

Learning to replace a human: A virtual performing agent

John McCormick

Motion.lab

Deakin University Australia

221 Burwood Highway

Burwood 3125, Australia

john.mccormick@deakin.edu.au

Steph Hutchison

Motion.lab

Deakin University Australia

221 Burwood Highway

Burwood 3125, Australia

steph.hutchison@deakin.edu.au

Adam Nash

RMIT University Australia

Bldg 14, Level 11, Room 12

Melbourne 3001

adam.nash@rmit.edu.au

Kim Vincs

Motion.lab / Centre for Intelligent
Systems Research

Deakin University Australia

221 Burwood Highway

Burwood 3125, Australia

kim.vincs@deakin.edu.au

Saeid Nahavandi

Centre for Intelligent Systems
Research

Deakin University Australia

75 Pigdons Road

Warun Ponds 3216, Australia

saeid.nahavandi@deakin.edu.au

Douglas Creighton

Centre for Intelligent Systems
Research

Deakin University Australia

75 Pigdons Road

Warun Ponds 3216, Australia

douglas.creighton@deakin.edu.au

ABSTRACT - In this paper we describe two artworks, *Recognition*, an outdoor interactive installation and *Instrumental* a live dance performance. In both works a performing agent has learnt sequences of movement from a dancer and uses these to stand in for a human performer. The agent uses an Artificial Neural Network to learn to dance from the human dancer and can perform in the human's stead. In *Recognition* the agent's movement is used when there are no humans present in order to keep continuity of the installation. In *Instrumental* the agent becomes a performing partner of a live human dancer, able to recognize the dancers movement and synthesize movement sequences based on the human dancer's movements.

Categories and Subject Descriptors- J5 [Arts and humanities]: Arts, fine and performing.

General Terms - Performance, Experimentation, Human Factors.

Keywords - Performing Software Agent, Machine Learning, Electronic Art, Interactive Installation.

1. INTRODUCTION

In this paper we introduce two artworks, *Recognition* and *Instrumental*. *Recognition* is an interactive installation which uses a performing software agent to effect change in the virtual environment when there are no humans present. *Instrumental* is a performance between a dancer and performing agent performed at Motion.lab in Melbourne, Australia. *Recognition* was exhibited at Cube 37, a glass fronted gallery in Frankston, Australia. The exhibition environment used the movement of passers-by to animate a morphing avatar which was projected onto the front of the gallery. As the exhibition ran through the night, there were often times when there were no pedestrians around from which to gather movement information, however there was still a lot of vehicular traffic passing the gallery, which was located on a major road. Rather than have a still screen at these times, a software agent was trained to move using the movement of a dancer, and this agent's movement was used by the projected avatar to

animate itself. The use of learned human movement allowed the software agent to quickly acquire the capability to stand in for a human when needed. The projected forms also used pictures of the dancer's iris as a texture for its body, giving the installation a uniquely organic signature. Borrowing appropriate material from a human to quickly generate capacity for the agent was one of the features of *Recognition*.

Recognition is part of an ongoing investigation into digital performing agents, in particular agents that can learn to dance with a human. We viewed the agent as a performing partner and as such decided to treat the relationship between the agent and dancer in a similar manner to what might occur between two dancers. The dancers might generate dance sequences, perhaps through improvisation initially, they share and learn these movements, then in performance can use this shared movement vocabulary to perform in unison, perform independently or take cues from each other to determine what parts of the movement they might perform (semi – improvised). This is a very simplified structure however it suggested a learning model as the basis for the agent if it were to take the role of a performer. The agent would also need to be able to both generate movement and recognize what the dancer is doing. These requirements influenced our decision as to the type of structure to use for the agent's learning model. *Recognition* is both an interactive installation and a performance, for the majority of the time it is available for pedestrians to interact with, however at certain times a dancer may perform within the installation. The installation is essentially the same, the only difference is the familiarity of the




Figure 1. Human participants providing movement data for the visible avatar that is projected onto the glass frontage of the Cube 37 gallery.




