

Obscuring Length Changes During Animated Motion

Jason Harrison

Ronald A. Rensink

Michiel van de Panne

University of British Columbia




Figure 1: The three frames shown above are from an animation where the arms and legs change in length over time. In order to investigate the extent to which changes in length can be obscured, we conducted five sets of experiments with the stimulus shown on the right.

Abstract

In this paper we examine to what extent the lengths of the links in an animated articulated-figure can be changed without the viewer being aware of the change. This is investigated in terms of a framework that emphasizes the role of attention in visual perception. We conducted a set of five experiments to establish bounds for the sensitivity to changes in length as a function of several parameters and the amount of attention available. We found that while length changes of 3% can be perceived when the relevant links are given full attention, changes of over 20% can go unnoticed when attention is not focused in this way. These results provide general guidelines for algorithms that produce or process character motion data and also bring to light some of the potential gains that stand to be achieved with attention-based algorithms.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation; H.1.2 [Models and Principles]: User Interfaces—Human Factors

Keywords: computer animation, change detection, visual attention, length perception, Weber fraction

1 Introduction

Computer animation is about communicating ideas and perceptions through sequences of images. The area of perceptually-based rendering and animation looks at how we can take advantage of known limitations of the observer’s visual perception to better utilize the limited resources available to animate and render the images. Successful applications of this approach include the use of three color primaries to match human color perception, and the use of 24-30 Hz motion blurred frames in film and video as an accepted rate for reproducing smooth motion.

This paper extends this approach into another area of animation: the perception of animated characters formed via rigid links. We

examine what constraints have to hold in order for links to be *perceived* as being rigid during animations. Changes in link length occur in hand-drawn animations and have recently been proposed in algorithms for cleanup of motion capture data [Kovar et al. 2002]. However, to exploit this fluidity effectively, it is important to know what factors are critical in the perception of length. As this paper will show, one of these factors is visual attention.

For static objects, sensitivities for length estimation are well known. Consider the following pair of lines:

Without the use of a ruler you may have difficulty determining that the right line is 2.7% longer than the line on the left. However, you probably will not have difficulty with the following pair:

But, as we will show, in some conditions observers could not reliably tell these two lengths apart. These limitations can be exploited in tools for processing or synthesizing animations, ones that will work in many of the conditions encountered in normal viewing.

Looked at more generally, a number of questions can be asked about the perception of length changes in the links of an animated object: What is the lower bound on the sensitivity of an observer to a change in length of a link?¹ What happens to the sensitivity if an observer is attending to another task in the same area of the visual field? Is the sensitivity to increases in length the same as the sensitivity to decreases? Does sensitivity depend on the duration over which it is carried out?

The main contribution of this paper is a set of experiments that provide answers to these questions. A second contribution, discussed in Section 4, is to show that visual attention is important in seeing change, and that attention must therefore play an important role in any perceptual metrics developed for animation. When combined with predictive models of attention (or other known means of directing attention) this paves part of the path towards new classes of perceptually-based animation algorithms.

2 Experiments

In all of our experiments we used the stimulus illustrated on the right in Figure 1. This two-link articulation was described to the observers as a “top-down view of a human arm resting on a table; the bottom most point is the shoulder, the top-left point is the elbow, and the top-right point is the wrist/hand.” We chose this sim-

¹Sensitivity thresholds are measured as a percentage change in length, thus a low threshold means that the observer is highly sensitive.

ple stimulus rather than a fully-rendered 3D character for our experiments because it introduces fewer confounding factors, has the potential to be the most general, and additional features can easily be added to this model if desired.

In all of our experiments we fixed the shoulder, and moved the wrist/hand along a circular path. The elbow was moved such that the two links formed an articulated pair. We will refer to the arm segment between the shoulder and elbow as the “upper segment,” and the arm segment between the elbow and the wrist/hand as the “lower segment.” Both arm segments were typically 5° in length², the radius of the circular path was 1.5° , and the centre of the circle was offset 5.5° both upwards and to the right from the position of the shoulder. The hand moved along the circular path at a rate of 0.3 to 0.7 rev/sec.

Observers were naive to the goals of the experiments, and were either paid for their time or received credit for an undergraduate psychology course. Stimuli were presented on a CRT with an update rate of 89 Hz. Observers were seated so that the screen subtended a visual angle of 30° horizontally.

