Inferring Performer Skill from Aesthetic Quality Features in a Dance Game Gesture Corpus

Christopher Maraffi, Sascha Ishikawa, Arnav Jhala
University of California, Santa Cruz
[topherm, stishika, jhala]@soe.ucsc.edu

Abstract

In this paper, we describe experiments for inferring the artistic skill of performers by analyzing pose features in a gesture corpus. Poses were generated by having participants play a popular Kinect dance game in a markerless motion capture studio. Skeletal data was analyzed for features derived from both statistical analysis as well as arts and animation theory, and aesthetic metrics were designed to score pose features along three dimensions: balance, asymmetry, and readability. We applied our metrics to poses in a corpus of 10,080 annotated frames generated from 20 dance performances ranked according to the performing arts background of each participant. This work is the foundation of a computational performatology approach to quantifying artistic gesture in media by identifying aesthetic features that indicate figurative quality to viewers. The potential application of gesture analysis and feedback will be to inform the design of performative logics for virtual control of avatars and non-player characters in videogames.

Introduction

Our motivation for this study is to develop metrics that quantify aesthetic features of artistic gesture. We define artistic gesture as a performed series of non-natural poses and transition movements that are designed to maximize spectator affect, as opposed to the utilitarian nature of everyday gesture. Artistic gesture is represented by trained performers in the arts, including dancers, actors, puppeteers, and even animators. Although these figurative traditions have genre-specific techniques that express a variety of aesthetics, they also have the common goal of portraying a human body in such a way as to make the figure interesting to watch. Performing arts theorists have hypothesized the existence of "universals" for making an articulated body visually affective, sometimes associated with pre-expressive stage presence (Barba and Savarese 2005). Animators have employed general "principles" to make humanoid figures more appealing in a variety of non-interactive media, regardless of specific character or narrative content (Thomas and Johnston 1981). Though these arts practices have demonstrated success at portraying fictive characters, many of the figurative methods they employ have yet to be computationally studied to verify their operational logics, and to derive figure composition rules that would be useful for virtual character control. Optical motion capture (mocap) devices combined with analysis techniques from the field of computational aesthetics can reveal gesture patterns previously too subtle to empirically quantify. By applying such methods to the study of artistic gesture, we seek to correlate arts conventions with observable features of composed figures, to develop metrics that infer quality.

We are interested in identifying gesture features that generalize across the figurative arts, to derive composition rules that would be valid for the posing of any fictive biped character. Developing composition logics that simulate pre-expressive stage presence could serve as an aesthetic base for avatars and non-player characters (NPCs), to ultimately enhance the player experience through improved artistic quality of gestural interaction. Our experimental approach draws from recent computational aesthetic studies on image composition. Researchers have analyzed how photography rules influence viewer perception of quality, and may indicate the skill of the photographer (Joshi et al. 2011). Similar studies to quantify the aesthetic procedures of photographers have applied results to virtual cinematography applications in videogames (Swanson et al. 2012). Our study extends this work to the analysis and inference of figure composition quality in media. In our initial experiments, we collect pose mocap data from experienced and inexperienced performers playing a Kinect dance game to analyze aesthetic features associated with pre-expressive stage presence (see Figure 1). Pose metrics developed for this study will be tested against media consumer preferences in future work, and could inform the design of figure posing logics in videogames.

Figure 1: Data collection involved participants playing a popular Kinect dance game on a markerless motion capture stage.

We are working under the assumption that the convergence of natural user interface (NUI) devices with cinematic aesthetics in videogames will require more robust character logics and player affordances based on artistic gesture practices (Maraffi and Jhala 2013). Intelligent virtual characters modeled on how artists portray the body will need metrics that infer the most aesthetically interesting poses in relation to the virtual camera viewpoint, but so far no such tools exist. Computational aesthetics provides a phenomenological approach to understanding what arts practitioners do, by combining top-down feature selection using arts theory with bottom-up statistical analysis of performer data. Before we can model NPCs that generate better figurative aesthetics, we first have to quantify what spectators mean when they identify an artistically posed body as a "good" or skilled performer.