Figure 2. Human movement is tracked by Kinect and used to animate a live avatar (silver avatar on right). At the same time the performing agent (red avatar on right) is continuously generating movement based on what it has learnt from the dancer. The silver and red avatars are not visible to the viewers, their data is used by the morphing avatar (left image behind dancer) to animate itself.

dancer with the system and the movement choices available to the dancer due to her flexibility and experience. However even the dancer is improvising as are the pedestrians.

Recognition focuses on the agent's ability to generate movement from what it has learnt and perform this in parallel with a human participant in order to maintain a constant flow of movement data for the installation.

Instrumental is a performance where the agent is able to recognize the dancer's current movement and create its own movement sequences using the dancer's movement as a starting point. The dancer and agent create an improvised dance duet based on the movement the agent has learnt from the dancer. The movement learning and creation draws on the movement creation techniques developed in *Recognition* with the added ability for the agent to recognize the dancer's current movements.

2. RELATED WORK

The requirement for human-like movement is very common as seen in a myriad of games, films and animations. Non-Player Characters as agents in a game are often animated with motion capture data to bestow human-like movement qualities on the characters. These libraries of movement are often blended together to form variations on the stock movements. Other research seeks to go beyond pre-recorded motion to give agents more flexibility in their movement. Hsu et al. applied time warping to the learning of motion style from a library of different motion-captured walks in order to generate variations on walking style. [1] Brand and Hertzmann used a Hidden Markov Model to create a style machine capable of learning different walking styles from motion capture data. [2] The style machine could re-synthesize movement with different styles, for example applying an experienced dancer's smooth style to a novice ballet dancer's movement. Taylor and Hinton applied Conditional Restricted Boltzmann Machines to the same task of learning movement style in walks and generating new variations from the trained network. [3] The latter in particular showed that a learning model could be used instead of a pre-recorded motion model for movement generation. However we were interested in a model that could be used for both movement generation and recognition as well as for learning a substantial movement vocabulary rather than movement style. For use in our virtual environments, stylistic consistency was more important than stylistic variability as we

wanted the agent's movement to blend in with the live human's movement. We also wanted to develop a framework that could be used for both performance and installation work. There has been a great deal of work done in the area of movement recognition. Han et al. used a combination of Self-Organising Map (SOM) and Markov Model for the prediction of human positional movement for location based services. [4] Son et al. used a Recurrent SOM to detect movement patterns in the body of *Lumbriculus variegatus* (a type of worm) in response to the introduction of toxins to their environment. [5] Caridakis et al. successfully employed SOM and Markov Model to track the 2D point trajectory of the hand for gesture recognition. [6] We applied SOM to full-body 3-dimensional movement as the basis for our agent's learning with a novel approach to its use for both movement generation and recognition.

3. RECOGNITION DESIGN

The Cube 37 gallery has a glass front onto which imagery is rear projected, allowing it to be viewed by passing pedestrians and people in cars. A Kinect sensor behind the glass screen tracked people's movement on the street in front of the gallery. (Figure 1) The Kinect data is used to extract a skeleton representation of the pedestrian's movement, or at times, the dancer's movement, and this data is used to animate an avatar which is in the background and is unseen by the participants. The joint positions of the unseen avatar are used by the visible morphing eye to change its shape accordingly. There is also another unseen avatar representing the performing agent. When no humans are present the visible eye takes its movement data from the agent's avatar instead. Thus the avatar representing the live human and the avatar representing the performing agent work together to continuously provide movement data to the morphing eye. (Figure 2)

When a person enters the area in front of the gallery they face a giant morphing shape that looks out at them from the gallery window. It changes according to their movement, fluidly transforming like liquid or molten metal. When it solidifies into a single shape it appears like a giant eye, casting its gaze over the pedestrian. Depending on the movement of the pedestrian, it can break into smaller parts or components, leading to extremely varied morphology. It invites improvised participation, either from passers-by or at specific times, from the dancer who provided the agent's movement. It is meant to be playful and engage both participants and onlookers who can appreciate the myriad forms brought about by the participant's movements.