We will present all detection thresholds in terms of relative changes of length, specifically the *Weber fraction*. Weber fractions capture the relationship between the size of the estimation error and the size of the stimulus, which is expressed as a proportion. The Weber fraction for the length of lines drawn on paper is 2.9% [Coren et al. 1994]. That is, without the use of a ruler or alignment, for two lines to be reliably perceived as different in length they must differ in length by more than 2.9%.

To study the sensitivity of the visual system to length changes, we varied: the type of change (temporary increase, increase, or decrease); the arm segment(s) changed (upper, lower, or both); the duration over which the change is made; the velocity of movement along the circular path; the number of reversals of motion along the circular path; and the task to be performed while detecting the change in length. We did not vary the size, position, or orientation of the arm, nor did we vary the radius or offset of the circular path.

In all of our experiments we used a two alternative forced choice paradigm and Kaernbach’s “Unforced Weighted Up Down” algorithm [2001] to dynamically set the amplitude of the change in length following each observer’s response. We set the target performance to chance (75% correct). Observers received immediate feedback as to the correctness of their responses. Thresholds were then averaged across observers.

Figure 2 summarizes the thresholds found for differences in length as a function of the conditions of the task. We now describe the specific experiments in more detail.

2.1 Expectation

To determine the effect of expectation on sensitivity to length changes we asked 10 observers in our first experiment to judge which of two presentations of the arm contained a change. During each presentation the hand revolved twice around the circular path at 0.5 rev/sec. The links of the arm increased to length $l + \Delta l$ and immediately returned to length l using linear interpolation within a 500 ms duration on each revolution. Each trial in a block had the change in length occur centered at the same posture. We determined the 80% threshold for detection using PEST [Taylor 1967]

Observers could detect very small changes in length of the arm segments, achieving a threshold of 2.7%. Such a high level of sensitivity may have occurred because the trials were grouped into blocks and feedback was provided so that observers had an expectation about how the change in length would affect the motion of the arm. Note that this threshold is very close to that obtained for static lines, indicating that movement did not affect the ability to detect length changes.

²Lengths are measured in terms of degrees of visual field.

Several studies have shown that expectation about the direction of motion, or other stimulus features, temporarily boosts the sensitivity to that feature [Sekuler et al. 2002]. In order to test the effects of expectation, we recruited another 7 naive observers and presented them randomly interleaved changes of length of the entire arm with only one revolution of the wrist/hand around the circular path in each presentation. This treatment produces thresholds of 5.6%, a worsening of sensitivity by a factor of 2.2 (statistically significant, Kruskal-Wallis, $\chi^2(1, n = 102) = 37.28, p < 0.00001$).

2.2 Task Interference

In a second experiment we required observers to detect changes in length while carrying out an unrelated primary task. The primary task was to count the number of times the wrist/hand changed direction of movement along the circular path, and the secondary task was to detect a length change.

We randomly selected a direction of length change (grow or shrink) and asked observers to report this direction of change only after correctly reporting the count of changes in direction of the wrist/hand. We used randomly interleaved conditions where: the upper, lower, or both arm segments changed length; and velocities were 0.3, 0.5 or 0.7 rev/sec. In each trial, 5 to 9 changes of direction occurred while the arm moved for four seconds (same duration as in the other experiments but duration of change of length was 400-666 ms). We attempted to “mask” the motion artifact induced by the change in length by introducing it as a linear increase or decrease from one randomly chosen wrist/hand change of direction to the subsequent change of direction.

In order that the two directions of change of length be comparable, we presented changes in length between extremes l and $l + \Delta l$ rather than between l and $l \pm \Delta l$. In both cases we computed Weber fractions for the detection thresholds as $\Delta l/l$.

The effect of the primary task is remarkable. In comparison to the threshold found using blocked trials (2.7%), we found a threshold of 10.5% for velocities of 0.5 rev/sec, an increase of almost 300% (statistically significant, Kruskal-Wallis, $\chi^2(1, n = 72) = 33.60, p < 0.00001$). Slower velocities made changes slightly easier to detect, 9.2% for 0.3 rev/sec and 10.7% for 0.7 rev/sec (Pearson correlation between average thresholds and \log_{10} velocities is $r_p = 0.95$). Unexpectedly, we also found a significant difference between the thresholds for the upper and lower arm segments: 10.7% and 14.5% respectively ($n = 14$, Wilcoxon signed-rank test, $p < 0.03$). The source of the difference in thresholds is probably due to the fixed position of the shoulder which allows the observers to make a more accurate judgement.

