Clone Attack! Perception of Crowd Variety

Rachel McDonnell Michéal Larkin Simon Dobbyn Steven Collins Carol O'Sullivan

*Graphics, Vision and Visualization Group, Trinity College Dublin.

Figure 1: *Example of a crowd used in the Appearance Variation Experiment with the maximum number of clones.*

Abstract

When simulating large crowds, it is inevitable that the models and motions of many virtual characters will be cloned. However, the perceptual impact of this trade-off has never been studied. In this paper, we consider the ways in which an impression of variety can be created and the perceptual consequences of certain design choices. In a series of experiments designed to test people's perception of variety in crowds, we found that clones of appearance are far easier to detect than motion clones. Furthermore, we established that cloned models can be masked by color variation, random orientation, and motion. Conversely, the perception of cloned motions remains unaffected by the model on which they are displayed. Other factors that influence the ability to detect clones were examined, such as proximity, model type and characteristic motion. Our results provide novel insights and useful thresholds that will assist in creating more realistic, heterogeneous crowds.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—[Animation]

Keywords: perception, variety, crowds, animation

1 Introduction

Realistic virtual environments depicting thousands of virtual humans can be challenging to create due to a number of limiting factors. For data-driven crowd systems, assets such as skinned human models and motion captured character animations can be expensive to purchase or commission, while system resources are a further constraint. Therefore, a fixed number of template characters are usually deployed to generate large crowds using instancing, thus producing *Appearance Clones (A_C)* that can often be easily noticed. Character motions also consume memory resources, which data compression can help to reduce [Arikan 2006]. As with the character models, the same set of animations is typically used multiple times, resulting in *Motion Clones (M_C)* that can give the disturbing impression of a crowd of people all moving identically.

For these reasons, some researchers have developed approaches aimed at increasing the visual variety of humans in a simulated crowd (e.g., [Tecchia et al. 2002; Ulicny and Thalmann 2002]). While compelling results have been achieved, the factors that affect the *perception* of variety in crowds have not been evaluated to date. Such information is however essential to allow for effective trade-offs between realism and resource wage, by ensuring optimal variety. Perhaps it is the case that a small number of walking motions applied to all individuals in a pedestrian crowd would be acceptable as long as their appearance is different, whereas it is unlikely that different motions applied to a single character will result in a heterogeneous crowd.

Using a library of twenty template models and motion captured gaits, we performed two sets of perceptual experiments. In our *Baseline Experiments*, we investigated the factors that affect people's ability to identify a single pair of clones amongst a number of characters or motions. In the *Multiple Clone Experiments* we examined the perception of variety when the number of a character's appearance or motion clones (i.e., its *multiplicity*) is increased.

Summary of effects:

- Appearance clones were easier to detect than motion clones
- Increasing clone multiplicity reduced variety significantly
- No appearance model was more easily detected than others
- Certain gaits were more distinctive than others
- Color modulation and spatial separation effectively masked appearance clones
- Combined appearance/motion clones were only harder to find than static appearance clones when their cloned motions were out-of-step
- Appearance clones were also harder to find when combined with random motions
- Motion clones were not affected at all by appearance, even with random appearances

*http://gv2.cs.tcd.ie/

Furthermore, from our results we derived some useful thresholds that could be used to balance the variety/resource trade-off mentioned above. While we concentrate in particular on pedestrian-like crowds in this paper, our results should also provide insights for many other types of crowds (e.g., stadium crowds or armies). Our results will be mostly useful for developers of games and other interactive applications. Furthermore, many of the issues raised are also applicable to non-realtime applications, where production deadlines may necessitate rapid prototyping of crowds.

2 Related work

In previous work, computer animation researchers have investigated the perception of animated human motion and have developed perceptual metrics for improving the quality of the animations (e.g., [Hodgins et al. 1998; Reitsma and Pollard 2003; Harrison et al. 2004; Wang and Bodenheimer 2004]). More specifically, perceptual metrics have also been used for designing crowds, with the focus on different level of detail representations and how effective they are at replicating the appearance and motion of crowd characters [Hamill et al. 2005; McDonnell et al. 2005; McDonnell et al. 2006].

Perception researchers have analyzed motion independent of appearance. Johansson [1973] was the first to introduce a stimulus called a "point light walker" to the community, where lights were attached to the joints of humans and their motion was analyzed with only the light sources visible. Furthering Johansson's work, Cutting [1977] found that by using motion cues alone, people are able to recognize individuals based only on their walks. Beardsworth [1981] found that even one's own walking pattern could be identified using this method. This shows that different individuals can have perceptibly different gaits, which implies that unique walks may be needed in order to create a realistic and varied pedestrian crowd.