When there are no human participants in front of the gallery, the great eye changes in color and texture from that borrowed from the iris of one of the artists, John, to that of the dancer, Steph. The movement behind its morphing forms also changes to the movement of the software agent, which has been learnt from Steph's improvisations. Thus, people in passing cars and pedestrians across the street are able to witness the continued dance of the morphing eye.

4. INSTRUMENTAL DESIGN

Instrumental used a 24 camera motion capture system housed at Motion.lab. The motion capture system tracked the movement of the dancer and streamed the data to the agent in real-time. The agent was able to use the data to recognize what the dancer's current movements were and then respond accordingly.



Figure 3. Dancer Steph Hutchison wearing a motion capture suit to capture movement with which to train the agent.

The agent and dancer each had their own humanoid avatars, which were projected onto a 10 metre screen covering the rear wall of the theatre. The 3D avatars and their environment were projected in stereoscopic 3D requiring the audience to wear passive stereo glasses. The dancer danced live in front of the screen and the two avatars were seen performing together on screen.

5. TEACHING THE AGENT

5.1 Artificial Neural Network

For *Recognition* we have chosen a particular type of Artificial Neural Network known as a Self Organising Map (SOM). [7] The SOM is an effective means of data mining as it is able to represent high dimensional data in fewer dimensions, allowing the data to be more easily visualized. It is also useful for clustering like segments of data into regions that help elucidate patterns inherent in the data. For *Recognition*, we are not so interested in the clustering capabilities of the SOM as much as its ability to adjust its internal weights to closely match those of the input data. By doing so it can create a map containing movement postures that can be traversed to generate movement sequences for the agent's avatar and in turn, for the giant eye, in the absence of humans. This feature of SOM is sometimes referred to as Associative Memory.

We have used a modified form of SOM that has multiple layers, the first layer contains information describing postures of the body, 79 weights equating to the position of the agent and its joint rotation angles. The second layer contains temporal information, potential pathways through the first layer to link the postures in order to produce movement. The SOM is an unsupervised form of ANN, the recorded movement data is presented to it to learn without any labelling or suggested outcome. Poses that are near identical will be encapsulated within the same neuron and neurons with similar poses tend to cluster together in the first layer. A single neuron may encapsulate a number of similar poses from the input data while other neurons may not have any. It is a competitive process accomplished during the learning phase. In order to provide the training data for the agent we used a motion capture system to record the movement of a dancer while

she was interacting live with the projected eye avatar. (Figure 3) The optical motion capture system used provided higher resolution data than the Kinect and resulted in movement that looked closer to the dancer's original movement. Recording the data as the dancer improvised with the installation resulted in movement that was similar to what a human would perform in the live installation, giving the agent an appropriate movement vocabulary to work with.

We used the same skeleton to record the data for the SOM to learn with as well as to animate the avatars. (Figure 4) The skeleton is relatively simple in order to keep the data size as small and optimal as possible so as to reduce the training time and allow the data to perform well in a live performance setting. The same skeleton is used for the avatars representing the live humans and the performing agent allowing their data to be interchangeable. The 19 joints of the skeleton produce 79 input vectors per frame of recorded movement. All rotations are local rotations and the positions are relative changes in position from the previous frame. The local rotations and relative positions allow the agent to generate new movements that follow on from its last position and posture so they are not dependent on where they are in the virtual world. Furthermore, as the agent's neural network can also be used for recognizing a human's movement [8] it is able to do so independent of the position of the human's avatar in the virtual world. The initial tests were done in Matlab and the final installation was developed using the Unity game engine. The trained SOM was imported for the agent to use as its movement memory from which it could generate appropriate sequences to animate its avatar.