2.3 Increases vs Decreases in Length

Another important issue is whether observers are more sensitive to increases in length than to decreases. We asked observers to detect changes in length, with direction of change of length as an experimental condition, a primary task of counting changes of direction of the wrist/hand (as in the previous experiment), and changes made over four seconds. Observers were allowed to report the direction of change in length only if correctly reporting the count of wrist/hand direction changes.

This condition produced a striking difference in sensitivity; for growing arms, the threshold is 6.9%, whereas for shrinking it is 19.7%, a much higher value ($n = 12$, Wilcoxon signed-rank test, $p < 0.0005$)³.

³We have replicated this result with an additional 34 observers, sensitivity to growing, 8.0%, was much better than sensitivity to shrinking, 18.4% ($n = 46$, Wilcoxon signed-rank test, $p < 0.0001$). For three observers, sensitivity to shrinking was slightly better than sensitivity to growing.




Figure 2: Average thresholds of detection (Weber fraction) for changes in length of an animated articulation as a function of the experimental conditions summarized in the table on the right. Error bars are ± 1 standard error.

2.4 Increases vs Decreases and Divided Attention

Since animated displays rarely consist of only one object, it is important to examine what happens when attention is divided between objects. Thus, in this experiment we added a very small, low contrast distractor to the display and used the same 13 observers from the last experiment. Their primary task was counting changes of direction of the distractor and their secondary task was reporting the direction of change in length. The distractor moved around a circular path in the same fashion as the wrist/hand, and made 5-9 changes of direction during the four second presentation. We placed the distractor in two positions, one “overlapping” the path of the wrist/hand, and the second “disjoint,” 11° below the wrist/hand to additionally test the effect of eccentricity.

The addition of the distractor further decreased sensitivity for slowly-growing arms, from 6.9% to 12.7% for the overlapping condition, and a further increase to 17.6% for the disjoint condition ($n = 13$, Wilcoxon signed-rank test, $p = 0.064$ and $p < 0.0001$ respectively). This is likely due to eccentricity (i.e., visual acuity) or the differing attentional resolutions between the upper and lower visual fields [He et al. 1997]. Sensitivity for slowly-shrinking arms was 19.4% for the overlapping distractor and 23.3% for the disjoint distractor ($n = 13$, Wilcoxon signed-rank test, $p = 0.97$ and $p = 0.09$ respectively). Both these values did not differ significantly from the value obtained without distractors (19.7%), suggesting that a limit of some kind may have been reached.

2.5 Duration of Length Change

A potentially important factor is the duration of time over which the change is made. To see if this is the case, we had two groups of 14 observers count the changes of direction of the disjoint distractor while detecting changes of length (as in the previous section). We randomly interleaved trials with durations of 50, 125 and 250 ms, or 500, 1000, or 2000 ms. As seen in Figure 3, which includes data from the 4000 ms duration condition from the last section, thresholds increase with duration of change, and shrinking allows larger changes than growing across a range of 50 to 4000 ms. Pearson correlations between average thresholds and \log_{10} duration of changes are $r_p = 0.89$ for growing, and $r_p = 0.99$ for shrinking, indicating a strong linear relationship between log duration and threshold of detection ($p < 0.0075$ for both). Interestingly, even at very small durations (50 ms), thresholds can still be as high as 10%.

Summary of Experimental Conditions and Results

Expectation, Section 2.1

- 1A Trials blocked on posture of change, $\bar{x} = 2.7 \pm 0.3\%$
- 1B Randomly interleaved trials, $\bar{x} = 5.6 \pm 0.4\%$

Task Interference, Section 2.2

- 2A Upper arm changes, $\bar{x} = 10.7 \pm 1.4\%$
- 2B Lower arm changes, $\bar{x} = 14.5 \pm 0.5\%$
- 2C Entire arm changes, $\bar{x} = 10.5 \pm 0.8\%$

Increases vs Decreases with Task Interference, Section 2.3

- 3A Increases in length, $\bar{x} = 6.9 \pm 1.4\%$
- 3B Decreases in length, $\bar{x} = 19.7 \pm 1.7\%$

Increases vs Decreases and Divided Attention, Section 2.4

- 4A Distractor overlapping, increases in length, $\bar{x} = 12.7 \pm 2.4\%$
- 4B Distractor disjoint, increases in length, $\bar{x} = 17.6 \pm 2.4\%$
- 4C Distractor overlapping, decreases in length, $\bar{x} = 19.4 \pm 2.5\%$
- 4D Distractor disjoint, decreases in length, $\bar{x} = 23.3 \pm 1.3\%$

3 Guidelines

The results of our experiments point to the following guidelines for obscuring length changes during animation: (1) Length changes of up to 2.7% will likely go unseen even when the length change is attended to and a change of length is expected; (2) Length changes should never be greater than 20%, even if the viewer is not attending to the figure; (3) Decreases in length are less noticeable than increases; (4) Slower changes are harder to detect; (5) Changes are more difficult to perceive during fast motions; (6) Having the viewer carry out a different task, even within the same figure, is an effective way to mask changes; (7) An expectation of a particular type of change results in a significantly greater likelihood of detection. (Please refer to Section 2 for the specific thresholds involved.)

