The Effects of Noise on the Perception of Animated Human Running

Bobby Bodenheimer, Anna V. Shleyfman, and Jessica K. Hodgins

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
{bobbyb|shleyf|jkh}@cc.gatech.edu

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 create 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 running motion produced either by dynamic simulation or from motion capture data. We evaluate our results through human subject testing.

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 [18]. These few cycles are then repeated to create 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. A run could be videotaped and rotoscoped with an animator adding in variations of the motion, but that requires an animator to edit the motion, and controlling the support of the foot becomes problematic. Simulations of running can also appear repetitive: 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 leverage some of the biomechanical research on variability and incorporate 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. 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. 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 a motion. Style is a measure conveying the difference between the motion of two subjects or the overall emotional expressiveness of a motion, while variability is a measure 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 be the arm occasionally brushing against the torso. Stylistic differences are very interesting from an animation perspective (see, for example, [17, 1, 21]), but our concern in this paper is with variability.

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 context of generating realistic animations.

2 Background

Variability in human movement has been explored by biomechanists, particularly in the study of motor control and skill acquisition [14]. Variability in movement has been a problem of long-standing interest: the first variability studies were performed by Woodworth [23] 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 [9] 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. It 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 [20, 24]. Schmidt *et al.* [20] 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 [24] 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, [13]), 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, although Arutyunyan *et al.* [2, 3] 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. Nonetheless, treating the degrees of freedom as independent probably represents a reasonable approximation for less precisely coordinated movements.

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 [6] or Witkin and Popović [22], but is more closely related to the works of Unuma *et al.* [21], Amaya *et al.* [1], and Rose *et al.* [17]. 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. 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 [15, 16] 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, and a generalization of our results might provide insight into what noise functions should be included with his system.

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 of a human runner. 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. Note that changing the trajectory of the arm swing can impart movement to the rest of the body; the control system will attempt to minimize this disturbance but it will produce some variability throughout the body. 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. The simulations our work is based on were described by Hodgins *et al.* [11].

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 uneven 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 for many applications. 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

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

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 we perturb the output torques, Eq. 3 becomes

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

where τ_p represents the perturbation. If we add sensor noise, then Eq. 3 becomes

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

where θ_p represents the sensor noise.

We experimented with using both Eqs. 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 definition differs from that 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 [12]). 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 [8]. Code to generate this type of noise can be found in [19]. Examples of both these types of noise are shown in Fig. 1.

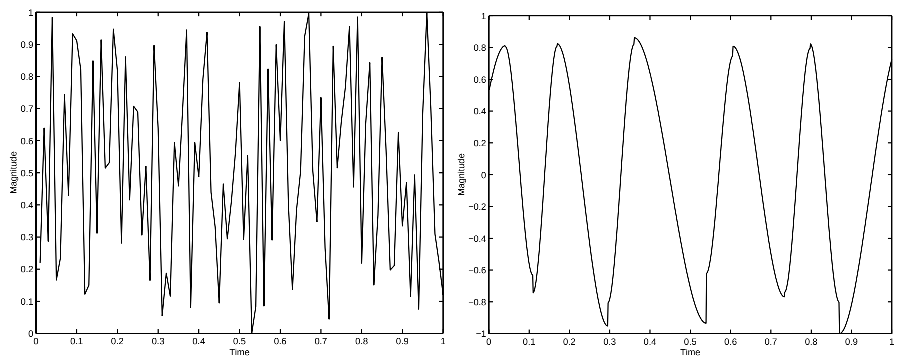

Fig. 1. Examples of white and sinusoidal noise. White noise is shown in the left plot and sinusoidal in the right.

A priori, both of these noise types seemed reasonable to add to a simulation via Eq. 5. However, if the magnitude is a significant fraction of the magnitude of the original motion, these noise types produce motion which appears odd because the perturbations have no contextual relation to the larger movements of the body, 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 [7]. Thus, our perturbations should also be correlated.

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 Fig. 2. In this figure, the solid line is the sum of the two non-solid lines and 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 tested our hypothesis with the user studies presented in Section 4.

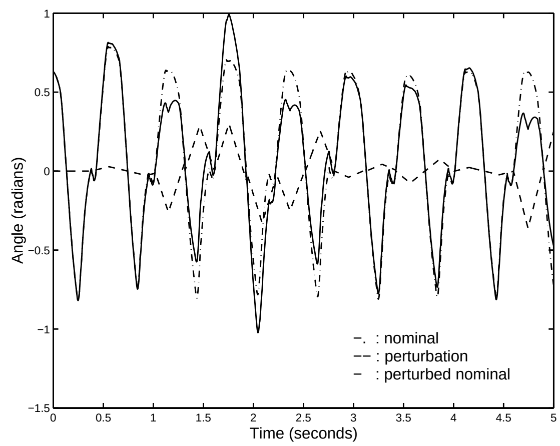

Fig. 2. An example of a perturbation in phase with the y (forwards/backwards) rotation of the shoulder.