Related Work

While there has been some computational work done on what actors do in the performing arts (Bates 1993, Seif El-Nasr 2004, Marsella 2006, Magerko 2010), such studies have not analyzed aesthetic features in figurative data to quantify general artistic skill associated with stage presence. Our study extends previous image composition work in computational cinematography to the related domain of figure pose quality. In media, cinematography and embodied performance share a common goal of sustaining viewer attention, and we hypothesize that similar rules for enhancing visual interest inform composition procedures in both. Physical performers must be constantly aware of how the lines of their body will read when viewed by a live audience or camera lens, and are trained to improvise real-time adjustments of their body positions to maximize visual affect. Similarly, animators employ figure principles that increase readability and appeal, which were originally derived from studying live actors (Thomas and Johnston 1985). In both cases, viewer interest is enhanced by the practitioner's real-time skill in exercising learned artistic conventions related to posing a figure. Since we are isolating "how" artists improvise character content beyond authorial constraints specified in the script or choreography, we employ Barba's concepts of pre-expressivity, stage presence, and developing a fictive body (Barba and Savarese 2005).

Barba's practice of inter-cultural theatre anthropology was intended to empirically study classical actors from a variety of eastern and western traditions to explain the "anatomical science of performer bios" (Barba 2003). He described the phenomenon of presence as "extra-daily tensions" that trained performers display to give life to a fictive body independent of the formal score, and which make spectators want to watch the characterization. Because our study requires more specificity than what performing arts theory provides, we also reference animation principles that were designed to give the "illusion of life" to any rendered figure (Thomas and Johnston 1985). Like puppeteers, animators are distanced from the fictive bodies they control, and through this abstraction reveal some of the mechanisms behind their craft. For instance, animators do not have to feel angry to render a convincingly mad character; they just have to exercise skill in artistic conventions for portraying anger in a series of poses. The practice of key-frame animation is consistent with classical performing arts practices like puppetry and external acting methods that adopt a character through learned postures. Since a full description of our performatology approach is beyond the scope of this paper, please see the authors' previous technical report (Maraffi and Jhala 2013). We consider our approach to be complementary to other computational work on choreography tools and method acting (Calvert et al. 2005, Tanenbaum 2011), because pre-expressive aesthetics and expressive content are layered together in practice.

Drawing from Computational Aesthetics

A tutorial on inference of quality in images provides a good overview of the emerging field of computational aesthetics (Joshi et al. 2011). Methodologies for this area of study have been influenced by psychology and visual art theory related to the aesthetics of image composition. Artistic conventions have been studied in the popular arts of painting, photography, and cinema to derive procedural rules for quantifying image quality. In several studies, Rule of Thirds was implemented as a quality-inferring feature used to distinguish professional from amateur photography. The Gestalt psychology concept of goodness configuration, where perception is organized according to properties like symmetry and simplicity, also influenced feature selection. Processing fluency theory suggests readability features are related to image appeal. Challenges in this field include an "aesthetics gap" that stems from the inherent semantics gap between low-level computable visual features and high-level subjective semantics. A core problem is in understanding individual preferences that make some images more appealing than others.

Drawing from Computational Cinematography

Our current study closely follows previous work done in computational cinematography that quantified the aesthetic features of images generated through a photography videogame (Swanson et al. 2012). The top-down feature selection of image composition rules was combined with bottom-up statistical analyses to correlate aesthetic quality with viewer preferences. The game scored and annotated the player's photos according to aesthetic rules derived from photography texts. Real-time visual feedback functioned as an aesthetic meter, encouraging players to take pictures that rated highly for Rule of Thirds, balance, and symmetry. Using a videogame to generate a corpus of images isolated the data set in a way that would be difficult using photos from image repositories on the Web, and gave them control over feature dimensions through design abstraction, such as minimizing cultural references in the game's representation of landscape photography. Subsequent preference studies on the viewer perception of image quality were conducted through crowdsourced ratings via Mechanical Turk. Workers classified images according to quality by using pairwise comparisons of four-alternative forced choice (4AFC) ratings to learn individual as well as general preferences. They reported relatively accurate predictions for an arbitrary user, and even better accuracy with less training examples for individually learned preferences.