Beyond the manipulations used in our experiments, it may also be possible to obscure length changes in other ways. For example, when using a rendered 3D character there is greater uncertainty about where the segment ends are, which may interfere with length estimates. It may also be possible to take advantage of the ambiguity between a length change or a 3D out-of-plane rotation. Since figure translations have been found to obscure length changes [Tomat et al. 1999], some camera motions, such as those along the direction of length change, may also allow the change to be obscured.

4 Discussion

The experiments described above have provided an interesting set of results concerning the degree to which length changes in animated figures can be obscured. Although these results can be put into the form of individual guidelines, a more powerful framework is possible based on the more general point that what is needed to see change is visual attention [Rensink 2002]. To see this consider: i) less attention is given to unexpected events than expected ones; ii) attending to one object or task reduces attention on the others [Simons and Chabris 1999]; iii) length decreases do not attract attention, whereas length increases do [Franconeri and Simons 2003]; and iv) the slower the change, the weaker the motion signal, and the less it draws attention [Simons et al. 2000].

Thus, by manipulating the salience of the change and the amount of attention available to the task, we were able to obscure larger and larger degrees of length change in a moving articulated figure. This is just one demonstration of the need for attention to detect changes




Figure 3: Average thresholds of detection for changes in length as a function of direction and \log_{10} duration of change when observers count the changes of direction of a disjoint very small, low contrast distractor. Error bars are ± 1 standard error.

in objects. This limitation has recently become a focus of research in vision science and can become a key element of perceptually-based animation algorithms. The extent to which these algorithms will be effective will be the extent to which they can incorporate knowledge of how perception operates. In particular, such algorithms will be effective to the extent that they can incorporate a knowledge of visual attention.

4.1 The Nature of Visual Attention

Recent work in vision science has resulted in a number of striking demonstrations that show the severe limitations of human vision in the absence of attention. Specifically, observers can be “blind” to unexpected stimuli that appear while they are monitoring the movements of other objects. For example, observers will miss the sudden appearance of a person in a gorilla suit, even when this person parades through the middle of the scene [Simons and Chabris 1999]. Observers will also often fail to see changes that are large and repeatedly made, even when they know that changes are present somewhere in the display [Rensink 2002].

Cognitive scientists are increasingly coming to the conclusion that much of the reason for these failures is due to limitations on visual attention. Rather than visual attention being “unlimited, everywhere, and quick,” as might be believed based on subjective impressions, the visual system appears to work very hard with limited resources to detect changes in scenes [Rensink 2002], track the movement of objects [Pylyshyn and Storm 1988], and discriminate even the simplest dynamic patterns [Cavanagh et al. 2001]. In other words, our perception of the world and the events in it is a construct built using a system with severe limitations. Limitations exist on spatial and temporal resolution [He et al. 1997; Verstraten et al. 2000], and speed of movement from object to object [Moore and Wolfe 2001].

4.2 The Control of Visual Attention

Given the above, the most effective perceptually-based algorithms would be those that could take advantage of knowledge about when and where the viewer sends their attention. Ideally, a computational model would take a description of the animation and the goals (if

any) of the viewer, and produce a prediction of which objects in the scene would be attended to, and when (at least for the near future). However, creating such a model is difficult, not least because verifying it must rely on a large amount of data from human observers [Horvitz and Lengyel 1997].

But progress has been made. To begin with, two forms of attentional control have been discovered: one based on high-level goals (or interests) that sends attention to particular objects, and one based on low-level visual factors (saliency) that draws attention to particular locations. In regards to the former, it is known that attention is first sent to those objects that the observer considers interesting [Rensink 2002], because, for example, the object is the goal of a particular action (e.g., it is a target), or because it is interesting in its own right (e.g., is a human). Computational models of low-level saliency are under constant development. For example, a recent model can predict up to 85% of the fixations that observers will make when searching for people in a scene [Oliva et al. 2003]. (See Mozer and Sitton [1998] for a broader review of computational models of spatial attention than we can provide here.)