There have been some recent advances in modifying the appearance of crowd clones so that they all appear different, by using accessories and texture modulation [Thalmann et al. 2007]. However, the simplest and least resource intensive way to add variety to cloned template meshes is to use hardware accelerated per body part color modulation. This is a very popular and effective technique used in the crowd research community [Tecchia et al. 2002; Gosselin et al. 2005; de Heras Ciechomski et al. 2005a; Dobbyn et al. 2005; Dobbyn et al. 2006; Maïm et al. 2007]. Color variety is based on texture color modulation, i.e., modulating the colors of the pixels on the model's texture map using look-up tables. The alpha channel of the human's texture map is manually encoded with different values indicating which body-part regions are to be colored differently. The fragment shader is then used to determine for each pixel which part of the body it is associated with and then color it appropriately. This can be done very efficiently in real-time using graphics hardware in a single rendering pass.

3 Assets and Framework

In order to analyze variety in the motion and appearance of crowds, a set of individual motions and a set of template models were required. In realtime crowd applications, typically between 3 and 10 templates are used. For example, the systems in Figure 3 use 4 and 7 templates respectively [Thalmann et al. 2007]. Furthermore, it is within the foreground characters of a crowd that lack of variety is most noticeable. Depending on viewpoint and crowd density, occlusions and perspective foreshortening mean that only a few characters are close enough for their details to be clearly visible. This is especially true for systems that show the environment from a first person perspective. Therefore, we chose twenty models as a reasonably conservative estimate of the number of foreground characters clearly visible, which also exceeds the number of templates used in most real-time applications.

3.1 Models

We acquired a set of 20 commercially available models that represented a variety of pedestrian types: six female and fourteen males; aged young, middle-aged or elderly; wearing both formal and casual attire. The diffuse texture maps for these models were photographed from real people, so their appearance was very natural. Each model had a single texture map which incorporated all of the textures for their clothing, skin and hair (Figure 2(a)). Similar to the technique used by De Heras Ciechomski et al. [2005b], we created an alpha map, which encoded each region with a unique greyscale value. At run-time, these values were then used to index an *outfit* map, that stored the amount of HSV color modulation necessary for each region. We manually created 32 unique outfits for each model with varying hair, skin, clothing, and shoe colors (Figure 2(b)).

In [McDonnell et al. 2007], we showed that the sex of a walker could be identified using a mannequin model, which was itself judged by all participants to be androgynous. We chose to use this representation in cases where we wished to analyze motion alone without a distracting appearance. This choice was made over point light walkers [Johansson 1973], since spatial cues are important in our target applications.

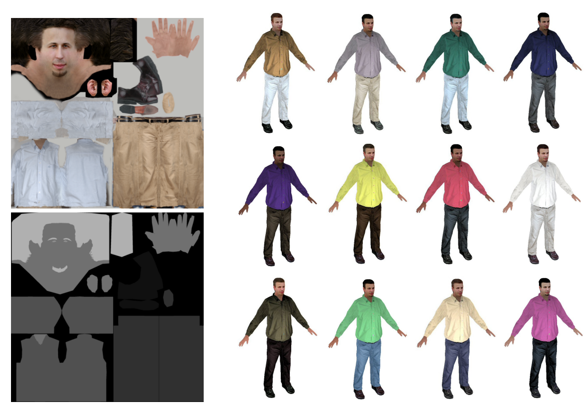

Figure 2: *(a) Texture map and Alpha map images, (b) Example of the HSV shader being used to create 9 outfits.*

3.2 Motions

Since people have unique walks in reality, we motion captured twenty different walks in order to have sufficient motion variety. We captured the motions of 14 male and 6 female volunteers of varying ages and body types - one for each of the models described above - using a 10 camera optical tracking system. Volunteers were not informed as to the purpose of the motion capture session. A straight line path was drawn on the floor in the capture area and they were asked to walk along this path up to 25 times. We captured a number of walks, but did not tell the volunteers when we were capturing them, since we wanted them to walk as naturally as possible. Three cycles of the walk were selected and extended to a longer motion clip by repeating the walk a number of times and blending between the transitions without artifacts. Another effective way of piecing together motion capture clips would be to use motion graphs [Kovar et al. 2002; Lee et al. 2002].

Figure 3: *Examples of crowds where the foreground characters are most visible [Thalmann et al. 2007].*

3.3 Framework

In the Baseline Experiments, characters were displayed in an orthographic matrix format - see Figures 4 and 6. In the Multiple Clone Experiments, all 20 characters were displayed without shadows on a grey ground plane in such as way as to optimize their visibility and facilitate picking - see Figure 1. The experimental crowd system was developed using an open-source renderer that used DirectX 9.0. HSV color modulation was implemented as a Higher Level Shading Language (HLSL) program. The experiments were run on a workstation with 2GB of RAM, an 8-series GeForce graphics card on a wide-screen LCD monitor. Participants sat at a distance of 57cm from the display and they were instructed to maintain this throughout the experiment.