Feature Design

Initially our gesture metrics focus on pose features in a single frame of mocap data, rather than features of continuous movement. Constraining our study to poses reduces feature dimensions and aligns well with aesthetic studies on image composition for still photography. In addition, cinematography conventions decouple scenic composition from transitions, which has enabled the application of image composition studies to virtual camera control in time-based media. We anticipate similar advantages to decoupling figure composition from transition movements in the study of performer data for virtual character control. Animation practices also emphasize the hierarchical difference between poses and transitions in the concept of key-frames, the primary poses designed by lead animators to convey the most meaning, with transition or in-between frames being a secondary consideration. Our pose feature selection is derived from performing arts theory on pre-expressive stage presence and animation principles: balance, asymmetry, and readability (BAR features see Figure 2).

Photography conventions like Rule of Thirds are utilized in a variety of professional fields that require image composition, including photography and graphic design that feature figures, implying structural rules that may generalize well. Other rules used to infer aesthetic quality in images, such as balance and symmetry, have comparable principles in embodied art practices like acting and dance. We suspect that composition rules in graphic and performing arts overlap because the artists employing them have the unified goal of catching and holding a viewer's attention through visual interest. Good composition through the arrangement of shapes and lines stimulates eye movement through the framed image space, which reflects visual processing theories from psychology (Reber et al. 2004, Peters 2007). In studying features associated with pose quality, we hope to better understand how an articulated body can engage viewer attention in order to establish a communication channel for delivering character and narrative content. Our assumption is that skilled players trained in gestural performance conventions will consistently score higher for our metrics than unskilled players, and therefore will produce poses that rank higher for quality by the general public.

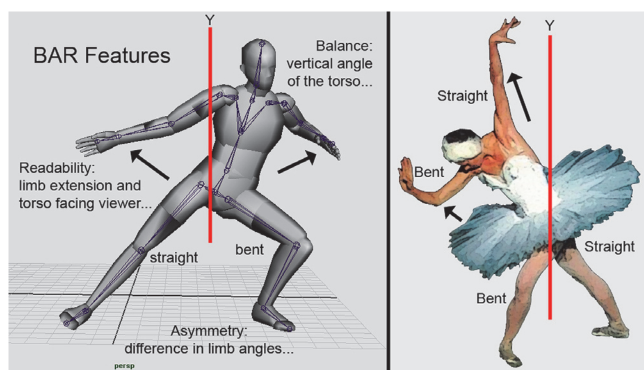

Figure 2:BAR features of balance, asymmetry, and readability in a pose from our corpus, and in a ballet dancer pose (Barba and Savarese 2005).

Balance

Trained performing artists develop skill at pushing the limits of their balance, which is the first BAR feature. Our metric is designed to score higher for what Barba calls "precarious balance" (Barba and Savarese 2005), or how close the figure's torso approaches imbalance. Poses that push the limits of balance by articulating the body into an aesthetically appealing "S" shape are considered more dynamic by performing arts theorists. To achieve poses with precarious balance, performers train at shifting their hips and shoulders in opposite directions, which counters the weight of the upper and lower body. Our metric measures the distance between torso sections to calculate the steepness of the curve in a standing position. Our balance metric has similarities to metrics in aesthetic studies that feature photo composition rules like Rule of Thirds and the Golden Triangle (Khan 2012, Swanson 2012), in that the head and hips of the figure are moved off the center line to power points in the thirds region. The assumption we are testing is that skilled performers will bend their torso off the vertical axis to more extreme angles than unskilled performers.