4.3 Relation to Prior Work

The framework proposed here can cast new light on prior work, indicating where it might be expected to succeed, and where it might fail. For example, O’Sullivan and Dingliana [2001] measured the ability of observers to detect anomalous collisions as a function of eccentricity from fixation point and number of distractors moving on the display. Their motivation for investigating the effect of eccentricity lies in the observation that the central 5° of the visual field is disproportionately represented from the retina through to the visual cortex [Coren et al. 1994]. Recent studies of the spatial resolution of visual attention [He et al. 1997] potentially extend O’Sullivan and Dingliana’s findings across the visual field to denser displays, where observers would not be able to track even a single object among “distractors.” The framework proposed here would predict that in conditions where observers are unable to track objects, anomalous collisions and changes in the “rules of physics” [O’Sullivan et al. 2003] would be undetectable.

Kovar et al. [2002] presented an algorithm to remove “footskate” artifacts from motion capture data. An element of their algorithm is the adjustment of the length of the leg in order to avoid unnaturally fast extensions/contractions of the knee when the leg is nearly straight and the foot position needs to be adjusted. Our results provide bounds on leg length adjustments as a function of the observer’s attentional focus.

Reitsma and Pollard [2003] performed psychophysical experiments to determine perceptual based metrics for measuring errors in ballistic human motion. Our results indicate that the use of a metric that does not take expectation and attention into account may over- or underestimate the sensitivity of observers.

5 Conclusions and Future Work

We have carried out five sets of psychophysical experiments to establish the sensitivity to changes in length as a function of expectation, task interference, direction of change, division of attention, and duration. The implications of these results are summarized in the guidelines given in Section 3. These results have implications for inverse kinematics and motion processing algorithms and may also be useful for computer animators.

We plan to replicate our results with experiments using more realistic three-dimensional renderings of characters as observers have been found to be more sensitive to differences in motions when the characters were rendered more realistically [Hodgins et al. 1998]. It is not known if this increase in sensitivity is limited to the task

of comparing movements, which requires selecting features and encoding them into short term memory for later comparison. In any case, we plan to determine how strong the effect of rendering style is relative to the manipulations of attention used in our experiments.

There are also a number of other interesting issues to consider. Does the perceived intent of a movement obscure large length changes? For example, the lengthening of an arm while reaching for an object that is just beyond reach versus an arm swinging while walking? Will changes in the rigidity of the articulation affect high-level judgements such as the gracefulness or age of the character? It has been demonstrated that increasing the animation level-of-detail of simulated soccer players increased observers' ratings of team skill level—even though the manipulation was not consciously noted by the observers [Oesker et al. 2000]. Thus, will supra or sub-threshold changes to an articulation induce shifts in high-level judgements?

Our experimental method can be extended to more complex articulations. Several fundamental perceptual issues, such as the nature of perceptual encoding of translating and rotating objects are being addressed by other researchers [Gysen et al. 2002]. These results all point to the critical role of attention for the perception of movement, in particular the need for attention to see change. Consequently, as the number of links increases, there are more places to attend, and thus more possibilities for hiding changes.

Perhaps the most general implication of our results is the clear need to “pay attention to attention.” Research into human perception shows that the constructs that represent unattended stimuli are much less descriptive than we think. In particular, changes of objects cannot be seen without attention. A framework has been proposed for the incorporation of such research into the creation of perceptually-based algorithms that can exploit these limitations to compute animations at a lower computational cost than algorithms which do not relax geometric or physics-based constraints.

Perceptually-based animation algorithms will be effective to the extent that they can take advantage of the characteristics of the human visual system, in particular, the characteristics of visual attention. The guidelines proposed here are another step towards this goal. And as our understanding of perception improves, such algorithms will be able to provide an increasingly realistic experience at an increasingly lower cost.

6 Acknowledgements

Thanks to David Loran who collected the data, Shirley Chen who created the character animations used in the video, and the UBC Visual Cognition Lab where the experiments were conducted. Funding was provided by Nissan Motor Company Limited, the Natural Sciences and Engineering Research Council of Canada, and the Canada Foundation for Innovation. Thanks also to the observers and to the anonymous reviewers for their helpful comments on the presentation and content of this paper.

References

CAVANAGH, P., LABIANCA, A. T., AND THORNTON, I. M. 2001. Attention-based visual routines: sprites. *Cognition* 80, 1-2, 47–60.