A mouse was used to control an onscreen cursor and participants clicked to select characters, whereupon arrows appeared above the characters' heads. If an incorrect pair was chosen, the arrows remained above their heads for a second and then disappeared, the scene did not change and this was recorded as a false positive. If a correct pair was chosen, the scene changed and the reaction time was recorded as the time from the start of the trial until clicking on the first of the correct pair. If the participant failed to complete the task within the allotted time, the full time was recorded as their reaction time. Between trials a fixation cross appeared in the center of the screen, upon which participants fixated until the next crowd scene was loaded, thus ensuring that they fixated on the same screen position for every new presentation.

4 Baseline Experiments

We wished to provide a baseline analysis of the factors affecting people's ability to identify a single pair of clones amongst our library of appearances and motions. In particular, we wanted to gain insights into whether certain models or motions had characteristics that made it particularly easy to detect their clones.

4.1 Appearance Baseline

Does color variation help in disguising an appearance clone? Are some model types more distinctive than others? To find answers to these questions, 15 naïve participants from different educational backgrounds (11M-4F) took part in an experiment. They were first informed what an appearance clone was (using an example of a model not used in the experiment). Then, three rows of four models were shown onscreen (Figure 4), two of the twelve slots contained the same character and participants were asked to click on them as quickly as possible. A maximum of 30 seconds was allowed for participants to make their choice.

Figure 4: *Example of the exact clone condition in the Appearance Baseline Experiment.*

We used a two-way, repeated measures design where the conditions were *model type* (20) and *clone type* (2). We separated the clone type condition into two separate blocks: block one showed a single clone with no color modulation, while block two showed a single clone with color modulation. Three repetitions of each condition were shown resulting in a total of 120 trials: 20 model types * 2 clone types * 3 repetitions. The 120 trials were viewed in random order by each participant.

4.1.1 Results

We averaged participants' reaction times over the three repetitions for each of the models, for the two clone types (exact or color modulated). A two-factor ANalysis Of VAriance (ANOVA) showed that there was a main effect of clone type ($F_{1,12} = 103, p < 0.0001$). The mean reaction time for identification of exact clones was 5.7 seconds, whereas for color modulated clones it was 12.3 seconds. This answers our first question, in that *the addition of color modulation significantly masked the appearance of cloned models*. We therefore decided to use color modulated clones in the Multiple Clone Experiments.

A main effect of model type was also found ($F_{19,228} = 1.9, p < 0.02$) and an interaction between model and clone type ($F_{19,228} = 2.4, p < 0.002$). Post-hoc analysis was then performed using a standard Newman-Keuls test to further investigate the differences in reaction times between the models. It was found that, for the exact clone condition, *there was no significant difference between reaction times for any of the models*. However, in the color clone condition, one model was detected significantly less often than some others, which was probably due to the outfits of that model being particularly different from each other. This implies that, notwithstanding the main effect of model type detected, this was not due to differences in model shapes. As a result, we improved some of the remaining outfit maps manually in order to ensure that they were sufficiently different for the following experiments.

Because of the matrix configuration, we were also able to analyze the effect of position of the clone pair. For each participant, we split their reaction times for distances along three axes: horizontal, vertical and diagonal. Each of these had a value for near and far; near being the average of all combinations that were one space away and far being the average of all combinations that were more than one model away from each other. A two factor ANOVA was conducted on this data where the conditions were *Axis* (3) and *Distance* (2). We found a main effect of axis ($F_{2,24} = 5.59, p < 0.02$), where horizontal pairs were identified most quickly. Vertical pairs were identified second quickest and diagonal pairs took the longest to identify. A main effect of distance was also found ($F_{1,12} = 54.13, p < 0.0001$), where near models were spotted on average in 7 seconds and far in 10 seconds. There was also an interaction between distance and axis ($F_{2,24} = 5.59, p < 0.02$). Here we can see that once models are far away, the axis has no effect, but the near models are seen more quickly when located horizontally and vertically than on a diagonal (Figure 5). Although we did not further investigate this factor, we felt that this may be of interest to developers of certain types of crowds, such as in a stadium.

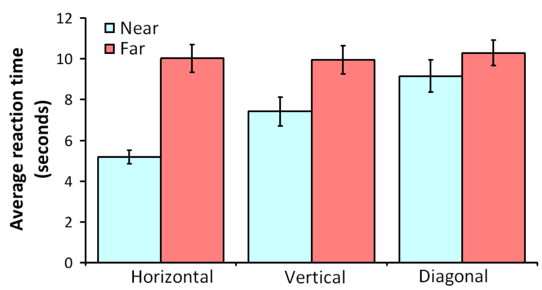

Figure 5: *Interaction between Axis and Distance.*

4.2 Motion Baseline

Are similar motions harder to find than similar appearances? Are certain gaits more individual than others? In this experiment, appearance was kept constant while gait was varied. As in the Appearance Baseline Experiment, initially three rows of four characters were shown onscreen - one of the gaits was cloned and participants were asked to click on both as quickly as possible. All other gaits displayed were unique and chosen randomly from the 20 motion captured walks. The cloned gaits were displayed in-step and all other motions were randomly out-of-step. However, the first two participants failed to find almost all pairs in the time allocated (60 seconds), leading to a duration of almost one hour each. Rather than increasing the exposure time, which would make the experiment unacceptably long and also change the nature of the task, we decreased the number of models onscreen to three rows of three characters. This proved to be equally difficult so we finally decreased the number to just six onscreen (Figure 6). Even with this simplification, participants still found the task of finding cloned motions out of six characters difficult, but the reaction times were within the time allotted of 60 seconds. Nine naïve participants (7M-2F) performed this version of the experiment.

