Collisions and Attention

CAROL O'SULLIVAN
Trinity College Dublin

Attention is an important factor in the perception of static and dynamic scenes, which should, therefore, be taken into account when creating graphical images and animation. Recently, researchers have recognized this fact and have been investigating how the focus of attention can be measured, predicted, and exploited in graphical systems. In this article, we explore some preliminary strategies for developing an automatic means of predicting and exploiting attention in the processing of collisions and other dynamic events. Recent work on the perception of causality has shown that attention can change the way in which a dynamic scene consisting of collision events is perceived. We describe a series of experiments designed to determine the source of biases in the perception of anomalous collision dynamics and, in particular, whether attention plays a role. Using an eye-tracker, eye-movements were recorded while participants viewed animations of simple causal launching events in 3D involving two colliding spheres. Results indicated that there was indeed a definite pattern to the allocation of attention based on the nature of the event, which is promising for the goal of developing a predictive metric. As a follow-up, a paper-based experiment was carried out in which participants were asked to sketch the predicted post-collision trajectories of the same two spheres printed on paper. These experiments demonstrated that attention alone was not sufficient in determining performance, but rather the *nature* of the dynamic event itself also played a role.

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

General Terms: Algorithms, Experimentation, Human Factors, Measurement

Additional Key Words and Phrases: Animation, collisions, perception, eye movements, attention

---

1. INTRODUCTION

Perceptual issues relating to computer animation have been receiving increasing interest in recent years. As methods and technology develop rapidly, greater realism is now possible in all forms of visual media and the need to optimize and evaluate the realism of images and animations is becoming ever more important. Human visual perception of motion and dynamic events is a complex cognitive process, which is still not fully understood. Therefore, studies that provide insights into motion and event perception through the use of synthetic imagery and animation can be extremely useful, with respect to furthering the development of both the graphics and perception fields. In this article, we consider the specific problem of perception of collision events. In the real world, our perception of dynamic and causal events is strongly dependent on the way in which we perceive objects touching or colliding and on the way that they behave postcollision. In graphics, collision handling is a core part of any animation system and consumes a large proportion of the computation time for each step of a simulation. Hence, maximizing the realism and efficiency of this particular aspect of animation is an area of active research, both from the computational and perceptual perspectives.

Author's address: Computer Science Dept., Trinity College Dublin, Dublin, Ireland; email: Carol.OSullivan@cs.tcd.ie.

Much recent research, outlined in Section 2, has shown that attention is one of the most important and interesting features of human perception, which can also be exploited for improving the efficiency and realism of graphical systems. The underlying principle of this work is: If the focus of attention can be measured or predicted, then computational resources can be concentrated on rendering or simulating those parts of the depicted scenes that will most contribute to the user's perception. One option is to use an eye-tracker to directly determine the fixation point of the viewer. Modern eye-trackers are becoming increasingly usable and their employment in the field of human computer interaction is becoming more prevalent [Duchowski 2003]. However, as eye-tracking may not always be either feasible or desirable, it is often necessary to find other means of predicting gaze and attention (which may not necessarily be the same, see Hunt and Kingstone [2003]).

In Section 2, we will review some situations in which knowledge of the focus of attention has been effectively exploited in some way, to improve graphical displays. Up to now, this approach has mainly been applied to graphical user interfaces, or to rendering for images, or precomputed animations. Strategies for predicting attention have either been based on bottom-up processing (e.g., using preattentive pop-out features, such as color or velocity), or have required user-defined lists of semantically important objects to be defined. The *nature* or *semantics* of the objects and events depicted have not really been taken into account in any generic or automatic way. In this article, we show that it is indeed sometimes possible to determine the focus of attention in scenes depicting dynamic collision events, based on the nature of the collision events themselves. Specifically, we show that in events in which a stationary ball is struck by a moving one, people *attend to* the struck ball more after the collision, while they are less interested in the postcollision behavior of the striking ball (see Section 3). Eye-tracking is used to evaluate this hypothesis. Furthermore, using a set of paper-based experiments (see Section 4), we examine whether attention alone is sufficient to explain previously recorded biases in people's ability to perceive dynamic anomalies in such collision events.