Asymmetry

While the balance metric measures how dynamically a figure's torso is posed, asymmetry is intended to test for dynamics of limb composition. Mirroring limb positions across the body's center-line is considered static and boring in figurative arts theory. An asymmetrically posed body has variations in limb rotations. For instance, one limb may be straight while the other is bent. Asymmetrical limbs create structural lines that may encourage spectator eye movement, increasing visual stimulation as consistent with visual processing theory (Reber et al. 2004). Recent psychology studies on symmetry and facial beauty have shown that functional asymmetries in the natural faces of models scored higher for attractiveness than digitally mirrored faces (Zaidel 2005), so we hypothesize that body asymmetries may also appeal to viewers. Symmetrical poses like the default T-pose used to model 3D characters may be highly readable because the same content is reflected on each side of the body, but asymmetrical poses are information rich and imply a higher level of limb control by the performer. Our assumption is that skilled performers will score higher for asymmetrical poses than unskilled performers.

Readability

Our third BAR feature is how well a pose can be read by a spectator. Performing artists are trained to be aware of where the viewer or camera is in relation to their bodies, and adjust their torso and limbs to produce a readable pose from a distance and in a variety of lighting conditions. Professional animators are taught to minimize occlusions by posing a character's limbs away from the torso, and to check readability by distancing themselves from a figure and blurring their eyes to see only the silhouette. Readability may influence how expressive a pose is by how well it is perceived by a viewer (Thomas and Johnston 1985). Psychology studies have shown that how easily an image can be cognitively processed is related to aesthetic pleasure (Reber et al. 2004). We hypothesize that if a pose is unreadable, it cannot convey enough semantic information to affect spectators. Part of traditional figurative arts training is to exaggerate natural gestures to more clearly articulate the performers intentions. Our metric measures how close limbs are to the body, to determine how much of the torso is visible from the spectator's point of view. The assumption is that trained performing artists will have more skill in articulating their body readability for a camera than unskilled performers.

Metrics

We made the assumption that poses with a greater sense of precarious balance, that is, shifting one's weight without falling, will be more difficult to execute and have greater aesthetic value. As a result, our balance metric was aimed towards capturing this intuitive notion by estimating the amount of "lean" in the hip, shoulder, and torso. The hip and shoulder lean are calculated by drawing a line between the left and right hip/shoulder and calculating its angle from horizontal. The torso lean is obtained in a similar fashion by drawing a line between the center of the hip to the base of the neck and calculating its angle to vertical. Figure 3A illustrates this.

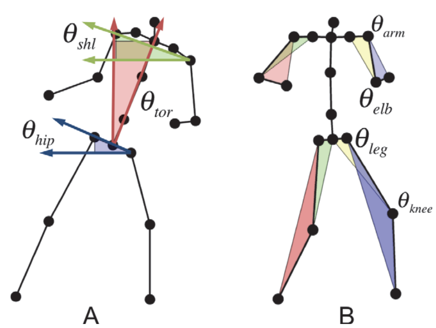

Figure 3: Illustration of the angles used for Balance (A) and Asymmetry (B) features.

The asymmetry metric captures the differences between the left and right sides of the body by considering (1) the angles formed at the leg, knee, arm, and elbow joints (Figure 3B), and (2) the orthogonal rotation of joints along the three spatial dimensions. For the latter, we consolidate the rotational angles into a single measure of rotation by treating it as a vector magnitude. The overall asymmetry of the joint rotations was obtained by calculating the difference in rotation between the left and right joints. Let $\theta_{h,a}$ denote the rotation of the joint, where the subscripts h and a specify the left/right side of the body and the axis along which the angle is measured. Then the rotational asymmetry of the j-th pair of joints is given by

$$a_r^{(j)} = \sqrt{\left(\theta_{L,x}^{(j)} - \theta_{R,x}^{(j)}\right)^2 + \left(\theta_{L,y}^{(j)} - \theta_{R,y}^{(j)}\right)^2 + \left(\theta_{L,z}^{(j)} - \theta_{R,z}^{(j)}\right)^2}$$

The rotational asymmetry of the pose is given by

$$A_r = \frac{1}{N_j}\sum_{j=1}^{N_j} a_j$$

where N_j is the total number of joint pairs. Finally, the rotational asymmetry and the joint asymmetry are summed and normalized over the total range of degrees they can vary.