The spatial layer of the SOM is visible in (Figure 5). It shows the results of a learning phase with neurons containing weights that describe the postures they have captured. Some neurons have many hits of similar postures, where a movement may have been held or repeated and so re-occurred in the input movement data. When a neuron is stimulated, its weights can be used to animate its avatar. By moving from neuron to neuron, the job of the temporal layer, the avatar is continuously animated by the SOM.




Figure 4. The skeleton used throughout the installation, from recording the movement data from the dancer to animating the agent and human avatars. The common skeleton allows the avatars and their data to be interchangeable within the installation.




Figure 5. SOM after training showing accumulated postures. The cells represent neurons and the numbers represent how many frames from the input movement data the neuron was able to match. (Best Matching Unit) The image at right shows what the weights of the neuron might look like when translated onto the avatar's skeleton.

6. RECOGNITION VISUALS AND SOUND

In *Recognition* we borrowed not only the movement from the dancer in order to develop the agent, but also other identifying features of the dancer and other humans. This allowed us to rapidly develop the agent while also giving it somewhat recognizable identifiers with which an audience could respond.

6.1 Clothing the Avatar

Besides providing human movement data to the agent in order to learn how to move, we also used images of Steph and John's irises as the textures for the main avatar. (Figure 6) This provided further connection to the human "donors" and a quick means of giving the main avatar a unique identity. When a human participant was present the main avatar used John's iris to clothe its body. When no humans were present and it was drawing on the agent's movement, it used Steph's iris. (Figure 7)



Figure 6. (Left) Steph's iris used by the main avatar when no humans are present and the agent's movement is in use, (Right) John's iris used when the human avatar movement is active.

The main avatar's body was produced using a marching cubes algorithm and tables courtesy of Paul Bourke. [9] This algorithm allowed the main avatar to constantly reform according to the movement data from the agent or live human.

6.2 Sound

Extending upon the manner in which the installation drew upon existing human data for its movement and visual sources, so too the sound for the installation drew upon the environment it was occurring in. The traffic passing by the front of the building created a peculiar undertone of noise with strong Doppler

characteristics. At dusk huge flocks of birds flew into the area and their calls were almost overwhelming. The sound emitted from the installation had similar qualities with a deep undercurrent of dopplered noise and short higher pitched overtones reflecting the calls of the birds. The sounds palette changed depending on whether the human or agent's movement data was being used and moved spatially according to the movement of the visible avatar. In the performance version of *Recognition* we used echocardiogram recordings of the heart. This was in keeping with the borrowing of human data to add capability to the agent. The recordings of the heart added a driving rhythm to the performance and gave the dancer another layer to interact with.

7. INSTRUMENTAL VISUALS AND SOUND

Instrumental began with the two different colored avatars on screen and the dancer performing in front. The avatars represented the live human dancer and the agent dancer respectively. The dancer would perform movements based on sequences of movements the agent had learnt. The agent would in turn synthesize a movement sequence using the dancer's current movements as a starting point. Thus the human and virtual dancer's movements would converge at times when the agent began a new movement sequence, and then diverge as its movements became different to the human dancer's.

As the performance of *Instrumental* progressed, the pathway of the agent was drawn as a continuous trail of crystal beams. The pathway changed every performance as it was a result of the interactive movement creation of the agent prompted by the improvisations of the human dancer. The pathway of crystal




Figure 7. (Left) The main avatar using Steph's iris and the agent's movement, (Right) using John's iris and live human (Kinect) movement




Figure 8. Recognition. The visible avatar alternating between using the human movement to animate itself and the agent's movement when the humans leave the area in front of the gallery.

beams later became a musical instrument as crystal spheres fell from above and created a soundscape of bells as they rebounded off the crystal structure. (Figure 9)

8. CONCLUSION

One of the goals in developing the performing agent was to borrow what we could from a human in order to rapidly develop the agent's capability. In *Recognition* we used movement and biological data from the dancer to allow the agent to develop quickly. Allowing the agent to learn from the dancer's movement which was captured while interacting with the installation gave the agent an appropriate vocabulary to use when there were no humans present to interact. Having the agent and human avatars co-existing and ready to provide movement data when necessary proved a reliable and seamless solution to the problem of the main avatar having no data to animate itself with if there were no humans present. The movement generated by the agent was visually similar to what a human would have produced.