A single factor (i.e., gait) design was used and each of the 20 gaits was cloned for 3 repetitions, resulting in a total of 60 trials. The clones were randomly placed in the scene for each repetition and the 60 trials were viewed in random order by each participant.

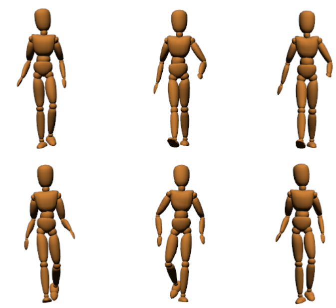

Figure 6: *Example of final positioning in the Motion Baseline Experiment.*

4.2.1 Results

Reaction times for each gait were recorded and we averaged the results over the 3 repetitions. Average reaction time for this task was 18 seconds. We found a main effect of gait ($F_{19,133} = 3.6, p < 0.0001$) where it was found that three *particular walkers were spotted more easily than others*. On examining these, we found that one was very characteristically male, with the palms of the hands facing inwards, which seemed to act as a discriminating visual cue. The other two had a distinctive head sway or ladylike steps. The most important result here is that *motion clones are detected with far more difficulty than appearance clones*, suggesting that appearance may be the dominating factor in our follow-on experiment using multiple clones.

5 Multiple Clone Experiments

In the Baseline Experiments, we examined participants' performance at detecting single clones in a simple orthographic matrix configuration. However, in most implementations of virtual crowds, several clones of the same template character will be visible. Therefore, we now describe two experiments that investigate the ability of participants to notice such multiple clones in a more realistic perspective configuration.

Throughout the experiments, all appearance clones had different outfits chosen from their set of 36. Participants were first instructed as to what an appearance or motion clone was, and then asked to click on the first two clones that they spotted, as quickly as possible. The trial was displayed onscreen until the correct pair was chosen, or 60 seconds had passed.

Our main hypothesis is that increasing the number of clones of a single model or motion will make clone pairs easier to find, and our results do in fact illustrate this, with a concomitant decrease in reaction times. A summary of the different conditions being tested in each of the experiments can be seen in Table 1.

	Appearance (A)	Motion (M)
A_C	varied, some cloned (C)	no motion
$A_C M_C$	varied, some cloned (C)	varied, some cloned (C)
$A_C M_R$	varied, some cloned (C)	random, all different (R)
M_C	all same	varied, some cloned (C)
$M_C A_R$	random, all different (R)	varied, some cloned (C)

Table 1: *A list of the different Multiple Clone experimental conditions.*

5.1 Appearance Clone Detection

Will combining appearance clones with motion clones ($A_C M_C$), so that each appearance clone has the same motion each time, hinder or aid recognition over the case where no motion is present (A_C)? Will random motions applied to appearance clones ($A_C M_R$) make this harder or easier? To explore these issues, twenty-eight participants were asked to match clones based on their appearance. Fifteen participants (11M-4F) viewed conditions A_C and $A_C M_C$ in counterbalanced order, while thirteen others (10M-3F) viewed condition $A_C M_R$. All participants were naïve to the purpose of the experiment and from different educational backgrounds.

For the $\mathbf{A}_C$**: Appearance Clones** condition, the factors were: *number of appearance clones* (10) and *orientation* (2). A crowd of 20 models was used throughout the experiment and we chose just one model to multiply clone each time. The minimum number of appearance clones dispersed among the crowd was 1 and the maximum was 10. Models were oriented either facing forward or facing in random directions.

Since we were evaluating appearance only and not motion, we used static meshes in a neutral pose. One of the 20 template models was randomly chosen at each trial to be cloned, since we were not interested in the differences between models in this experiment. All other template models in the scene were different. At each trial, the 20 models being used were placed randomly into the 20 locations we had allocated for them in a scene where occlusion was minimized (Figure 1). Three repetitions of each trial were conducted, resulting in a total of 60 trials: 10 numbers of clones * 2 orientations * 3 repetitions. Repetitions were not exact copies of each other, as models, outfits and positions of the models onscreen were randomized each time.

For the $\mathbf{A}_C\mathbf{M}_C$**: Appearance Clones with Cloned Motion** condition, the factors were: *number of appearance clones* (10) and *synchronization* (2). The two levels of synchronization meant that characters were animated with their characteristic motions either in-step or randomly out-of-step. Clones walking out-of-step is typical in many real-time systems, so we wished to know whether this aids in disguising the clones.