COREN, S., WARD, L. M., AND ENNS, J. T. 1994. *Sensation and Perception*, 4th ed. Fort Worth: Harcourt Brace.

FRANCONERI, S. L., AND SIMONS, D. J. 2003. Moving and looming stimuli capture attention. *Perception & Psychophysics* 65, 7, 999–1010.

GYSEN, V., DE GRAEF, P., AND VERFAILLIE, K. 2002. Detection of intrasaccadic displacements and depth rotations of moving objects. *Vision Research* 42, 379–391.

HE, S., CAVANAGH, P., AND INTRILIGATOR, J. 1997. Attentional resolution. *Trends in Cognitive Science* 1, 3 (June), 115–120.

HODGINS, J. K., O'BRIEN, J. F., AND TUMBLIN, J. 1998. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics* 4, 4, 17–25.

HORVITZ, E., AND LENGUEL, J. 1997. Perception, attention, and resources: A decision-theoretic approach to graphics rendering. In *Proceedings of the 13th Conference on Uncertainty in Artificial Intelligence*, Association for Uncertainty in AI (AUAI), 238–249.

KAERNBACH, C. 2001. Adaptive threshold estimation with unforced-choice tasks. *Perception & Psychophysics* 63, 8, 1377–1388.

KOVAR, L., SCHREINER, J., AND GLEICHER, M. 2002. Footskate cleanup for motion capture editing. In *ACM SIGGRAPH Symposium on Computer Animation*, ACM SIGGRAPH, 97–104.

MOORE, C. M., AND WOLFE, J. M. 2001. Getting beyond the serial/parallel debate in visual search: a hybrid approach. In *The Limits of Attention: Temporal Constraints in human Information Processing*. Oxford: Oxford University Press, ch. 9, 178–198. Kimron Shapiro (Ed.).

MOZER, M. C., AND SITTON, M. 1998. Computational modeling of spatial attention. In *Attention*, H. Paschler, Ed. London: UCL Press, 341–393.

OESKER, M., HECHT, H., AND JUNG, B. 2000. Psychological evidence for unconscious processing of detail in real-time animation of multiple characters. *J. Visualization and Computer Animation* 11, 105–112.

OLIVA, A., TORRALBA, A., CASTELHANO, M. S., AND HENDERSON, J. M. 2003. Top-down control of visual attention in object detection. In *Proceedings of the IEEE International Conference on Image Processing*, vol. 1, 253–256.

O'SULLIVAN, C., AND DINGLIANA, J. 2001. Collisions and perception. *ACM Transactions on Graphics* 20, 3, 151–168.

O'SULLIVAN, C., DINGLIANA, J., GIANG, T., AND KAISER, M. K. 2003. Evaluating the visual fidelity of physically based animations. In *Proceedings of ACM SIGGRAPH 2003*, ACM Press / ACM SIGGRAPH, 527–526.

PYLYSHYN, Z. W., AND STORM, R. W. 1988. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision* 3, 179–197.

REITSMA, P. S. A., AND POLLARD, N. 2003. Perceptual metrics for character animation: Sensitivity to errors in ballistic motion. In *Proceedings of ACM SIGGRAPH 2003*, ACM Press / ACM SIGGRAPH: New York, ACM SIGGRAPH, 537–542.

RENSINK, R. A. 2002. Change detection. *Annual Review of Psychology* 53, 245–277.

SEKULER, R., WATAMANIUK, S. N. J., AND BLAKE, R. 2002. Perception of visual motion. In *Stevens Handbook of Experimental Psychology*, 3rd ed., vol. 1. New York: Wiley, ch. 4.

SIMONS, D. J., AND CHABRIS, C. F. 1999. Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception* 28, 1059–1074.

SIMONS, D. J., FRANCONERI, S. L., AND REIMER, R. L. 2000. Change blindness in the absence of visual disruption. *Perception* 29, 1143–1154.

TAYLOR, M. M. 1967. Pest: Efficient estimates on probability functions. *Journal of the Acoustical Society of America* 41, 4, 782–787.

TOMAT, L., SCAMARDI, P., AND VICARIO, G. B. 1999. The perception of length changes in moving objects. Presented at ECVP'99.

VERSTRATEN, F. A., CAVANAGH, P., AND LABIANCA, A. T. 2000. Limits of attentive tracking reveal temporal properties of attention. *Vision Research* 40, 3651–3664.