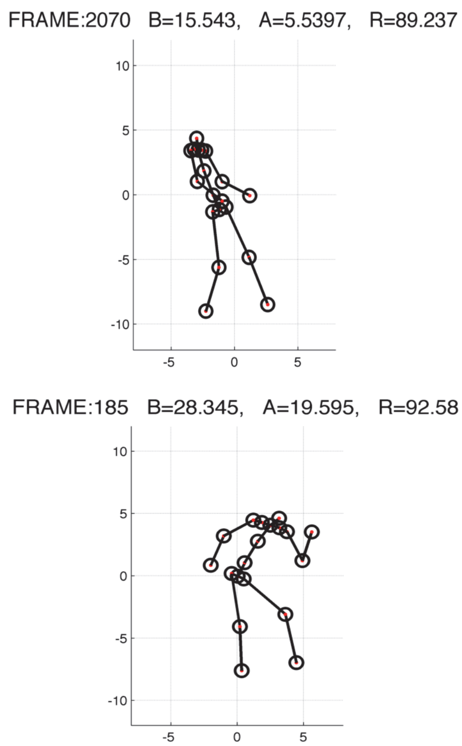

Figure 4: Application of our metrics to score poses for BAR features. Note the readability in the top pose is lower due to occlusion whereas asymmetry in the bottom pose is higher.

Figure 4 shows two poses with different aesthetic qualities. Note that the BAR features provide a qualitative distinction between the precarious balance, asymmetry, and readability. For example, in the top pose it is harder to distinguish the arms from the torso (low readability), which is reflected in the lower readability measure relative to the pose shown below, which is much easier to read.

The readability of the pose depends, to some extent, on the location of the viewer relative to the camera. Therefore, a component of readability can be represented by measuring the amount of foreshortening that occurs in the projection of the skeletal features onto the camera plane. Since the motion capture system normalizes the skeletal features of each performer to standard lengths, an estimate of readability can be obtained by calculating the ratio of foreshortened lengths to the actual lengths of the skeletal features.

Experiment

After an initial pilot study, we applied our metrics to poses in a gesture corpus generated through playing a Kinect dance game in our mocap studio. Dance-themed videogames have been designed to be accessible to players from a variety of backgrounds, not requiring specialized skills to play. Dance is also popular with both amateurs and professionals, and is the most abstract of the performing arts, so we can focus on structural rules that do not require a narrative component. Participants played the Michael Jackson Experience (MJE) Kinect dance game. We chose MJE because the songs and style of dance are known to a large section of the population, and the game is designed to feature several iconic Michael Jackson poses.

Pilot Study

An initial qualitative assessment was made in order to determine the efficacy of using the BAR features to characterize human poses. BAR features for three different motion sequences, namely the "lateral stretch," full-body rotation, and the "arm wave," were inspected manually. These are compared to the baseline pose of simply standing while facing the camera. Figure 5 plots each BAR feature throughout the sequence of frames. The motion sequences examined are highlighted as vertical bands on the left, middle, and right of the figure.

We begin by establishing a baseline for the BAR features during a simple standing pose: the actor stands in a relaxed position, with limited amount of motion, while facing the camera. This is captured by the frames 1400-1600 (no highlight) in Figure 5.

Frames 190-335 (highlighted far left) correspond to movements resembling the "lateral stretch" where the actor stands facing the camera while the torso sways laterally. The actor sways three times, which is captured by the periodic motion in the balance and asymmetry plots.

Frames 675-850 (highlighted middle) capture the full-body rotation, which produces corresponding oscillations in the readability plot. As expected, high readability values coincide with frontal views whereas lower values coincide with side views where the joints are heavily occluded.