In order to create appearance clones with cloned motion, we applied characteristic gaits to the 20 template models from the set of 20 motions captured. One motion was always associated with the same model, and was chosen as a congruent match. This meant that we matched the model's sex with the sex of the actor that performed the motion and we also chose the motion based on similarity in size and age between the model and the actor that performed the motion. Adding characteristic motions to models is common practice in real-time crowd simulations. This is due to the fact that lower level representations often need to contain pre-baked animations so that they can be cloned very often. This is particularly true for impostors [Tecchia et al. 2002; Dobbyn et al. 2006; Pettré et al. 2006], since images from multiple viewpoints are pre-computed in advance for each frame of animation. Three repetitions of each trial were conducted, resulting in a total of 60 trials: 10 numbers of clones * 2 synchronization levels * 3 repetitions.

For the $\mathbf{A}_C\mathbf{M}_R$**: Appearance Clones with Random Motion** condition, the factor being tested was *number of appearance clones* (10), and three randomized repetitions of each condition were shown. We chose a model to be cloned and each clone of this model had a random different motion applied to it (from the 20 captured gaits). We used 3D Studio Max to retarget motions from one character to another. Using this technique, foot-plants are maintained and gravitational acceleration is altered based on the height of the character. Some artifacts occurred when female motions were applied to male characters, which could have been perceived as unnatural. However, since this was only particularly noticeable for one of the female walks, due to the randomization of trials we are confident that the results were not affected. Furthermore, no participant commented on unnatural postures when interviewed post-experiment.

Every model in the scene (including the clones) had a different random motion. The reason for testing this factor was that we felt that it might be usable as a strategy for increasing variety. This would not be possible to reproduce for a crowd of pre-generated impostors, where multiple viewpoint textures are stored for each frame of animation, since texture memory consumption would be too high. However, it could be a viable option for other lower level virtual human representations used in real-time crowds, such as low resolution polygonal models, or characters used in non real-time crowds. Three repetitions of each trial were conducted resulting in a total of 30 trials: 10 numbers of clones * 3 repetitions.

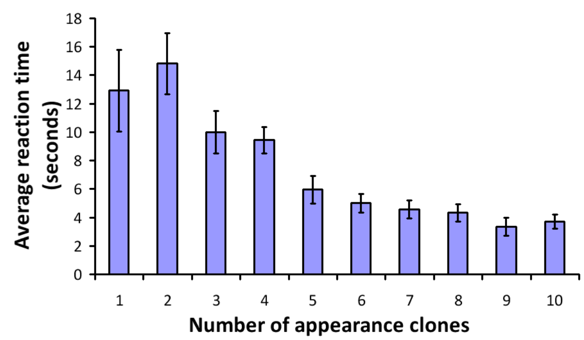

Figure 7: *Average reaction times for static appearance clones A_C (forward facing).*

5.2 Results

Two-factor ANOVAs were conducted, where number of appearance clones was a factor in all conditions. A main effect ($p < 0.0001$) was found in all cases, with an overall decreasing trend in reaction times with increasing numbers of clones. Figure 7 shows an example of this effect on A_C. A summary of the other conditions, averaged over the number of clones can be seen in Figure 8.

For $\mathbf{A}_C$, there was a within-group main effect of orientation ($F_{1,14} = 18.4, p < 0.0001$). This implies that *models facing in random directions made clones more difficult to spot*. For $\mathbf{A}_C\mathbf{M}_C$ there was a within-group main effect of synchronization ($F_{1,14} = 5.6, p < 0.04$), which showed that *playing motions out-of-step made clones more difficult to spot than when they were played in-step*.

Comparing $\mathbf{A}_C$ with $\mathbf{A}_C\mathbf{M}_C$, we tested reaction times for identifying appearance clones. Firstly, we looked at the case where characters were static compared to the case where in-step characteristic motion was applied (all facing forward). Surprisingly, we found

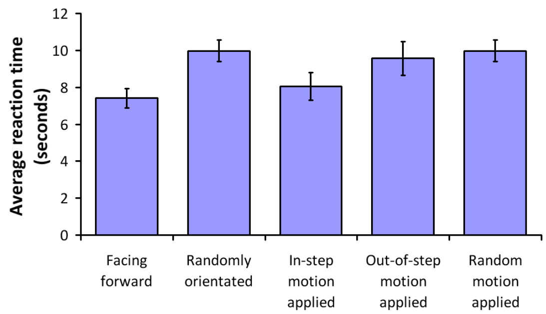

Figure 8: *Appearance Clone effects: Reaction times averaged over the 10 number of clone levels.*

that the presence or absence of in-step motion did not have an effect. However, there was an interaction ($F_{9,126} = 6, p < 0.0001$). Post hoc analysis using Neuman Keuls comparisons showed us that this effect was due to the fact that with the minimum number of clones onscreen (1), motion made it much more difficult to identify the clone than when no motion was present ($p < 0.0001$). The average reaction time for identification of 1 clone was 13 seconds when the characters were static, and 27 seconds when the characters were moving (Figure 9).

