The Effects of Noise on the Perception of Animated Human Running

Bobby Bodenheimer
bobbyb@cc.gatech.edu

Anna V. Shleyfman
shleyf@cc.gatech.edu

Jessica K. Hodgins
jkh@cc.gatech.edu

College of Computing
Georgia Institute of Technology

May 20, 1999

Abstract

Cyclic motions such as running or walking are difficult to animate because limitations in time and technology often result in only a small number of distinct cycles being produced. These few cycles are then repeated to produce an animation of the desired length. Unfortunately, the repetitiveness of the resulting motion often appears unnatural. This paper seeks to fix this problem by determining how to introduce natural-looking variability into cyclic animations of human motion. We construct a noise function based on biomechanical considerations that introduces natural-looking perturbations in a base motion produced either by dynamic simulation or with motion capture data. We evaluated our results through human subject testing. Our tests show that users prefer some amount of variability in cyclic motions, and quantify how much variability should be introduced into our test sequences.

1 Introduction

Creating appealing humanoid animations is a difficult and time-consuming task. Cyclic motions such as running or walking are particularly hard to animate because limitations in time and technology often result in only a small number of distinct cycles being produced, whether the method is keyframing, motion capture, or simulation [21]. These few cycles are then repeated to produce an animation of the desired length. Unfortunately, the repetitiveness of the resulting motion is often noticeable and no matter how high the quality of the individual cycle, the resulting stream of animation appears unnatural.

For example, if a running motion is created through motion capture, it should look realistic, because it was obtained from a real person. But the limited field of view of most motion capture equipment means that only a few strides of running over ground can be obtained. If an extended sequence of running is required, the repeating cycle will quickly become obvious. Simulations of running also share this difficulty: the control system attempts to maintain a steady-state running pattern and therefore it reduces the variability of the running motion. Even with a complicated simulation such as a full-body dynamic model with many controlled degrees of freedom, a good control system produces cyclic motion with little variability between strides.

This paper seeks to fix this problem by determining how to introduce natural-looking variability into cyclic animations of human motion. We assume that variability is noticeable and we hypothesize that a motion that possesses it looks more natural. From a biomechanical perspective, a fundamental attribute of biological systems is that motion varies from performance to performance. We attempt to address the problem of the repetitiveness of long animations by leveraging some of the biomechanical research on variability and incorporating it into our work. In this framework, introducing variability into animations equates to modifying the trajectories of the joint degrees of freedom from their nominal values, and can be viewed as adding a type of noise to the animation. Through human subject testing, we seek to answer the question of whether adding noise produces a more natural-looking animation, and if so, how much and what type of noise should be added.

We make a distinction in this paper between variability in a motion and the style of the motion. Style is a coarse-grained measurement conveying the difference between the motion of subjects or the overall emotional expressiveness

of a motion, while variability is a fine-grained measurement conveying the changes between repetitions of a task. A clear example of a stylistic difference in motion would be a dejected walk as compared with a buoyant walk; an example of variation in a walk would the arm occasionally brushing against the torso. Stylistic differences are very interesting from an animation perspective (see, for example, [20, 2, 24]), but our concern in this paper is with variability. Stylistic variations tend to occur when the physical constraints of the task do not constrain the motion to occur in a certain way. For example, people walk in very different ways, and the speed of these walks varies considerably. Studies have shown that in small villages the pace of walking is about one-half what it is in large cities; if we constrain people to walk at a certain speed or cadence, then people will adopt their stride length to walk in a minimum energy fashion [1]. Yet variability in the walk still exists.

Our experiments are confined to running motions, in part because animating a convincing run is extremely difficult, and in part because running is a more constrained activity than walking. Animating a motion captured run is challenging because the limited field of view of most motion capture equipment would require a very expensive setup to capture many strides. The run could be videotaped and rotoscoped, but that is also expensive as it requires an animator to edit the motion, and controlling the support of the foot becomes problematic. Running is hard to dynamically simulate because of the difficulty of constructing a control system that produces natural-looking motion.

In the next section, we review the background literature in this subject. We then describe and compare various types of noise which might be introduced to a simulation or motion capture data to add variability to the motion, and explain how we introduce the noise into our system. Section 4 discusses the experimental design we used to evaluate the resulting animations with human subjects. Finally, we evaluate our results and try to place them in the larger perspective of generating realistic animations.

2 Background

Variability in human movement has been explored extensively by biomechanists, particularly in the study of motor control and skill acquisition, e.g., [17]. Variability in movement has been a problem of long-standing interest: the first variability studies were performed by Woodworth [26] in 1899, who studied the variation in the back-and-forth movement people made when repeatedly moving a pencil through a slit to the beat of a metronome With their eyes closed, subjects had errors that were approximately constant with respect to velocity; with their eyes open, the errors increased with velocity.

In 1954, Fitts [11] conducted experiments in which people moved a stylus between two targets as rapidly as possible, where the distance between the targets and the size of the targets varied. His experiments demonstrated a logarithmic relationship between movement duration, accuracy, and movement time,

$$T = a + b \log_2 \left(\frac{2D}{W} \right) \tag{1}$$

where T is the movement duration, D is the movement distance, W is the target width (thus providing the notion of accuracy), and a and b are experimentally determined constants. Equation 1 is known as Fitts' Law. Fitts' Law holds for a large class of aiming-related movements, but it does not hold for all classes of movements: for some types of movement there is a linear relation between speed and accuracy [23, 27]. Schmidt *et al.* [23] refined Fitts' Law by performing experiments where subjects were asked to minimize the variability of their movements in a prescribed amount of time. Schmidt and his collaborators observed that the standard deviation of the movement of the endpoints, i.e., the variability, denoted by W_e, increased with the distance of the movement and decreased with the time of movement:

$$W_e = k \frac{D}{T} \tag{2}$$

where k is an experimentally determined constant. Zelaznik [27] observes that movements that do not use feedback-based corrections will not exhibit a logarithmic relationship between speed and accuracy, but rather a linear one. So, for example, most aiming movements or movements where precise end effector placement is required obey Fitts' Law; movements where the velocity is constrained obey a linear speed-accuracy tradeoff. Interestingly enough, Woodworth's original experiments were velocity-constrained. These relationships represent fundamental limitations on the control of human movement; while biomechanicists use them to investigate the processes of skill development and task performance (for example, [15]), we will use them to introduce realistic variability into our simulations.

We will treat the variability of an individual degree of freedom as independent of other degrees of freedom. Unfortunately, in human movement this is not always the case. Arutyunyan *et al.* [3, 4] have shown that in experienced marksmen aiming a pistol, the shoulder and wrist joint degrees of freedom are reciprocally covarying, resulting in lower endpoint variability of the pistol tip. In inexperienced marksmen, the covariation is weaker, which leads to higher endpoint variability. They conclude that appropriate coordination of participating joints permits low endpoint task variability while allowing higher variability in the angles of the individual joints. Nonetheless, treating the degrees of freedom as independent represents a reasonable approximation for less precisely coordinated movements. We assume this approximation is a reasonable one in our work.