In *Instrumental* the SOM was also used for movement recognition, the neuron containing the closest match to the human's current live movement will fire and we can act upon this recognition with appropriate events such as the movement generation performed by the agent. This utilized both the movement synthesis and recognition capabilities of the ANN to allow the agent to actively engage with the human dancer.

In *Recognition* we were able to develop a virtual performing agent that could stand in for a human in a public outdoor display, allowing the installation to have a visual continuity in the absence of human intervention. In *Instrumental* this was extended to enabling the agent to stand in as a virtual performer that could interact with the human dancer to create a collaborative improvised dance duet.

ACKNOWLEDGMENTS

The artists and authors would like to thank the following people for their much appreciated assistance:

Peter Divers and Daniel Skovli from the Deakin Motion.lab, Chee Peng Lim and Matthew Watson from the Centre for Intelligent System Research and Paul Bourke. Thanks to Cube 37 for their support in presenting Recognition.

Additional information and video is available at:

<http://www.johnmccormick.info/?s=recognition>




Figure 9. Performance, Instrumental. Left. Agent (light green avatar) and dancer (dark green avatar) The agent uses the dancer's current movement as a starting point for its own movement from which it will diverge from the dancer. Right. The crystal trail formed from the agent's pathway becomes a musical instrument, crystal spheres falling through the structure generate a soundscape of bells. Image © John McCormick 2014

9. REFERENCES

- [1] Hsu, E.; Puli, K.; Popovic, J., *Style translation for human motion*. ACM Transactions on Graphics (TOG) -Proceedings of ACM SIGGRAPH 2005 2005, 24 (3), 1082 - 1089.
- [2] Brand M, Hertzmann A, *Style Machines*. SIGGRAPH 2000; 2000 July 23-28, 2000; New Orleans, Louisiana, USA.
- [3] Taylor, G.; Hinton, G. *Factored Conditional Restricted Boltzmann Machines for Modeling Motion Style*, Proceedings of the 26th International Conference on Machine Learning,, Montreal, Canada, 2009; Montreal, Canada, 2009.
- [4] Han, S.-J.; Cho, S.-B., *Predicting User's Movement with a Combination of Self-Organizing Map and Markov Model*. Artificial Neural Networks – ICANN 2006 Lecture Notes in Computer Science 2006, 4132, 884 - 893.
- [5] Son, K. H.; Ji, C. W.; Park, Y. M.; Cui, Y.; Wang, H. Z.; Chon, T. S.; Cha, E. Y., *Recurrent Self-Organizing Map implemented to detection of temporal line-movement patterns of Lumbriculus variegatus (Oligochaeta: Lumbriculidae) in response to the treatments of heavy metal*. WIT Transaction on Biomedicine and Health 2006, Vol. 10, pp. 77-91.
- [6] Caridakis, G.; Karpouzis, K.; Drosopoulos, A.; Kollias, S., *SOMM: Self organizing Markov map for gesture recognition*. Pattern Recognition Letters 2010, 31 (1), 52-59.
- [7] Kohonen, T. 1997 *Self-organizing maps*, 2nd ed, Springer.
- [8] McCormick, J. Vincs, K. Nahavandi, S. Creighton, D. *Learning to Dance with a Human*, in Cleland, K., Fisher, L. & Harley, R. (2013) Proceedings of the 19th International Symposium of Electronic Art, ISEA2013, Sydney.
- [9] Bourke, P. 1994 *Polygonising a Scalar Field*, <http://paulbourke.net/geometry/polygonise/> (accessed December 13 2014)