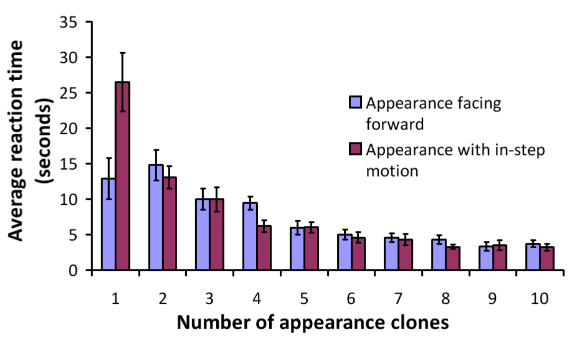

Figure 9: *Interaction between number of clones and the presence or absence of motion.*

We then tested reaction times for identifying clones in the case where characters were static compared to when they had out-of-step characteristic motion applied (all facing forward). We found that there was a main effect of out-of-step motion ($F_{1,14} = 10.6, p < 0.006$), where the addition of *out-of-step motion made it more difficult to identify appearance clones*. There was also an interaction ($F_{9,126} = 6.3, p < 0.0001$), which was again due to a single clone making it much more difficult to identify appearance clones.

Comparing $\mathbf{A}_C$ and $\mathbf{A}_C\mathbf{M}_R$ we found a between-groups main motion effect ($F_{1,26} = 19.68, p < 0.0002$), where *random motions made it more difficult to identify clones than when the clones had no motion applied*.

Comparing $\mathbf{A}_C\mathbf{M}_C$ with $\mathbf{A}_C\mathbf{M}_R$ we found a between-groups main effect of motion type ($F_{1,26} = 8.82, p < 0.008$), where *random motions disguised appearance clones more than in-step characteristic motion*. The same test was performed to compare out-of-step characteristic motion to random motion and we found no effect. Therefore, adding *random motion would not warrant the extra storage and motion capture time, since out-of-step characteristic motion is equally effective at helping to disguise appearance clones*. This is a useful result for crowd systems that use pre-baked animations, such as systems using impostors.

Figure 10: *Models that were most often misclassified as clones.*

False Positives

We counted the number of false positives for all pairs, to investigate if certain models were confused more often than others. We found that confusion occurred mainly due to clothing. For example, we found that our two executive models wearing suits were confused with each other often, even though their faces and bodies were very different (Figure 10). Similar confusion occurred for our two models with checked shirts, and for three of the female models wearing jeans. We then counted the number of false positives for each participant for each number of clones onscreen, and performed an ANOVA on the data. We found a main effect of number of clones ($F_{9,126} = 18, p < 0.0001$). Post-hoc analysis showed that the number of false positives decreased with increasing numbers of clones (ranging from an average of 1.3 false positives when one clone was onscreen to 0.1 when 10 were onscreen). We also looked at the number of failed attempts that occurred, and found only 2 failures out of all responses. These failures occurred when just one clone was onscreen.

5.3 Motion Clone Detection

Will motion clone matching at different levels of multiplicity be a much harder task than appearance clone matching? Will motion clones be disguised when appearance is varied? To answer these questions, 15 naïve participants (9M-6F) from different educational backgrounds were asked to match clones based on their motion. The order in which they viewed conditions M_C and $M_C A_R$ was counterbalanced.

For the $\mathbf{M}_C$**: Motion Clones** condition, the factor being tested was: *number of motion clones* (10). As for appearance, we tested a minimum of one motion cloned once and a maximum of one motion cloned ten times. Three randomized repetitions were shown of each condition. Motion clones were randomly dispersed among the crowd. All motions were played in-step and the participants were asked to use the mouse to click on the first pair of motion clones that they spotted, as quickly as possible.

Here we wished to analyze the effect of motion variety independently of appearance, as we wanted to see how many motion clones could be placed on the screen before the user noticed them. The 20 captured walk motions discussed in Section 3 were used. We applied these motions to the mannequin as we did not want appearance to be an influencing factor. The 20 animated mannequins were then placed in a scene, facing forward (Figure 11). We did not wish to test the random orientation condition in this case since we found the task too difficult in pilot runs and due to the fact that reaction times were so long for the Motion Baseline Experiment.