Modifying existing animation to produce different characteristics is not new. Our focus in this work is not on motion editing such as that of Bruderlin and Williams [8] or Witkin and Popović [25], but is more closely related to the works of Unuma *et al.* [24], Amaya *et al.* [2], and Rose *et al.* [20]. They alter existing animation to produce different stylistic or emotive motions. While these researchers produce animation that is quite compelling, for a fixed style of movement their animations will also be repetitive. Our work does not aim to produce the emotional styles of animation that theirs does. Instead, we are focused on improving an existing cyclic stream of animation, and thus this work could be added as a postprocess to the works above.

Perlin [18, 19] has used noise to provide personality to animations. Through the addition of what he terms "coherent noise," he is able to convey small motions such as those of a character blinking, directing their gaze around a room, or maintaining balance. This work is inspirational for our own, but differs from it in that it is a system for scripting motions. An animator tunes the noise functions to convey the proper gestures or emotion. Perlin's results are impressive, but they lack the ability to respond to changes in the environment the way that a simulation can.

The simulations our work is based on were described by Hodgins *et al.* [13]. Hodgins and Pollard [12] extended this work by interpolating over the space of control laws to enable variability in the physical characteristics of the simulated figure. Thus, they were able to apply the running action to a child or an adult male or female. With the physical characteristics fixed, however, the running motion has very little perceived variability. Although the physically correct nature of the simulation and imperfect nature of the control system imply that there is some variation stride to stride, the variations are too small to be readily perceived and the motion appears repetitive. Our work seeks to remedy this repetition.

3 Construction of the Noise Function

We consider only types of noise that can be added to articulated rigid-body figures. Thus, we consider only types of noise that can be used to change a joint angle's trajectory over time, and ignore variations that would change a figure's limb lengths, for example. This limitation is not too restrictive, because it applies to most dynamic simulations, most uses of motion capture, and many types of keyframed animation. Additionally, we only add noise to the arm degrees of freedom. Noise could be added to other joints, although adding motion to the legs of a simulated runner might interfere with balance, or lead to foot slippage in a motion captured or keyframed runner. Because dynamic simulations produce physically correct motion, they generally place the most constraints on the ways a joint angle trajectory can be perturbed. Thus, we first consider the types of noise that can be added to a simulation, and later explore how these types apply to the other methods of animation.

Noise can be added into a simulation in a variety of different ways, e.g., by adding noise to the sensors, to the control gains, to the output torques, or via disturbances in the surrounding environment such as rough terrain. Adding disturbances to the surrounding environment is important for applications where characters need to interact with their environment, but changing the terrain model randomly is unappealing and reduces the realism of the simulation. Perturbing the control gains has the *a priori* objection that it will change the robustness of the system, possibly leading to instability (the character will "fall down"). As a result, we only add noise to the sensors or perturb the output torques.

The equation for the output torque for each degree of freedom of the simulation is given by

$$\tau = k(\theta_d - \theta) + k_v(\dot{\theta}_d - \dot{\theta}) \tag{3}$$

where τ represents the torque of the internal joint, θ is the joint angle, θ_d is the desired joint angle, $\dot{\theta}$ and $\dot{\theta}_d$ are their respective velocities, and k and k_v are control gains. Equation 3 describes a typical proportional-derivative servo. If

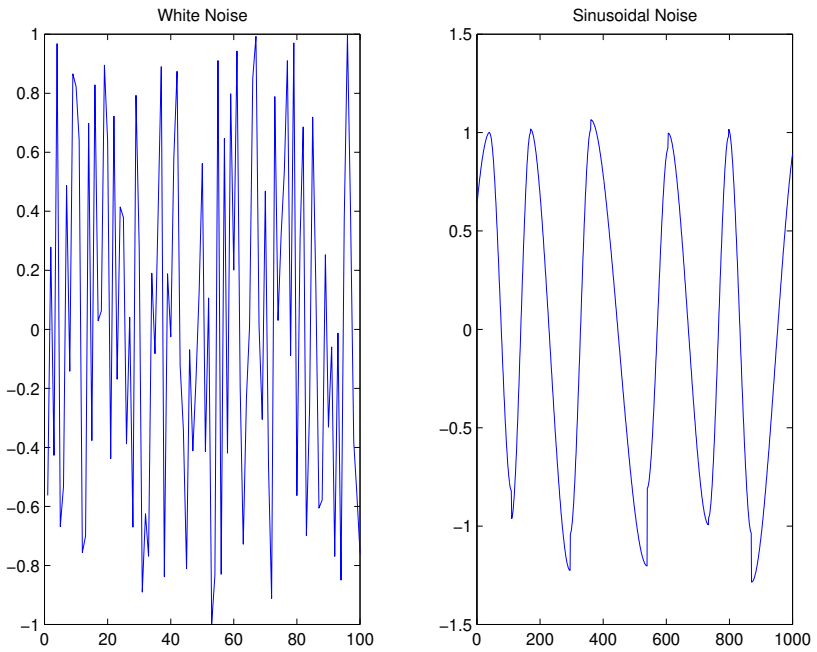

Figure 1: Examples of white and sinusoidal noise. White noise is shown in the left plot and sinusoidal in the right.

we desire to perturb the output torques, Equation 3 becomes

$$\tau = k(\theta_d - \theta) + k_v(\dot{\theta}_d - \dot{\theta}) + \tau_p \tag{4}$$

where τ_p represents the perturbation. If we desire to add sensor noise, then Equation 3 becomes

$$\tau = k(\theta_d - \theta - \theta_p) + k_v(\dot{\theta}_d - \dot{\theta}) \tag{5}$$

where θ_p represents the sensor noise. Note that changing the trajectory of the arm swing can impart moment to the rest of the body; the control system will attempt to minimize this disturbance but it will produce some variability throughout the body.

We experimented with using both Equations 4 and 5 in generating simulations, and found that perturbing the torques directly did not lead to satisfactory motion. An intuitive explanation of why this occurs may be that the control system is constantly correcting for the effect of the torque perturbation on the previous time step. To be visible, the torque perturbations must be substantial and the resulting motion appears jerky and unnatural. Equation 5, on the other hand, changes the desired angles and therefore works with the control system to produce motion that is much smoother and more natural looking.

3.1 Types of Noise

Many different types of noise exist in nature, and we considered three types. White noise is noise that is uniformly distributed and uncorrelated, of a limited amplitude. This differs from the definition of white noise commonly employed by physicists and engineers, where the distribution is most often normal (a consequence of the central limit theorem) and the amplitude is unlimited (for a discussion, see [14]). Examples of white noise abound in nature; one such example is the noise of a waterfall. Another type of noise is sinusoidal noise similar to that used by Perlin. This noise varies sinusoidally in both frequency and amplitude and can be thought of as a modulated wave pattern. Examples of this type of noise in nature would be sand ridges formed by wind or currents; see, for example, the pictures and discussion in [10]. Code to generate this type of noise can be found in [22]. Examples of both these types of noise are shown in Figure 1.