2. BACKGROUND

2.1 Graphics and Attention

Baudisch et al. [2003] discuss "*attentive displays*" that use the focus of attention to concentrate resolution or rendering effort on particular regions of a display, thus reducing overall computation and memory costs. Such techniques typically exploit the fact that an individual viewer can only focus on a small spatial region of a display at any time, with fixation being determined by an eye-tracker. Gaze-contingent displays degrade the resolution in the peripheral image regions. The high-resolution area moves with the user's focus, so the area under scrutiny is always rendered at a higher resolution. An imperceptible degradation is difficult to achieve, but often in visual search tasks, the reduction in quality has no effect on performance even if it is noticeable. In terms of interactive graphics systems, attentive three-dimensional (3D)-rendering engines use a viewer's gaze position to vary the level of detail (LOD) at which an object is rendered at the object geometry rather than the image level (see also Luebke et al.)

Rather than tracking a user's gaze explicitly, some researchers have used a model of visual attention to *predict* fixation locations instead. For example, Yee et al. [2001] used Itti et al.'s [1998] bottom-up attention model to accelerate global illumination calculations in precomputed animations. A saliency map was used to predict the focus of attention while viewing dynamic scenes, thus eliminating the need for an eye-tracker. Parafoveal degradation of the peripheral regions was then implemented, which led to a significant speedup in the production of the animation. However, after several repeat viewings of the animation, they found that people started to look away from these predicted regions to explore the field of view (FOV) more fully. Consequently, in interactive applications such as games, where a 3D virtual environment is navigated and explored in real time, it may not be possible to predict where a user might look using such a bottom-up model. In fact, some researchers have shown that this model does not accurately predict human fixations in such a virtual reality environment. Using a string editing methodology, Marmitt and Duchowski [2003] compared real (human) and artificial [Itti et al. 1998] scanpaths in a virtual reality setting and found that the attentional model predicted fixations over a wide area of the display, whereas real observers concentrated more on the center.

Clearly, in order to predict attention for more interactive systems, top-down factors must be taken into account. One approach is to consider the *task* that the user is performing. Concentration on a task can give rise to *inattentional blindness* [Mack and Rock 1998], i.e., the failure to see an object because attention is not focused on it *even if it is within the foveal region of the viewer*. Simons and Chabris [1999] demonstrated this in a particularly compelling fashion, where 50% of people told to count the number of players' passes in a videoed basketball game failed to notice a gorilla walking across the court! Humans clearly have limited attentional resources, which are usually focused on the task at hand. Cater et al. [2003] take some first steps toward predicting this allocation strategy, which could be exploited in order to focus rendering or simulation effort on task-specific regions, objects, or events. When performing a counting task on a still image rendered at three resolutions: high, low, and selective (where only the target objects were rendered at high resolution and the rest of the scene was at low), subjects consistently failed to notice any difference between the high- and the selective-quality image. Furthermore, subjects could easily tell the difference when there was no task involved. This demonstrates that subjects primarily attend to task-related objects and the authors postulate that such objects can often be identified in advance, depending on the task. They, therefore, conclude that task-driven attention can override low-level visual attention when it comes to noticing rendering artifacts. However, this approach requires a user-defined list of semantically important objects, which may not be possible for more complex tasks.

The above methods are focused, to the most part, on reducing the rendering complexity of unattended parts of graphical displays or objects, by degrading either display resolution or geometric complexity. However, attention is also important when attempting to reduce the computational effort required to generate the *motions* of simulated objects. Examples of this approach include Carlson and Hodgins' [1997] system, where simplified simulation rules are used for legged creatures that are not central to a game; or Chenney and Forsyth's [1997] system, where the dynamics of objects outside the view frustum were calculated at a lower level of accuracy. In quite a simple way, attention was used in both of these cases to determine the importance of simulated objects and their motions: the former by using task semantics and the latter using visibility. Harrison et al. [2004] also show that attention is crucial in order to detect anomalies in a situation where the lengths of links in an articulated structure may be changed during an animation. Such situations may arise when an animator edits a human motion in order to meet some specific constraints. Harrison et al. showed that, if attention can be captured by a task, quite significant changes in link extents go unnoticed. However, apart from the simple criteria discussed above, it is not entirely clear how attention could be determined in some generic or automatic way for these and other simulation scenarios. We propose that much more could be possible by examining the influence of the dynamic properties of the events themselves on the capture of attention.