For the $\mathbf{M}_C\mathbf{A}_R$**: Motion Clones with Random Appearance** condition, the factor being tested was *number of motion clones* (10). This time we wished to test motion clones in a more realistic scenario than with the mannequins, so we used our set of 20 models (discussed in Section 3). Three randomized repetitions were shown of each condition. For every trial, a motion to be cloned

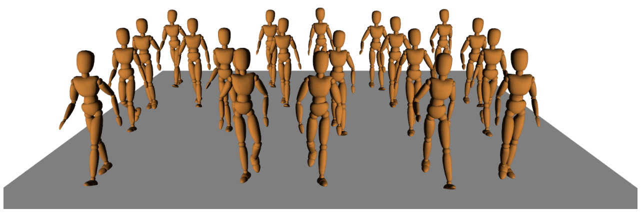

Figure 11: *Example of a crowd from M_C.*

was randomly chosen. The twenty template models were displayed for each trial, some with the cloned motion applied and the others with different random motions. As before, motions were played in-step, and participants were asked to click on the first pair of motion clones they spotted as quickly as possible.

5.4 Results

Comparing $\mathbf{M}_C$ and $\mathbf{M}_C\mathbf{A}_R$, we found a main effect of number of clones ($F_{9,126} = 107, p < 0.0001$), see Figure 12. The average *reaction time for motion clones was far slower than for appearance clones* (28 seconds). Surprisingly, *motion clones were not disguised when appearance varied*, as reaction times for the blocks with neutral and varied appearance were the same. We felt that adding variety in appearance would increase overall variety, as the same motion could look different when applied to a different body shape. However, this was not the case since there was no difference in reaction times for M_C and M_CA_R, so the level of motion variety required in a crowd seems to be generalizable and independent of the appearance of the models to which it is being applied.

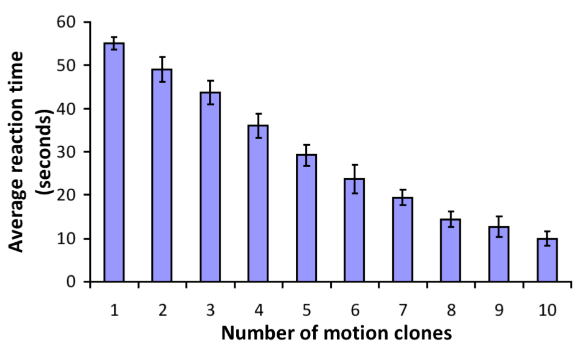

Figure 12: *Main effect of number of motion clones in the Motion Clone Detection Experiment.*

6 Conclusions

Our main finding is that appearance clones can be spotted much quicker than motion clones. This would imply that practitioners should concentrate on ensuring variety of appearance. Perhaps the addition of accessories, texture variation, or decals would aid in disguising clones [Thalmann et al. 2007; Dobbyn et al. 2006; Gosselin et al. 2005]. Based on the average reaction times of participants in the different experiments, we summarized the thresholds for motion and appearance clones that would be imperceptible at 5, 10, 15 and 20 seconds in Table 2.

From our Baseline Experiments, we found that all of the models in our set were spotted as quickly as each other and therefore it would appear that the model being cloned does not affect perceptibility of clones. However, in our Multiple Clone Experiments, we did find more false positives between models with similar clothing, so this may be a factor that could be taken into consideration when choosing template models. In all cases tested, we found that increasing the number of clones of a single model or motion will make clone pairs easier to find, with a concomitant decrease in reaction times.

Exposure	# Appearance clones	# Motion Clones
5 seconds	8	10
10 seconds	4	10
15 seconds	2	9
20 seconds	none	7

Table 2: *Summary of thresholds.*

Interestingly, we found that the presence of in-step motion did not make the task of finding appearance clones easier. Also, having appearance clones with their characteristic motion out-of-step is as effective a strategy for increasing variety as applying many different motions to the cloned models. This is a useful result, since adding random motion would not be possible to reproduce in a crowd created using pre-baked animations.

Another interesting finding was that random orientation and color modulation makes it more difficult to spot clones, so it would be advisable to adopt these techniques where possible. We also found that positioning appearance clones close to each other in the horizontal plane made them easier to see than the vertical or diagonal plane (for characters displayed in an orthographic matrix). Close-by clones on the vertical were next easiest to see, while all far and diagonal clones were most difficult to spot.

7 Future Work

In this paper, we focused on the cloning of single characters visible in the foreground. It would be interesting to examine if this would scale up when more characters are onscreen. We found that appearance dominated for the perception of our foreground crowd of twenty, but perhaps the opposite effect may be true for crowds at a distance, where the appearance is less distinguishable and the motion may become more dominant.

Using an eye-tracker to analyze what it is that participants focus on when doing these experiments would be very interesting. This may allow for perceptually guided metrics to be devised for appearance and motion variation, to focus on varying parts of the body that are visually attended to the most when identifying clones. The motions used in this work were all walking; it would be interesting to examine if similar effects occurred for different motions, such as those typically performed in crowds of characters in a stadium.

Simple pedestrian crowds were the main focus of this paper. However, more complex behaviors may affect the perception of variety, which we intend to study further. Also, future studies to analyze the effect of texture modulation and the addition of accessories would help to further our knowledge of the techniques that are most useful at disguising clones.

Acknowledgements

This work was sponsored by Science Foundation Ireland and the Higher Education Authority of Ireland.