A priori, both of these noise types seemed reasonable to add to a simulation via Equation 5. However, they produce motion which appears odd because the perturbations have no contextual relation to the larger movements of the body,

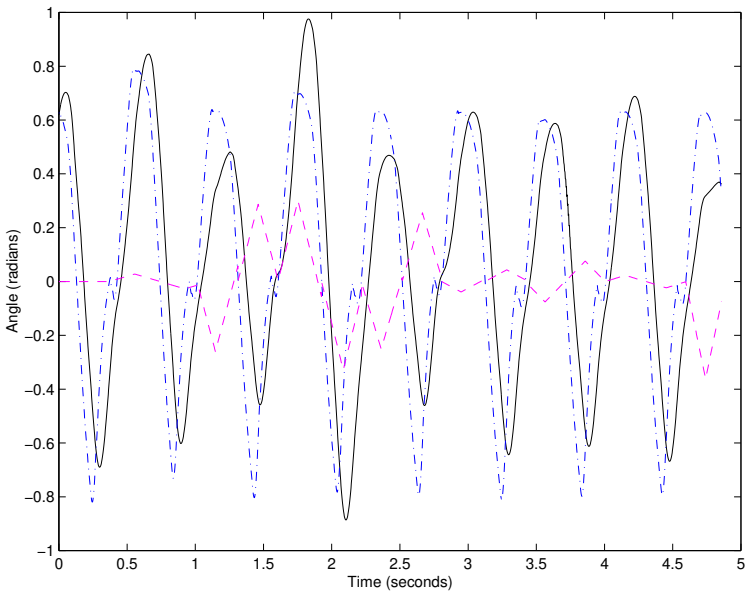

Figure 2: An example of a perturbation in phase with the y (forwards/backwards) rotation of the shoulder. The nominal trajectory is shown in the blue dash-dot line, the perturbation is shown in the magenta dashed line, and the perturbed trajectory is shown in the black solid line.

such as arm swing. People do not usually exhibit sizeable twitches and jerks when they are moving; human variability is usually correlated to the larger movements of the body [9]. Thus, our perturbations should in some way be correlated with the movements of the body.

The final type of noise we considered was a continuous noise function with its maximum amplitude occurring at the extrema of a degree of freedom during a cycle. The value of the maximum amplitude is a white noise process as described earlier. This perturbation has the advantage of always occurring in phase with the movement of the body. The perturbation could have any wave form with such a characteristic amplitude; we chose to use either a sinusoid or a triangle wave. For a simulation, the control system smoothes the triangle wave so that there is no perceptible difference between it and a sinusoid; a sinusoid is more appealing for use with motion capture data where no smoothing occurs. An example of this type of noise together with the shoulder trajectory is shown in Figure 2. In this figure, the sum of the two non-solid lines represents the desired value that the control system is attempting to track. Informal tests led us to believe that this type of noise leads to visually appealing motion, and we quantify this with the results of users studies in Section 4.

3.2 Scaling the Noise

We must determine the appropriate maximum amplitude of the perturbation for each degree of freedom. To facilitate tests with human subjects, we also need to have one scaling control that determines the amount of noise for each degree of freedom. Several types of scalings are possible, and the choice of scalings represents a secondary control on how variability is added into the simulation. A uniform scaling will not work because the amount of rotation about each degree of freedom is not the same. For example, the forward and backward swing angle (y) of the shoulder rotates through a much greater angle than the left and right rotation represented by x (see Figures 3 and 4). Adding the same amount of noise to each degree of freedom would either result in no noticeable effect on the y rotation, or an abnormally large effect on the x rotation. Thus, we base our scaling law on the biomechanical observation expressed by Equation 2. This equation says that the greater the movement distance, the greater the variability in the movement. As a result, noise is scaled for each degree of freedom based on the magnitude of the trajectory of the unperturbed degree of freedom.

This scaling law was applied to a simulation of a male runner, a female runner, and male motion capture data. These examples are discussed further in Section 5 and 6, but the scalings for the examples are shown in Table 1. The

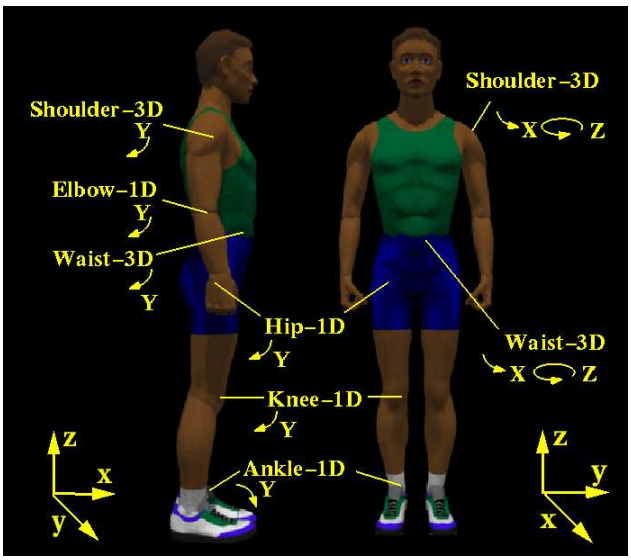

Figure 3: The coordinate systems of the joints in the simulated runner.

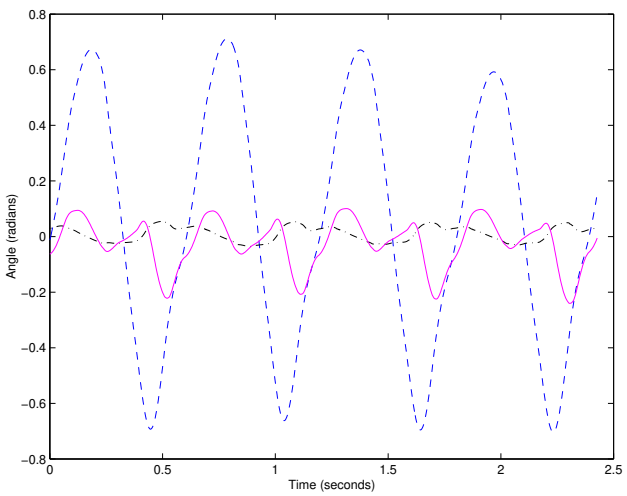

Figure 4: The zxy Euler angles for the right shoulder of the male runner simulation. The z rotation is shown in the black dash-dot line, the x rotation is shown in the magenta solid line, and the y rotation is shown in the blue dashed line.

	Male Simulation	Female Simulation	Motion Capture
Shoulder x	0.23	0.21	0.25
Shoulder y	1.00	1.00	0.42
Shoulder z	0.07	0.09	1.00
Elbow y	0.11	0.09	0.00

Table 1: The scaling relationships for the arm degrees of freedom for the male simulation, female simulation, and male motion capture example.

wrist is simulated with two degrees of freedom. Its joint angle trajectories would have given it slightly more noise than the elbow. However, adding noise to the wrist produced no perceptible variation in the motion, and reduced the maximum amount of noise that could be added to the simulation before it became unstable, so we neglected it.

4 Experimental Design

Our goal was to design an experiment that gave the user the ability to vary the amount of noise in the simulation and decide what level of noise appeared most "natural." This experimental design is opposed to a forced-choice (or A/B comparison) experiment in which users would be presented with several sets of two example motions and asked which looked more natural.

