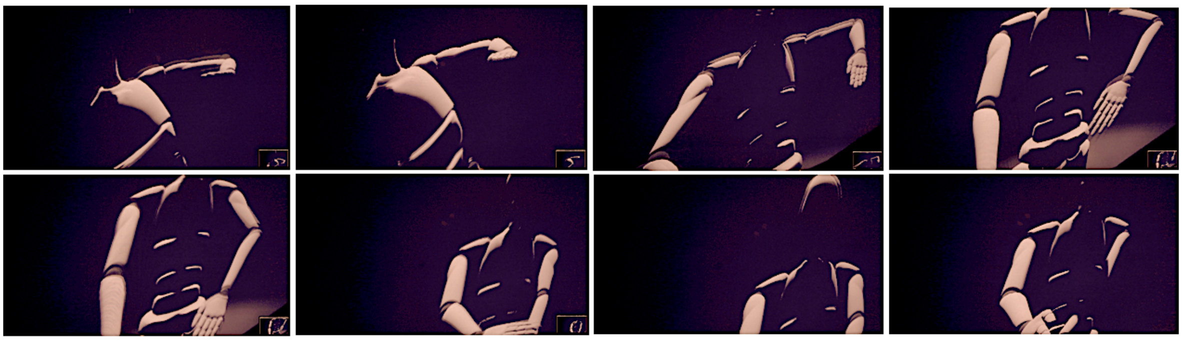

Figure 5: A pre-trained (using LSTM) virtual agent responding to the movement of a dancer, from the audience view.

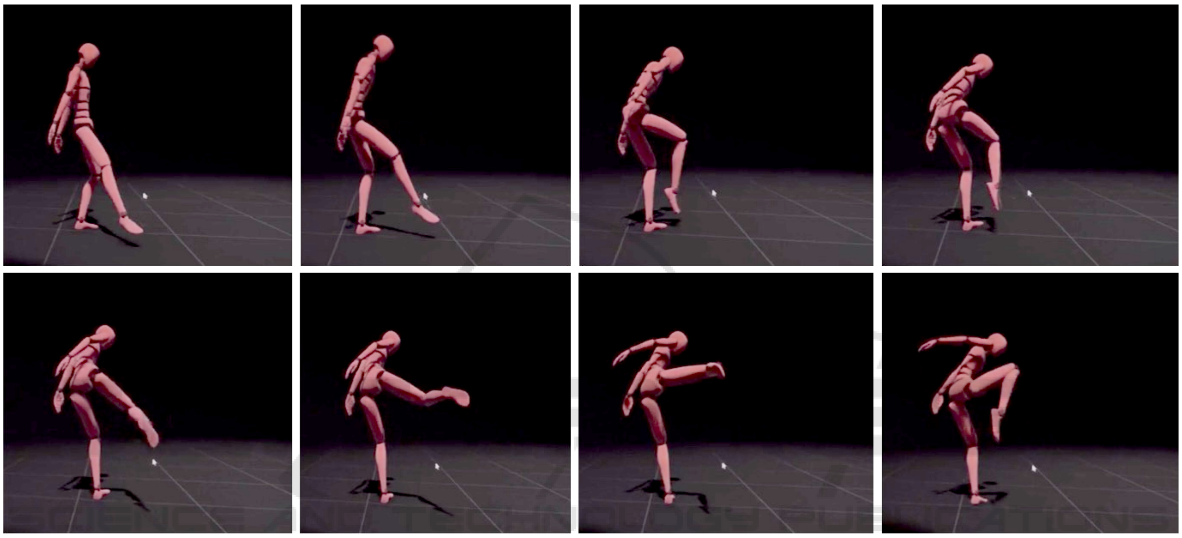

Figure 6: A dancer in a free improvisation, represented by a sequence of pose estimations captured from the motion data suit and remapped the default avatar. The extracted quaternions (n=15), which are not explicitly shown on the figure, are used to train all three neural network architectures. All frames correspond to the dancer, and none to the neural network.

Table 1: A comparison of the three models in producing choreographic phrases, based on motion's expressive qualities.