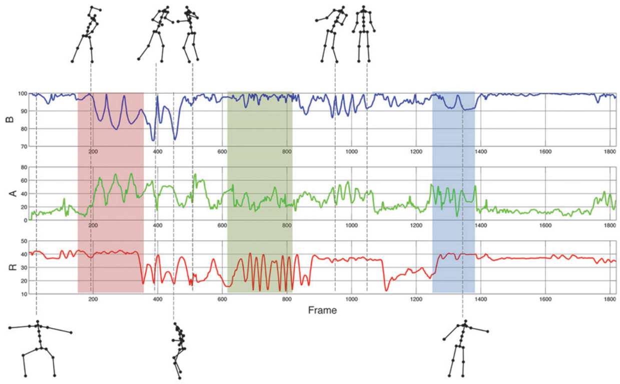

Figure 5: Application of our metrics to score poses for BAR features in an initial pilot study.

Frames 1290-1400 (highlighted far right) capture the "arm wave." In this motion sequence the actor stands facing the camera and begins a wavelike motion in one arm that travels across to the other arm, then back again. The actor leans into the direction of the wave, which produces a small, but steady, change in balance that is clearly visible as a "spike" in the balance plot.

We used the C4.5 decision tree implementation in the Weka machine learning library (Hall at al. 2009) to learn the participant's preferences given binary ratings for each pose along with its BAR features. Using a 10-fold cross-validation procedure, the resulting decision tree was able to rate "unseen" poses with an accuracy of 79.7 percent. The results are summarized in the following confusion matrix in Table 1. Note that the decision tree that was generated supports our hypotheses that "good" poses will have higher measures of precarious balance, asymmetry, and readability.

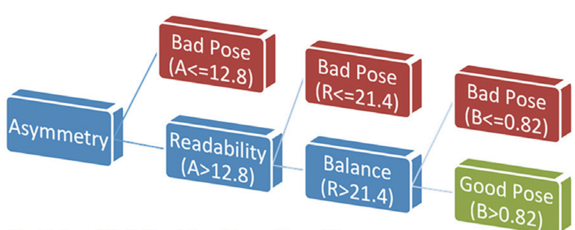

Results Using C4.5 Decision Tree Algorithm

	Good Pose	Bad Pose
Good Pose	45	17
Bad Pose	9	57

79.7% Correct

Table 1: Summary of our pilot study results.

Data Collection

Mocap data was collected from 20 performances, composed of male and female participants ages 18-47 who were classified based on their survey responses as 10 skilled and 10 unskilled in the performing arts. Classification was done through using a 4-choice self survey to determine how much formal training each participant had in performing arts practices like dance, combined with their performance as judged by the game scoring and an expert dancer. Participants were first taken through a warm-up and training phase on the MJE videogame. Then they danced the MJE choreography for the popular song Beat It twice through, once as a practice round to learn the choreography, followed by a final round where we recorded their data and high game score. We captured their gesture playing the game with a markerless Organic Motion system running at 120fps, streaming directly into Autodesk Motion Builder software. The captured data was retargeted to a normalized skeleton in Motion Builder and plotted to 6,059 individual frames, and imported to Autodesk Maya for final processing and rendering. We normalized the final captures to align with each other within a 10 frame threshold, and sampled the global translation and local rotation values for all the major skeletal joints to create a corpus of 10,080 annotated frames, ignoring in-between frames that contained transition movements.

Data Analysis of BAR Features

A total of 56 key-frames were identified and analyzed along with their 10-frame neighborhood to accommodate the likely scenario that the subjects did not arrive at their choreographed poses at exactly the same time. We performed principal components analysis on the data in order to visualize and reduce its dimensionality. The top-three principal components, or PCs, accounted for 89.1 percent of the overall variance in the data (51.5, 26.9, and 10.7 percent, respectively). Figure 6 shows a plot of the pose data in a two-dimensional space consisting of the top-two (PCs). There is reasonable separation between frames corresponding to skilled (green) and unskilled (red) performers, suggesting that there may exist some linear combinations of the 114 pose parameters that are capable of distinguishing between skilled and unskilled poses. The top-ten PCs were used, collectively representing 99.6 percent of the overall data variance. We employed a C4.5 decision tree classifier that was trained on data from half of the performers, of which an equal number were either skilled or unskilled. The decision tree was able to predict the skill level of the remaining (unseen) performers' poses with up to 81.3 percent accuracy (see Table 2). Note that the algorithm is better at classifying skilled performers than unskilled ones.