References

ARIKAN, O. 2006. Compression of motion capture databases. *ACM Transactions on Graphics 25*, 3, 890–897.

BEARDSWORTH, T., AND BUCKNER, T. 1981. The ability to recognize oneself from a video recording of ones movements without seeing ones body. *Bulletin of the Psychonomic Society 18*, 1, 19–22.

CUTTING, J., AND KOZLOWSKI, L. 1977. Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the Psychonomic Society 9*, 5, 353–356.

DE HERAS CIECHOMSKI, P., SCHERTENLEIB, S., MAÏM, J., AND THALMANN, D. 2005. Real-time shader rendering for crowds in virtual heritage. In *Proceedings of the 6th international Symposium on Virtual Reality, Archeology and Cultural Heritage (VAST 2005)*, 1–8.

DE HERAS CIECHOMSKI, P., SCHERTENLEIB, S., MAÏM, J., AND THALMANN, D. 2005. Reviving the roman odeon of aphrodisias: Dynamic animation and variety control of crowds in virtual heritage. In *VSMM: Virtual Systems and Multimedia*, 601–610.

DOBBYN, S., HAMILL, J., O'CONOR, K., AND O'SULLIVAN, C. 2005. Geopostors: a real-time geometry / impostor crowd rendering system. In *SI3D '05: Proceedings of the 2005 symposium on Interactive 3D graphics and games*, 95–102.

DOBBYN, S., MCDONNELL, R., KAVAN, L., COLLINS, S., AND O'SULLIVAN, C. 2006. Clothing the masses: Real-time clothed crowds with variation. In *Eurographics Short Papers*, 103–106.

GOSSELIN, D., SANDER, P., AND MITCHELL, J. 2005. Drawing a crowd. *ShaderX3*, 505–517.

HAMILL, J., MCDONNELL, R., DOBBYN, S., AND O'SULLIVAN, C. 2005. Perceptual evaluation of impostor representations for virtual humans and buildings. *Computer Graphics Forum (Eurographics 2005) 24*, 3, 623–633.

HARRISON, J., RENSINK, R. A., AND VAN DE PANNE, M. 2004. Obscuring length changes during animated motion. *ACM Transactions on Graphics 23*, 3, 569–573.

HODGINS, J., O'BRIEN, J., AND TUMBLIN, J. 1998. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics 4*, 4, 307–316.

JOHANSSON, G. 1973. Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics 14*, 2, 201–211.

KOVAR, L., GLEICHER, M., AND PIGHIN, F. 2002. Motion graphs. *ACM Transactions on Graphics 21*, 3, 473–482.

LEE, J., CHAI, J., AND REITSMA, P. 2002. Interactive control of avatars animated with human motion data. *International Conference on Computer Graphics and Interactive Techniques*, 491–500.

MAÏM, J., HAEGLER, S., YERSIN, B., MUELLER, P., THALMANN, D., AND VANGOOL, L. 2007. Populating ancient pompeii with crowds of virtual romans. In *Proceedings of the 8th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST 2007)*, 26–30.

MCDONNELL, R., DOBBYN, S., AND O'SULLIVAN, C. 2005. LOD human representations: A comparative study. *Proceedings of the First International Workshop on Crowd Simulation*, 101–115.

MCDONNELL, R., DOBBYN, S., COLLINS, S., AND O'SULLIVAN, C. 2006. Perceptual evaluation of LOD clothing for virtual humans. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 117–126.

MCDONNELL, R., JÖRG, S., HODGINS, J. K., NEWELL, F., AND O'SULLIVAN, C. 2007. Virtual shapers & movers: form and motion affect sex perception. In *APGV '07: Proceedings of the 4th symposium on Applied perception in graphics and visualization*, 7–10.

PETTRÉ, J., DE HERAS CIECHOMSKI, P., MAÏM, J., YERSIN, B., LAUMOND, J.-P., AND THALMANN, D. 2006. Real-time navigating crowds: Scalable simulation and rendering. *Computer Animation and Virtual Worlds (CASA 2006) 17*, 445–455.

REITSMA, P., AND POLLARD, N. 2003. Perceptual metrics for character animation: Sensitivity to errors in ballistic motion. *ACM Transactions on Graphics 22*, 3, 537–542.

TECCHIA, F., LOSCOS, C., AND CHRYSANTHOU, Y. 2002. Visualizing crowds in real-time. *Computer Graphics Forum (Eurographics 2002) 21*, 4, 753–765.

THALMANN, D., O'SULLIVAN, C., YERSIN, B., MAÏM, J., AND MCDONNELL, R. 2007. Populating virtual environments with crowds. In *Eurographics Tutorials*, vol. 1, 25–124.

ULICNY, B., AND THALMANN, D. 2002. Towards interactive real-time crowd behavior simulation. *Computer Graphics Forum 21*, 4, 767–775.

WANG, J., AND BODENHEIMER, B. 2004. Computing the duration of motion transitions: an empirical approach. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 335–344.