HUMAN EXPRESSIVENESS - MODEL	**MDN**	**AE**	**LSTM**
Interaction Irregularities	fair	poor	fair
Posture Prediction and Temporal Coherence	fair	poor	good
Overall Appearance and Aesthetics	poor	fair	good
Motion Consistency	fair	fair	good

4 RESULTS AND DISCUSSION

Human expressiveness should be assessed based on affective and cognitive cues (Billeskov et al., 2018). Visual inspection provides an immediate, though subjective, means of qualitatively comparing the three machine learning algorithms from the perspective of human expressiveness. The generative choreographic model, tested with three different modules, and obtained a variety of choreographies which then inspected in terms of posture prediction, action irregularities, motion consistency and overall appearance and aesthetics 1. The proposed system was capable to generate expressive choreographies following the motion style represented in the training data. In a kinematic level, the system was capable to reveal basic joints rotational relationships while in a more abstract level, choreographic semiotics and symbolic qualities were being communicated, derived from the training dancing moves. Moreover, the system was capable to form hybrid compositions, based on different datasets, corresponding to a variety of dancing samples.

The three modules depicted diverse behaviors and results. The MDN was capable of extracting the periodicity of the human performer, which can be attributed to the fact that an MDN can generate samples conforming to a Gaussian distribution. The use of a mixture coefficient, as a hyper-parameter, revealed the potential of generating dance moves triggered from a single choreographic pose. This led us to conclude that the MDN can be a useful tool to explore a vast space of possible choreographic solutions. However this advantage might come at a cost. If the standard deviation of the distribution became large, the method tends to generate a series of movements manifested as uncoordinated motion, highly variable, mostly perceived as jittery motion. The variational auto-encoder, as a self-supervised network doesn't require an explicit mapping between the input data and the generated motion data. This is a very attractive approach but also tend to over-fitting quite often, thus it requires a very large set of training data to generate meaningful outputs. This model was accurate on reproducing known series of poses but also produced wrong results for sequences it had not encountered before, possibly due to over-fitting. Thus, small variations in the input was producing big variations in the output which in turn produced perpetual oscillatory motion. The LSTM network considered to be the most stable from all the cases in this study. This network depicted the most consistent temporal behavior generating movements that were interpretations of the trained data, rather than reproductions. This can be explained due to the fact that recurrent topologies are robust on predicting with less over-fitting.

4.1 Case Study: Singularity

Singularity was a live, interactive performance that took place in the Multi-sensory Experience lab [4] of Aalborg University Copenhagen.

The concept of this performance, was to represent the acceleration of machine learning and AI towards the so-called "Singularity Threshold", where AI's will surpass humans in every task. By making a machine learning model that copies, and predicts, human choreography it is shown that even artistic expression is not safe from this advance. By immersing both the dancer and the guest in the digital world, through head-mounted displays, they can both see the effects of the dancers movements on the AI. How, when she accelerates, the AI does too. And how the AI starts to split and multiply, outnumbering the humans at the end of the performance. The dance itself was a ritual, a sort of abstract dance, which the guest also partook in. Furthermore, the changing soundtrack, played through the 360 degree speaker system, reinforces the immersion of the guest and dancer in the world. Audiences outside could see this performance from a third person perspective, through a screen as shown in figure 5

[4]https://melcph.create.aau.dk

5 CONCLUSIONS AND FUTURE WORK

In this paper, we detailed a generative animation system and compared three independent neural network architectures. The system was used to generate dancing animations of a virtual agent interacting with a human dancer with the main effort focusing on the machine learning approach to generate dancing sequences. Comparing the three neural network architectures, it can be hypothesized that a fusion of these approaches might result in better solutions that share the best characteristics from each individual approach. By developing a fusion of these models and combining them in linear combination, we expect to generate more complicated and more expressive motion phrases. Another interesting future direction could be to compare our machine learning approaches to the evolutionary methods used in behavioural simulation for autonomous characters based on motion or choreographic rules. More experiments by artists could lead to user experiences data to be worked on and also to feed the AI engine for enhancing the learning process.