	Skilled	Unskilled
Skilled	2872	208
Unskilled	946	2134

Table 2: Classification results of trained decision tree using C4.5 algorithm on unseen frames.

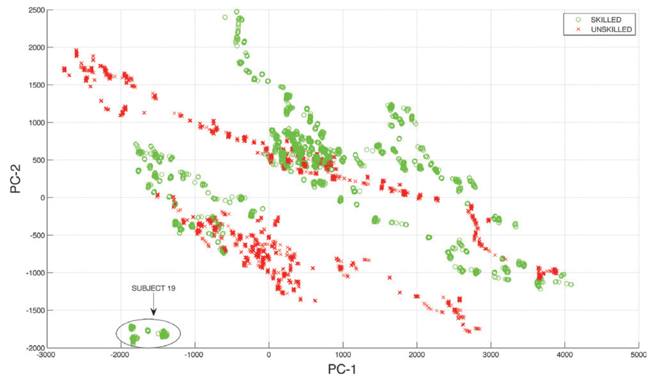

Figure 6: Distribution of selected key-frame poses in principal component feature space.

To further investigate the physical interpretation of each principal component, we counted the number of times each parameter was used as a term in the linear transformation of the top 10 PCs (see Figure 7). These were determined, in order of decreasing frequency, as `legUpperRight_Py`, `legUpperRight_Pz`, `legLowerLeft_Ry`, `footRight_Pz`, `footRight_Py`. These findings suggest that, given the choreography used in the study, the positional and rotational parameters of the legs and feet are the most diagnostic of skill.

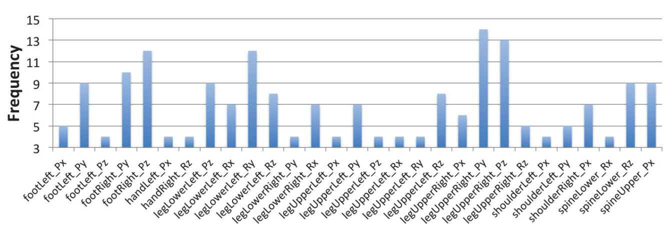

Figure 7: Frequency of pose parameters used in top-10 PC transformations.

We also checked to see if the in-game score was indicative of self-reported skill level of players, which seemed likely. For instance, subject 19 had poses clustering in the lower right portion of the PCA plot (Figure 6). The subject was self-ranked to be highly skilled, and received the 3rd highest score according to the videogame. To check this, we labeled players with skill ratings from the questionnaire, and then looked at scores across both groups. The scores in the skilled group ranged from 110K to 144K, and in the unskilled group ranged from 84K to 128K. To determine if these scores came from different distributions, we ran a two-tailed test. The test showed that performers who were considered "unskilled" and "skilled" (according to our interpretation of their survey) had scores from two distinct distributions (p-value = 0.0016 shown in Figure 8).

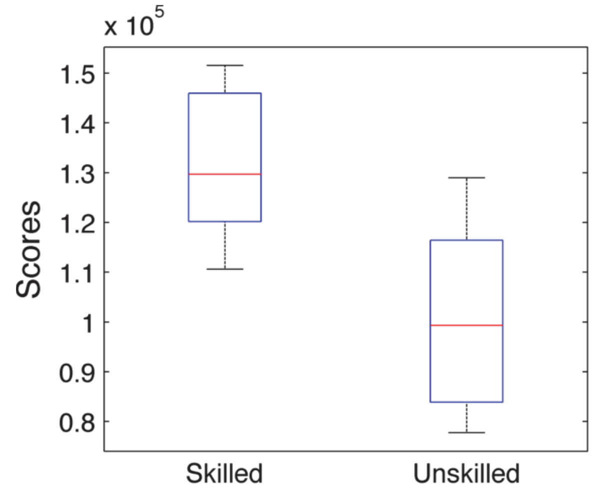

Figure 8: Distribution of in-game scores for players.