For each simulation we constructed 10 MPEG animations of running motion, each 10 seconds long. The amount of noise present in each animation was linearly increased starting from zero noise to an experimentally determined maximum. This maximum was the amount of noise that a 10 second simulation could tolerate without falling down. The noise was generated using a random number generator that was seeded at the beginning of each set of ten animations. The animations were presented to the users on an SGI R10000 machine and ran at 30 frames per second. The users could view the animations as many times as they wished, in any order, before answering which animation appeared the most natural. Samples of the animations are shown in Figures 5 and 6 for the male and female simulations. In most of the comparison frames, the arm configuration varies due to the addition of the noise. However, in the middle images of Figure 5, the legs have a different configuration even though they are not directly perturbed.

5 Results

Let n be the number of experiments conducted, and let $\hat{x}$ denote the mean value of the set, i.e.,

$$\hat{x} = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{6}$$

where x_i denotes the answer a particular subject provided. The standard deviation of the sample is

$$S_x = \sqrt{\frac{(\sum_{i=1}^{n} x_i^2) - n\hat{x}^2}{n-1}}. \tag{7}$$

Assume the underlying distribution of the samples is Gaussian. The 95% confidence interval is defined as the interval within which the true mean value ξ of the population will lie with probability 0.95. The confidence interval is given by

$$\hat{x} - t_{0.025,n}\frac{S_x}{\sqrt{n}} < \xi < \hat{x} + t_{0.025,n}\frac{S_x}{\sqrt{n}} \tag{8}$$

where $t_{0.025,n}$ is computed from the t-distribution for a set value of n. Values of n of interest to us are shown in Table 2. We compare two experiments based on their means with a t-test of the sample means, using (see [5]):

$$t = \frac{\hat{x}_1 - \hat{x}_2}{\sqrt{(S_{x_1}^2/n_1) + (S_{x_2}^2/n_2)}} \tag{9}$$

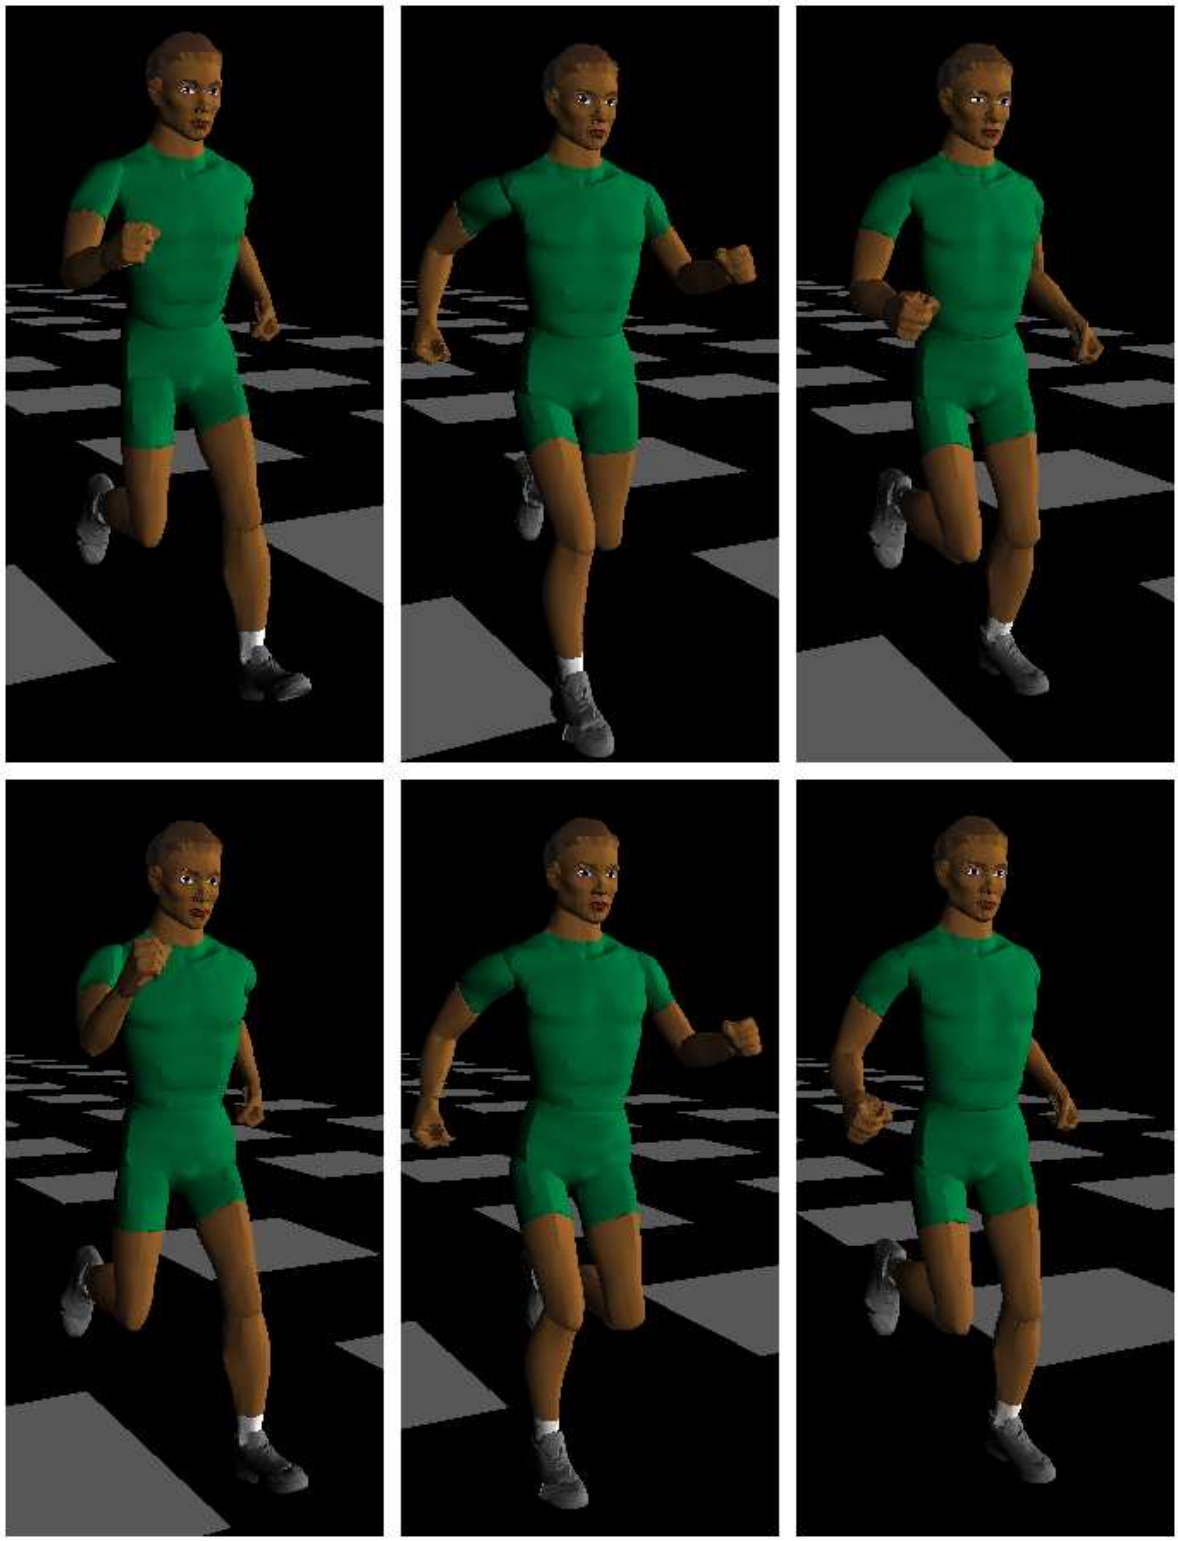

Figure 5: Example frames of the male simulation with and without noise. The top row shows various poses without noise; the bottom row shows analogous poses with noise.