2.2 Collisions and Attention

Collision is a causal event, i.e., a situation in which one action triggers another. For example, when one ball moves toward a stationary one, which, in turn, immediately starts to move upon impact, the second ball's motion is clearly perceived to be *caused* by the collision (also referred to as a *causal launch*: see Figure 1, top). Scholl and Tremoulet [2000] provide a comprehensive overview of recent research on perceptual causality, which began with the work of Michotte [1946]. Under certain circumstances,

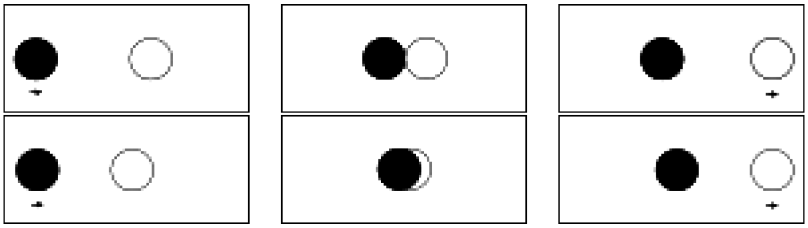

Fig. 1. (*Top*) A typical causal launching event; (*Bottom*) a typical noncausal passing event.

however, the perception of causality can be reduced. For example, the introduction of a time delay at the point of collision can induce the illusion of objects moving under their own volition. Partial or full overlap of the two objects at the point of collision means that the two objects are not perceived as colliding and moving apart, but rather as passing each other out (called a *noncausal passing*; see Figure 1, bottom). This occurs even if the objects are of different color, where it seems as if the moving object takes on the appearance of the static one during the passing event. However, Scholl and Nakayama [2002] have shown that *causal capture* can occur. By showing a causal launch event alongside a noncausal passing event, the perception of causality for the noncausal event can be strengthened. Furthermore, Choi and Scholl [2004] have shown in recent work that the proximity of the events, the way in which objects are grouped or linked, and the *focus of attention* can also either enhance or attenuate perceived causality. These research results show that attention is clearly an important factor in collision perception and, if we can determine how attention is allocated during the collision event itself, this can actually affect perception of the entire dynamic scene and not only the specific collision being viewed.

Another important reason for determining the focus of attention in scenes depicting *physical* simulation scenarios is to help mask any potential artefacts that may be visible. Such dynamic anomalies can occur in computer animation for a number of reasons. *Aesthetic* distortions can be introduced intentionally by the animator in order to achieve a certain end result or to satisfy contraints, while *unavoidable* distortions can be caused by reducing, in some way, the accuracy of the simulation. In our previous work, we investigated the perceptual impact of anomalous collision dynamics on the perception of the viewer. In [O'Sullivan and Dingliana 2001], we described some experiments that showed how certain factors can affect the perception of collisions, namely:

- *Separation:* how far the objects are apart when they collide is a particularly noticeable anomaly, which is less obvious with decreasing visibility of the point of collision,
- *Eccentricity:* the degree to which the collision occurs in peripheral vision affects the ability to notice spatial artefacts, such as separation,
- *Physical accuracy:* surprisingly enough, the amount of error in the calculation of physical response, while important, was not always the dominant factor in anomalous collision identification. In fact, it was possible to ameliorate this effect by adding random spin after collision (thus echoing and Gilden and Proffitt's [1989] earlier findings that people used heuristics to determine the accuracy of collision events).
- *Causality:* the introduction of a time delay at the moment of collisions had a strongly degrading effect on the perception of causality, as in previous research [Scholl and Tremoulet 2000].
- *Distracters:* other nontask objects moving in the scene had a significant effect on people's ability to notice anomalous collisions, in particular, when the distracters were visually similar to the colliding entities.