5.1 Communication and Aesthetics

Communication through semiotics of body movement and attitude is a very important factor, that is also closely associated to perception of aesthetics. The evaluation of computational aesthetics is a difficult problem, with most current paradigms drawing insights from models of human aesthetics such as Arnheim or Martindale. Galanter, (galanter2011) suggests putting the focus on empirical evidence of human aesthetics and the emergence of complex behaviors rather than algorithmic complexity. Neural networks are capable of building inner representations, through an abstraction space consisting of hierarchical layers to abstract the source into a semantic, low dimensional, representation (Hinton 2014). By simplifying a very rich source of data, such as human motion, in which unimportant details are ignored, a set of semiotics could possibly emerge conveying messages through non-verbal means, such as gestures and body expressions.

5.2 Multimodal Approach

The input data could be extended, beyond motion capture data to include other sensorial input, both direct such as a music score or indirect such as using Electroencephalography (EEG) or other bio-signals from a participant on stage (Hieda, 2017). By exploring this sensorial diversity new choreographic forms and practices can emerge, redefining the role of the performer and their artistic relationships.

REFERENCES

Alemi, O. and Pasquier, P. (2019). Machine learning for data-driven movement generation: a review of the state of the art. *CoRR*, abs/1903.08356.

Billeskov, J. A., Møller, T. N., Triantafyllidis, G., and Palamas, G. (2018). Using motion expressiveness and human pose estimation for collaborative surveillance art. In *Interactivity, Game Creation, Design, Learning, and Innovation*, pages 111–120. Springer.

Bishop, C. M. (1994). Mixture density networks. Technical report, Aston University.

Crnkovic-Friis, L. and Crnkovic-Friis, L. (2016). Generative choreography using deep learning. *arXiv preprint arXiv:1605.06921*.

Dorin, A., McCabe, J., McCormack, J., Monro, G., and Whitelaw, M. (2012). A framework for understanding generative art. *Digital Creativity*, 23:3–4.

Eckersall, P., Grehan, H., and Scheer, E. (2017). Cue black shadow effect: The new media dramaturgy experience. In *New Media Dramaturgy*, pages 1–23. Springer.

Feng, Q. (2019). Interactive performance and immersive experience in dramaturgy-installation design for chinese kunqu opera “the peony pavilion”. In *The International Conference on Computational Design and Robotic Fabrication*, pages 104–115. Springer.

Grba, D. (2017). Avoid setup: Insights and implications of generative cinema. *Technoetic Arts*, 15(3):247–260.

Hieda, N. (2017). Mobile brain-computer interface for dance and somatic practice. In *Adjunct Publication of the 30th Annual ACM Symposium on User Interface Software and Technology*, pages 25–26. ACM.

Hochreiter, S. (1998). The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.

Kakoudaki, D. (2014). *Anatomy of a robot: Literature, cinema, and the cultural work of artificial people*. Rutgers University Press.

Liou, C.-Y., Cheng, W.-C., Liou, J.-W., and Liou, D.-R. (2014). Autoencoder for words. *Neurocomputing*, 139:84–96.

Lister, M., Giddings, S., Dovey, J., Grant, I., and Kelly, K. (2008). *New media: A critical introduction*. Routledge.

Mokhov, S. A., Kaur, A., Talwar, M., Gudavalli, K., Song, M., and Mudur, S. P. (2018). Real-time motion capture for performing arts and stage. In *ACM SIGGRAPH 2018 Educator's Forum on - SIGGRAPH '18*, page nil.

Pavllo, D., Feichtenhofer, C., Auli, M., and Grangier, D. (2019). Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision*.

Schedel, M. and Rootberg, A. (2009). Generative techniques in hypermedia performance. *Contemporary Music Review*, 28(1):57–73.

Seo, J. H. and Bergeron, C. (2017). Art and technology collaboration in interactive dance performance. *Teaching Computational Creativity*, page 142.

Shu, Z., Sahasrabudhe, M., Alp Guler, R., Samaras, D., Paragios, N., and Kokkinos, I. (2018). Deforming autoencoders: Unsupervised disentangling of shape and appearance. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 650–665.

Socha, B. and Eber-Schmid, B. (2014). What is new media. *Retrieved from New Media Institute http://www. newmedia. org/what-is-new-media. html*.

Villegas, R., Yang, J., Ceylan, D., and Lee, H. (2018). Neural kinematic networks for unsupervised motion retargetting. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Webb, A. M., Wang, C., Kerne, A., and Cesar, P. (2016). Distributed liveness: understanding how new technologies transform performance experiences. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, pages 432–437. ACM.