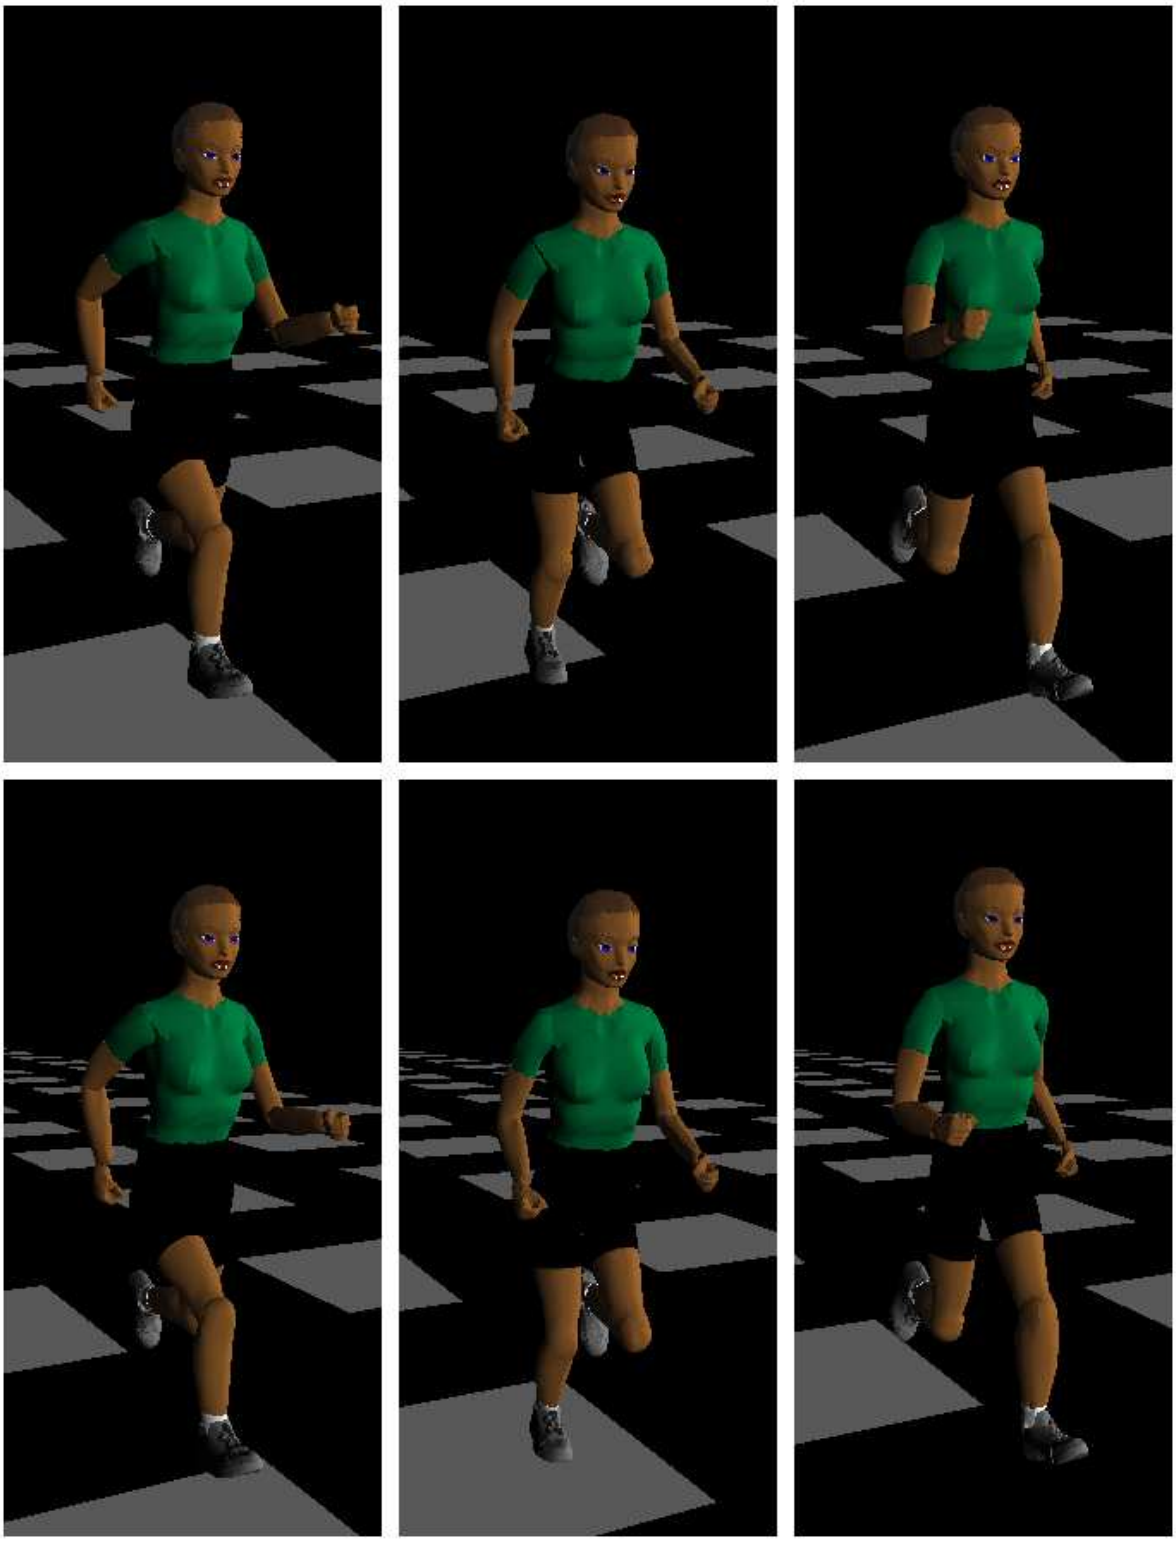

Figure 6: Example frames of the female simulation with and without noise. The top row shows various poses without noise; the bottom row shows analogous poses with noise.

n	$t_{0.025,n}$
10	2.2281
30	2.0423
50	2.0086

Table 2: Values of $t_{0.025,n}$ for various values of n, from [16].

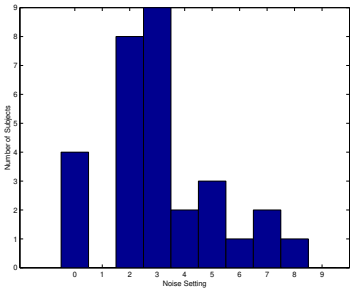

Figure 7: Distribution of 30 subject choices for one set of example simulations of the male runner: $\hat{x} = 3.13$ and $S_x = 2.06$.