Discussion and Future Work

Though we were able to get enough pose data from dancers playing MJE in our mocap studio, we ran into some inherent limitations for our study using an off-the-shelf Kinect game. We were constrained to poses required for the MJE Beat It choreography, and the fast movements required to play a dance game sometimes resulted in noisy data and limb flipping. In addition to ongoing studies to quantify viewer preferences of artistic gesture, future work includes implementing our own pose generation Kinect game that will feature a real-time aesthetic meter that scores for our metrics. This will allow us to generate a pose corpus from gameplay that is annotated for our BAR feature set, which can then be used for crowd sourced preference studies. Lastly, we intend to implement experiments that explicitly measure the aesthetic connection of cinematography to artistic gesture. In real-world media production, performer and camera composition often happen simultaneously through real-time collaboration. Automatically syncing virtual character posing with camera framing would theoretically produce the most aesthetically appealing visuals for character interaction in videogames.

Modeling artistic gesture through computational aesthetics can potentially transform the creative quality and personalization level of NPCs and avatars. Most research on procedural content generation techniques focuses on the analysis of algorithms in terms of functional value of artifacts rather than a match to the player's aesthetic preferences. Since experienced actors and animators are known to generate an aesthetic affect in viewers, studying and quantifying their embodied representations holds potential for simulating artistic skill in virtual characters that will also generate viewer affect.

References

Barba E. 2003. The Ripe Action. In Mime Journal '02-'03.

Barba, E., and Savarese, N. 2005. A Dictionary of Theatre Anthropology: The Secret Art of the Performer. New York, NY: Routledge.

Bates, J. et al. 1993. Dramatic Presence. Presence'93.

Calvert, T., et al. 2005. Applications of Computers to Dance. Computer Graphics and Applications, IEEE , vol.25, no.2, March-April'05.

Hall, M., et al. 2009. The WEKA Data Mining Software: An Update; SIGKDD Explorations, Volume 11, Issue 1.

Joshi, D., et al. 2011. Aesthetics and Emotions in Images: A Computational Perspective. Tutorial in IEEE Signal Processing Magazine, Sept'11.

Khan, S. and Vogel, D. 2012. Evaluating Visual Aesthetics in Photographic Portraiture. In Computational Aesthetics in Graphics, Visualization and Imaging '12.

Magerko, B, et al. 2010. Bottoms Up: Improvisational Mico-Agents, In 3rd Workshop on Intelligent Narrative Technologies, FDG' 10.

Maraffi, C, and Jhala, A. 2013. Raising the Aesthetic Quality of Character Interaction in Cinematic Videogames, Technical Report, UCSC-SOE-13-01, UCSC.

Marsella, S. et al. (2006). An Exploration of Delsarte's Structural Acting System. Proceedings of IVA'06.

Peters, Gabriele. 2007. Aesthetic Primitives of Images for Visualization. In Proceedings of IEEE'07.

Reber, R., et al. 2004. Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience? In Persp. Social Psychol. Review, vol8, no4.

Seif El-Nasr, Magy. 2004. A user-centric adaptive story architecture: borrowing from acting theories, In ACE '04 Proceedings of the 2004 ACM SIGCHI.

Swanson, R., et al. 2012. Learning Visual Composition Preferences from an Annotated Corpus Generated through Gameplay. In Proceedings of IEEE'12.

Tanenbaum, J. 2011. Being in the story: readerly pleasure, acting theory, and performing a role. In Interactive Storytelling '11.

Thomas, F. and Johnston, O. 1981. *Disney Animation: The Illusion of Life.* New York: Disney Editions.

Zaidel, D. and Cohen, J. 2005. The Face Beauty, and Symmetry: Perceiving Asymmetry in Beautiful Faces. In Intl. Journal of Neuroscience, vol115, no8.