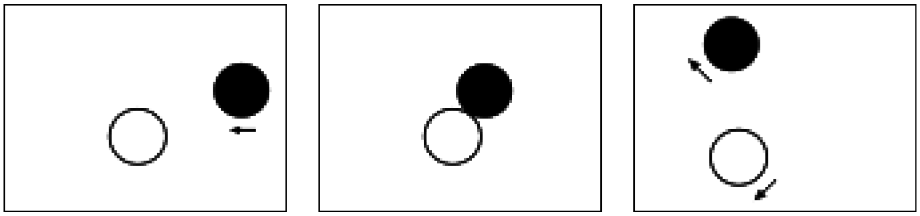

Fig. 2. A depiction of the type of causal events shown in Kaiser and Proffitt [1987] and O'Sullivan et al. [2003].

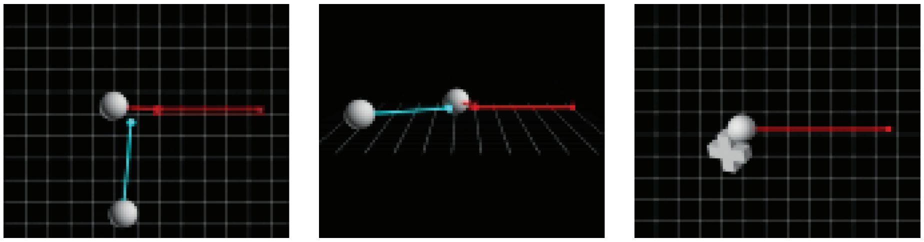

Fig. 3. Screenshots of the causal events shown in earlier studies [O'Sullivan et al. 2003], with the paths of the objects shown in red (striking) and blue (struck).

In order to build on these results and to establish some thresholds for the perception of dynamic anomalies in animation, we ran a further set of psychophysical studies [O'Sullivan et al. 2003]. Emulating some earlier experiments in two dimensions by Kaiser and Proffitt [1987] (see Figure 2), we displayed a series of simple causal launching effects in three dimensions (see Figure 3), in which a stationary ball was hit by an initially moving one. The postcollision trajectories of both balls were distorted in a number of ways, by manipulating the angles so that they moved on expanded or contracted paths. Momentum distortions and spatiotemporal distortions were also investigated. Some interesting biases were identified that indicated nonsymmetric responses to these distortions. For example, expansion of the angle between postcollision trajectories was preferred to contraction and increases in velocity were preferred to decreases. In particular, distortions to the striking ball's trajectory were accepted more extensively than to the struck balls. Some major questions remain to be answered before we can extrapolate from these results to more general cases involving multiple moving objects with complex dynamical properties. The important role of distracters and eccentricity revealed in our earlier study [O'Sullivan and Dingliana 2001] clearly points to a strong attentional effect in collision perception, even when all collisions are visible in the scene. In this article, we start this investigation by posing two specific questions: (1) Were the lower thresholds for perception of distortion to the struck ball's trajectories due to attention, i.e., were people *looking at* the struck ball while making their decision as to the accuracy of the collision event? (2) . . . and, if so, is attention enough to explain performance in this scenario or does the *nature* of the collision event itself play a role?

3. EYE-TRACKING EXPERIMENTS

Although there is an ongoing debate in the eye-movement literature as to whether attention can always be accurately measured by the location of eye fixations, with some claiming that *covert* attention can also be assigned to objects that are not currently being foveated (see Hunt and Kingstone [2003] for a brief review), eye tracking is nevertheless a powerful tool in assessing the focus of attention. We recorded the eye movements of participants while viewing similar events to those depicted in O'Sullivan et al.

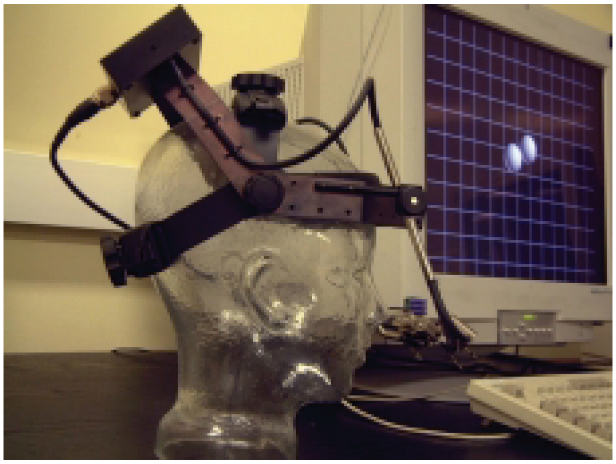

Fig. 4. Eye movements were recorded while participants viewed animations of simple collision events.