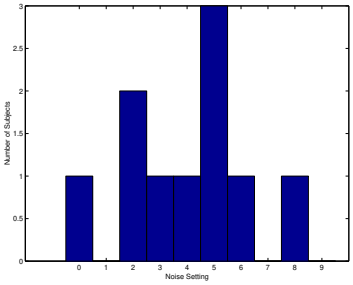

Figure 8: Distribution of 10 subject choices for a second set of example simulations of the male runner: $\hat{x} = 4$ and $S_x = 2.31$.

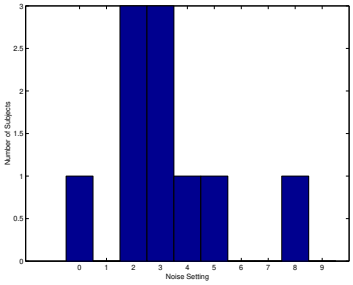

Figure 9: Distribution of 10 subject choices for a third set of example simulations of the male runner: $\hat{x} = 3.2$ and $S_x = 2.15$.

where the subscripts 1 and 2 reference the relevant statistics of the two experiments, and the value of t obtained is compared to the interval $[-t_{\alpha/2,\nu},\ +t_{\alpha/2,\nu}]$ obtained from the t-distribution for a given level of confidence α. The degree of freedom ν is approximately calculated with the formula

$$\nu = \text{floor}\left(\frac{\left(\left(S_{x_1}^2/n_1\right) + \left(S_{x_2}^2/n_2\right)\right)^2}{\frac{(S_{x_1}^2/n_1)^2}{n_1-1} + \frac{(S_{x_2}^2/n_2)^2}{n_2-1}}\right) \tag{10}$$

where $\text{floor}(x)$ denotes the greatest integer less than x.

For a given seed of the random number generator, we created ten animations of the male runner with varying degrees of noise, as described above. We then conducted 30 subject experiments. The results are shown in Figure 7; the mean $\hat{x}$ of the answers is 3.13, the standard deviation S_x is 2.06, and the mode is 3. The 95% confidence interval for this experiment is $[2.376, 3.890]$.

Using a second seed for the random number generator, we generated an additional ten animations of the runner and conducted experiments using ten subjects. The results are shown in Figure 8. The mean for this experiment was 4, the standard deviation 2.31, and the mode 5. The 95% confidence interval for this experiment is $[2.372, 5.627]$. Although the distributions appear different, the confidence interval of the 30 subject test lies within the confidence interval of the ten subject test, and a t-test of the means reveals that there is not a significant difference of the means at a 95% confidence level.

Generating a third set of animations for a third seed and conducting a ten subject experiment for this data resulted in the distribution shown in Figure 9. The mean for this experiment was 3.2, the standard deviation was 2.15, and the mode was tied at 2 and 3. The 95% confidence interval for this test is $[1.685, 4.715]$ and there is not a significant difference in the means of this test and the other tests at a 95% confidence level. This distribution has much the same shape as the distribution of Figure 9.

Combining these three experiments into a single result gives us the distribution shown in Figure 10. The mean for this distribution is 3.32, the standard deviation 2.11, and the mode is tied at 2 and 3. The 95% confidence interval for the entire distribution is $[2.720, 3.920]$.

The numbers chosen by the subjects can be translated into amplitudes of our noise function and related to the peak-to-peak amplitudes of joint rotations. For the male running simulation, the largest joint rotation is the y rotation of the shoulder (from Table 1). The y shoulder trajectory with no noise averages a peak-to-peak value of 1.5 radians. Noise is input directly into the y shoulder rotation without scaling (its scale factor is 1.0). The maximum amount

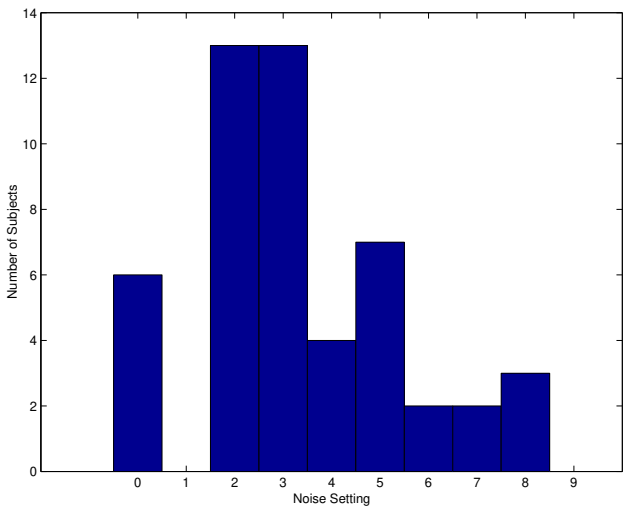

Figure 10: Distribution of all subject choices (50) for simulations of the male runner: $\hat{x} = 3.32$ and $S_x = 2.11$.

of noise the simulation could tolerate for 10 seconds was 0.65 radians peak-to-peak, or about 45% of the nominal peak-to-peak value. Thus, the answer "3" in our experiment corresponds to a maximum noise amplitude of about 15% of the peak-to-peak value of the largest joint rotation, or about 0.23 radians peak-to-peak. If we assume that perception of noise is approximately linear, then the mean of the 50 subject test corresponds to a noise level of 17%, or about 0.25 radians peak-to-peak.

The scaling law can also be applied to motion capture. The motion capture example we constructed used dynamic simulation and the control system of the male running simulation to move the lower body, but the joint angles of the arms were specified by the motion capture data without any simulation. With motion capture data, we can add noise of an arbitrarily large size; however, at some noise level the arms will start intersecting with the rest of the body. We set our maximum noise level so that self-intersection does not occur. This maximum is at about 35% of the peak-to-peak value of the z rotation. Although we have not yet conducted formal user studies, a setting of 15% of the peak-to-peak value of the z rotation appears to produce appealing motion for the motion capture data, just as it did for the simulated data.

The benefit of using a hybrid motion capture-simulation approach is that it prevents foot slide from occurring, a common problem when using motion capture data [6]. The obvious disadvantage is that we are not using motion capture for the leg motion, and the style of arm motion may not match the leg motion. From Table 1, we observe that the motion capture runner is running in a fundamentally different style than both the simulations, as the motion capture rotates around the z-axis much more than the x or y axes. This style is depicted in Figure 11, for different levels of noise. Notice that the middle images of this figure are identical, highlighting an attribute of our noise function: since the perturbation is continuous and has its maximum at the extremum of a joint trajectory, there is very little noise when the joint angles are close to zero.