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 all degrees of freedom. The choice of scaling 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 (y) of the shoulder rotates through a much greater angle than the left and right rotation represented by x (see Figs. 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 Eq. 2: 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 and a female runner. These examples are discussed further in Section 5 and 6, but the scalings for the examples are shown in Table 1. The 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.

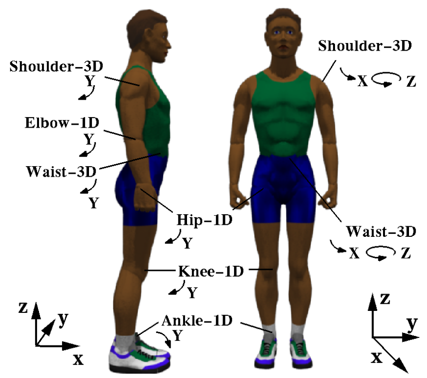

Fig. 3. The coordinate systems of the joints in the simulated runner.

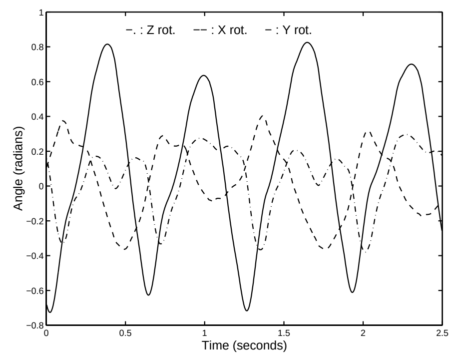

Fig. 4. The zxy Euler angles for the right shoulder of the male runner simulation.

	Male Simulation	Female Simulation
Shoulder x	0.23	0.21
Shoulder y	1.00	1.00
Shoulder z	0.07	0.09
Elbow y	0.11	0.09

Table 1. The scaling relationships for the arm degrees of freedom for the male and female simulations.

4 Experimental Design

Our goal was to design an experiment that gave a test subject the ability to vary the amount of noise in the simulation and decide what level of noise appeared most "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 from zero 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 O2 machine and ran at 30 frames per second. The users could view the animations as many times as they wished, in any order, before deciding which animation appeared the most natural. Samples of the animations are shown in Figs. 6 and 7 (see Appendix) for the male and female simulations, and the MPEG movies can be found at [10]. In most of the comparison frames, the arm configuration varies due to the addition of the noise. However, in the middle images of Fig. 6, the legs have a different configuration even though they are not directly perturbed.

5 Results

We will use $\hat{x}$ to denote the mean of an experiment with n subjects, and S_x to denote the standard deviation of the experiment. Assuming 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, and is computed from the t-distribution [4]. Using a t-test, two experiments can be compared based on their sample means [4], or on their confidence intervals [5].

Exp. No.	No. of Subjects	$\hat{x}$	S_x	Mode	95% confidence interval
1	30	3.13	2.06	3	[2.376,3.890]
2	10	4.00	2.31	5	[2.372,5.627]
3	10	3.20	2.15	2,3 (tie)	[1.169,4.715]
All	50	3.32	2.11	2,3 (tie)	[2.720,3.920]
Female	30	3.07	2.43	1,6 (tie)	NA

Table 2. The statistics for three different experiments with the male running simulation, the combined results, and an experiment with the female running simulation.

For each of three different seeds of the random number generator, we created ten animations of the male runner with varying degrees of noise, as described above. We then conducted three different experiments with different subjects. The mean $\hat{x}$, standard deviation S_x, mode, and 95% confidence intervals for the individual experiments and for the combined results are shown in Table 2. Although the distributions appear different, for example, between experiments 1 and 2, the confidence interval of the 30 subject test lies within the confidence interval of the 10 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.

All experiments are consistent based on their confidence intervals, and this measure of correspondence provides a stronger measure than simply comparing the means [5].

The numbers chosen by the subjects can be translated into amplitudes of the 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. The scale factor for the y shoulder rotation is 1.0. The maximum amount 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 in a local region, then the mean of the 50 subject test corresponds to a noise level of 17%, or about 0.25 radians peak-to-peak.

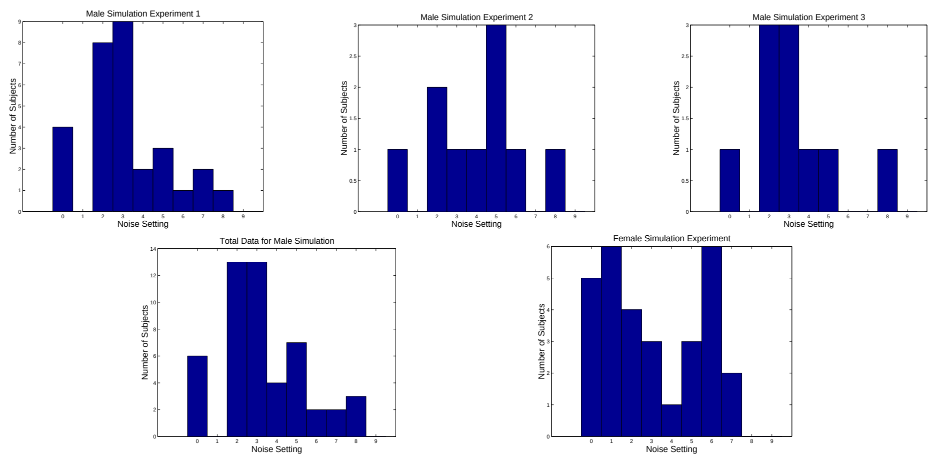

Fig. 5. Distributions of subject choices for the experiments. See Table 2.