6 Discussion

For the male running simulation, we have devised a noise function that introduces variability into the simulation and have shown that most users prefer a simulation with some noise (about 30% of the stability limit or 17% of the peak-to-peak amplitude of the largest joint trajectory) over no noise or higher amounts of noise. We justify this claim by comparing the confidence intervals of all of our experiments and note that they are all consistent. Comparing confidence intervals provides a more reliable test of significance than simply comparing means, see for example [7]. This result is encouraging because the control system for the running simulation is robust enough to be stable (run without falling down for 100 seconds) at this noise setting, yet the noise setting is enough to produce perceptible variability over shorter lengths of animation than the ones used in the subject test. Anecdotal evidence from the comments of the subjects indicates that many people find that no noise or low levels of noise look too stiff, while large

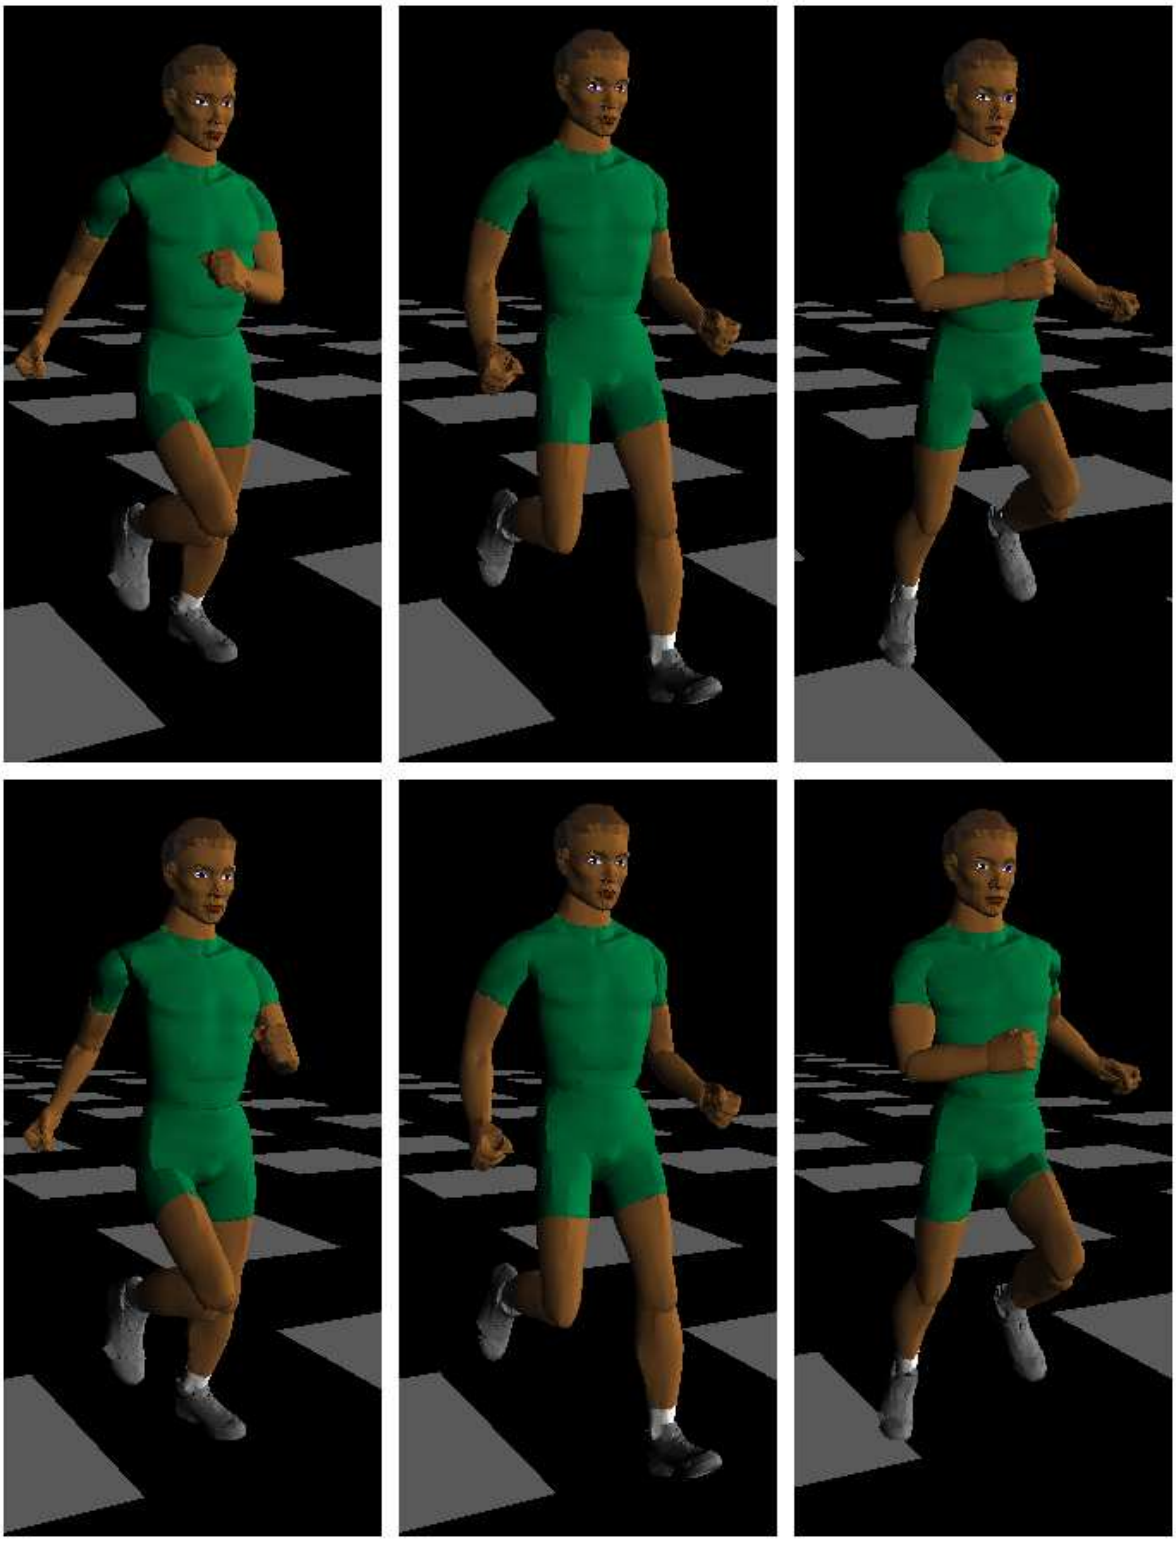

Figure 11: The motion capture-simulation hybrid with and without noise. The top row has no noise; the bottom row shows analogous poses with noise.

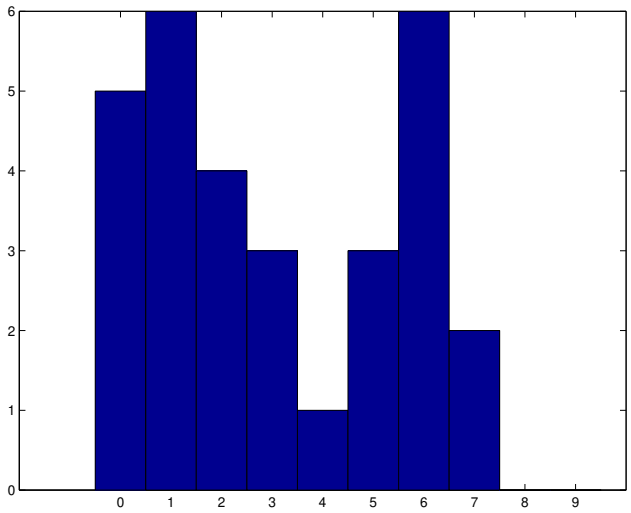

Figure 12: Distribution of 30 subject choices for one set of example simulations of the female runner: $\hat{x} = 3.07$ and $S_x = 2.43$.

levels of noise appear unnatural. When this noise function was applied to motion capture data with a different scaling, it also produced appealing results.

A good additional test of the validity these results would be to use χ^2 tests to determine how closely our distribution approximates a normal distribution [7]. Unfortunately, these tests are rather limited for the numbers of subjects we have tested and we would need to conduct more user studies to make them applicable. Other experimental designs, such as A/B comparisons, or perhaps A/B/C comparisons, would further validate our results.

An interesting avenue of research would be to try to incorporate variability into stylistic animations such as those of [20], although any animation system producing stylized animation cycles would work as well. Based on our experience with the male running simulation and, to a lesser extent, the motion capture example, we would start by perturbing these stylized animations using our scaling law and adding about 17% of the average peak-to-peak amplitude of the largest joint rotation to the animation. We believe that this would represent a good starting point for further tuning.

Unfortunately, these results do not seem to extend to the simulation of the female runner. We generated ten animations of the female runner and conducted experiments using thirty subjects. The results are shown in Figure 6. The mean for this experiment was 3.07, the standard deviation 2.43, and the mode 1 and 6. Thus, this distribution is bimodal. For this simulation, many subjects prefer a slight amount of noise but many subjects also prefer a large amount of noise. This result is perplexing, for though it tells us that users tend to prefer some variability in their animations, it does not provide any guidelines for how much noise should be reasonably added, and indicates that intermediate values of noise may be deleterious. One hypothesis is that the control system for the female runner produces less natural-looking motion and the addition of noise is not addressing the fundamental problem with the motion. We are currently investigating this result further.

At the conclusion of taking the tests with the male and female runners, some subjects made remarks like "This (number) looks best for a trained runner; this (other number) looks best if I knew the runner were out of shape." Hence, there is some overlap between variability and style, and an investigation of the boundaries between variability and style might prove interesting.

There are other types of perturbations than those we have considered here. For example, we could model the variability people actually exhibit by using several motion capture sequences from a single runner or sequences from multiple runners. Additional user studies would be required to assess whether these data-based models produced more natural-looking motion.

References

[1] ALEXANDER, R. M. *The Human Machine*. Columbia University Press, New York, 1992.

[2] AMAYA, K., BRUDERLIN, A., AND CALVERT, T. Emotion from motion. In *Graphics Interface '96* (May 1996), W. A. Davis and R. Bartels, Eds., pp. 222–229.

[3] ARUTYUNYAN, G. H., GURFINKEL, V. S., AND MIRSKII, M. L. Investigation of aiming at a target. *Biophysics 13* (1968), 536–538.

[4] ARUTYUNYAN, G. H., GURFINKEL, V. S., AND MIRSKII, M. L. Organization of movements on execution by man of an exact postural task. *Biophysics 14* (1969), 1162–1167.

[5] BECKWITH, T. G., MARANGONI, R. D., AND LIENHARD V, J. H. *Mechanical Measurements*. Addison-Wesley, Reading, Mass., 1993.

[6] BODENHEIMER, B., ROSE, C., ROSENTHAL, S., AND PELLA, J. The process of motion capture: Dealing with the data. In *Computer Animation and Simulation '97*, D. Thalmann and M. van de Panne, Eds. Springer NY, Sept. 1997, pp. 3–18. Eurographics Animation Workshop.

[7] BOX, G. E. P., HUNTER, W. G., AND HUNTER, J. S. *Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building*. John Wiley, New York, 1978.

[8] BRUDERLIN, A., AND WILLIAMS, L. Motion signal processing. In *Computer Graphics* (Aug. 1995), pp. 97–104. Proceedings of SIGGRAPH 95.

[9] DARLING, W. G., AND STEPHENSON, M. Directional effects on variability in upper limb movements. In *Variability and Motor Control*, K. M. Newell and D. M. Corcos, Eds. Human Kinetics, Champaign, IL, Apr. 1993, pp. 65–88.

[10] DUBÉ, R. L. *Natural Pattern Forms*. Van Nostrand Rheinhold, New York, 1997.

[11] FITTS, P. M. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology 47* (1954), 381–391.

[12] HODGINS, J. K., AND POLLARD, N. S. Adapting simulated behaviors for new characters. In *Computer Graphics* (Aug. 1997), pp. 153–162. Proceedings of SIGGRAPH 97.

[13] HODGINS, J. K., WOOTEN, W. L., BROGAN, D. C., AND O'BRIEN, J. F. Animating human athletics. In *Computer Graphics* (Aug. 1995), pp. 71–78. Proceedings of SIGGRAPH 95.

[14] JAZWINSKI, A. H. *Stochastic Processes and Filtering Theory*, vol. 64 of *Mathematics in Science and Engineering*. Academic Press, New York, 1970.

[15] KIM, K., MCMILLAN, M., AND ZELAZNIK, H. N. Behavioral analysis of trajectory formation: The speed-accuracy trade-off as a tool to understand strategies of movement control. In *Advances in Motor Learning and Control*, H. N. Zelaznik, Ed. Human Kinetics, Champaign, IL, 1996, pp. 1–12.

[16] MEYER, P. L. *Introductory Probability and Statistics*. Addison-Wesley, Reading, Mass., 1970.

[17] NEWELL, K. M., AND CORCOS, D. M., Eds. *Variability and Motor Control*. Human Kinetics Publishers, Champaign, IL, 1991.

[18] PERLIN, K. Real time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics 1*, 1 (Mar. 1995), 5–15.

[19] PERLIN, K., AND GOLDBERG, A. Improv: A system for scripting interactive actors in virtual worlds. In *Computer Graphics* (Aug. 1996), pp. 205–216. Proceedings of SIGGRAPH 96.

[20] ROSE, C., COHEN, M., AND BODENHEIMER, B. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Computer Graphics and Applications 18*, 5 (1998), 32–40.

[21] ROSE, C. F., GUENTER, B., BODENHEIMER, B., AND COHEN, M. Efficient generation of motion transitions using spacetime constraints. In *Computer Graphics* (Aug. 1996), pp. 147–154. Proceedings of SIGGRAPH 96.

[22] SCHLICK, C. Wave generators for computer graphics. In *Graphics Gems V*, A. W. Paeth, Ed. Academic Press, New York, 1995, pp. 367–374.

[23] SCHMIDT, R. A., ZELAZNIK, H. N., HAWKINS, B., FRANK, J. S., AND QUINN, J. T. Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review 86* (1979), 415–451.

[24] UNUMA, M., ANJYO, K., AND TEKEUCHI, R. Fourier principles for emotion-based human figure animation. In *Computer Graphics* (Aug. 1995), pp. 91–96. Proceedings of SIGGRAPH 95.

[25] WITKIN, A., AND POPOVIĆ, Z. Motion warping. In *Computer Graphics* (Aug. 1995), pp. 105–108. Proceedings of SIGGRAPH 95.

[26] WOODWORTH, R. S. The accuracy of voluntary movement. *Psychological Review 3* (1899), 1–114.

[27] ZELAZNIK, H. N. Necessary and sufficient conditions for the production of linear speed-accuracy trade-offs in aimed hand movements. In *Variability and Motor Control*, K. M. Newell and D. M. Corcos, Eds. Human Kinetics, Champaign, IL, 1991, pp. 91–115.