6 Discussion

We have devised a noise function that introduces variability into the simulation and have shown that for the male running simulation most users prefer 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. 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 levels of noise appear unnatural. Additionally, 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, variability may be dependent on style, and an investigation of the boundaries between variability and style could prove interesting.

A good additional test of the validity of these results would be to use χ^2 tests to determine how closely our distribution approximates a normal distribution [5]. Unfortunately, these tests are limited for the numbers of subjects we have tested. Other experimental designs, such as A/B comparisons, or perhaps A/B/C comparisons, might also provide insights into what noise function creates a natural-looking motion.

An interesting avenue of research would be to try to incorporate variability into stylistic animations such as those of [17]. Based on the test results from the male running simulation, 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 amount of noise would represent a good starting point for further tuning.

Unfortunately, the results with the male simulation do not seem to generalize to the simulation of the female runner. We generated ten animations of the female runner and conducted experiments using thirty subjects. Sample poses are shown in Fig. 7 (see Appendix) and the results summarized in Table 2. This distribution is bimodal. Many subjects prefer a slight amount of noise but many subjects also prefer a large amount of noise. This result is perplexing, for although 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.

There are other types of perturbations than those we have considered here. For example, we could model the variability people actually exhibit by analyzing 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. Finally, we have only tested running with these experiments, and broadening our study of variability to other types of motions is important.

References

1. 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.
2. ARUTYUNYAN, G. H., GURFINKEL, V. S., AND MIRSKII, M. L. Investigation of aiming at a target. *Biophysics 13* (1968), 536–538.
3. 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.
4. BECKWITH, T. G., MARANGONI, R. D., AND LIENHARD V, J. H. *Mechanical Measurements*. Addison-Wesley, Reading, Mass., 1993.
5. 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.
6. BRUDERLIN, A., AND WILLIAMS, L. Motion signal processing. In *Computer Graphics* (Aug. 1995), pp. 97–104. Proceedings of SIGGRAPH 95.

7. 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.
8. DUBÉ, R. L. *Natural Pattern Forms*. Van Nostrand Rheinhold, New York, 1997.
9. 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.
10. Found at http://www.cc.gatech.edu/gvu/animation/Areas/variability.
11. 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.
12. JAZWINSKI, A. H. *Stochastic Processes and Filtering Theory*, vol. 64 of *Mathematics in Science and Engineering*. Academic Press, New York, 1970.
13. 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.
14. NEWELL, K. M., AND CORCOS, D. M., Eds. *Variability and Motor Control.* Human Kinetics Publishers, Champaign, IL, 1991.
15. PERLIN, K. Real time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics 1*, 1 (Mar. 1995), 5–15.
16. 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.
17. ROSE, C., COHEN, M., AND BODENHEIMER, B. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Computer Graphics and Applications 18*, 5 (1998), 32–40.
18. 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.
19. SCHLICK, C. Wave generators for computer graphics. In *Graphics Gems V*, A. W. Paeth, Ed. Academic Press, New York, 1995, pp. 367–374.
20. 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.
21. 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.
22. WITKIN, A., AND POPOVIĆ, Z. Motion warping. In *Computer Graphics* (Aug. 1995), pp. 105–108. Proceedings of SIGGRAPH 95.
23. WOODWORTH, R. S. The accuracy of voluntary movement. *Psychological Review 3* (1899), 1–114.
24. 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.

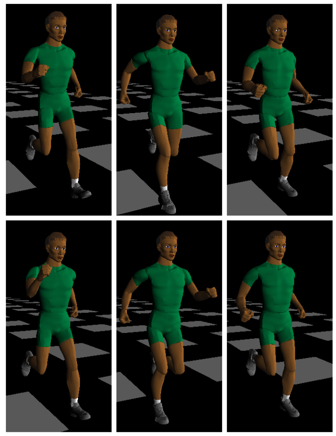

Fig. 6. Example frames of the male simulation with and without noise. The top row shows various poses without noise; the bottom row shows analogous poses with noise (Bodenheimer *et. al.*).

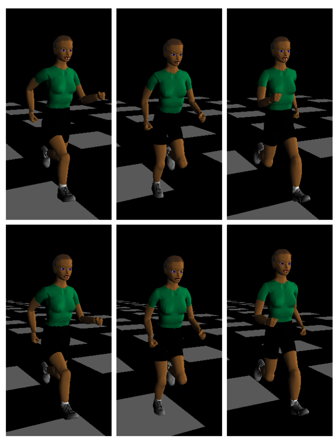

Fig. 7. Example frames of the female simulation with and without noise. The top row shows various poses without noise; the bottom row shows analogous poses with noise (Bodenheimer *et. al.*).