Fig. 5. Eye fixations (in red) for animated one-moving experiments (five participants): simulations as in O'Sullivan et al. [2003]. Blue shows the trajectories of the struck ball B and green the trajectories of the striking ball A.

[2003], in a pilot study we used the SMI Eyelink I 250 Hz eye-tracker and, in a second, more complete study, we used the SR Research Eyelink II 500 Hz device [see Figure 4].

In the first study, five participants viewed a series of collision events in which the angles of the post-collision trajectories were both distorted. Two of the participants had participated in previous experiments. The striking ball **A** hit the initially stationary struck ball **B** either from above or below and varying amounts of distortion were added to the trajectories at each viewing of the event. Participants were requested to indicate whether the event was natural or unnatural and several randomly interleaved staircases were carried out, as described in our earlier paper. The resulting trajectories with overlaid fixations are shown in Figure 5, and the first 4000 observations were recorded for each participant. A metric was needed to estimate which set of trajectories received the most attention. Initially, we used the average distance of the point of fixation to the center of each ball, calculated in world-space at each frame of the animation. However, it became clear that this was biasing in favor of the striking ball, as the struck ball travelled a longer distance due to the angle at which it was being hit and conservation of momentum. Hence, we divided the distance to the struck ball at each frame by the difference in trajectory length, which provided a fairer estimate. For the simple cases depicted in the pilot experiments, attention was clearly focused on the struck ball for each participant. We determined if this difference in mean distances from each ball was significant for each participant by performing a single-factor Analysis of Variance (ANOVA) on each participant's set of 4000 pairs of observations, i.e., $F[1, 7998] = 198, 571, 27, 209, 122$, respectively, $Fcrit = 3.84$, $p = 0$. The null hypothesis is rejected when the F statistic is greater than the critical value Fcrit, and the probability p is less than 0.05 (for 95% significance).

Despite this seeming affirmation of our postulation that the struck ball was attended more, several issues raised some doubts. First, participants seemed to focus almost exclusively on the upper half of the display and only one type of event was depicted, i.e., the striking ball always came from the right and an equal amount of distortion was applied to each ball's postcollision trajectory simultaneously. Hence, we decided to run a more complete set of experiments where a greater range of event types could be viewed.

In the second study, seven naive participants aged 20 to 25 viewed a series of collision events while wearing the Eyelink II eye-tracker. None had taken part in any previous experiments and all were undergraduate or graduate students from the author's institution. All had self-reported normal or corrected-to-normal vision. A full 12-point calibration process was carried out and the experiments only started once good accuracy was achieved. During intervals, a fixation dot was presented on the screen and drift correction was carried out by the eye-tracking system. Data from two participants were subsequently eliminated because of gaps in the data, probably due to excessive blinking or temporary losses of calibration. There was no loss of data for the remaining five participants.

The display and stimuli were as described in O'Sullivan et al. [2003], but this time the striking ball did not only appear from the right-hand side of the display. Instead, it appeared from the right, left, top or bottom of the screen. In each of these cases, it hit the struck ball on the left, right, bottom or top, giving rise to eight possible "directions" (see Figure 6). In each case, the collision point was positioned randomly at a point in the central region of the screen, so that the location of the collision event was different during each interval. All eight directions were displayed as a set of randomly interleaved staircases, one ascending and one descending. Since we were not interested in the results this time, we simply cut the staircase short after two repetitions of each condition. Again, participants indicated natural or unnatural events by pressing a button on a game pad. We had two conditions for each staircase: distortion of the striking ball A only and distortion of the struck ball B only. Therefore, there were 32 staircases in total and two viewings of each, some of the natural case and some of the distorted case, giving rise to a total of 32000 observations per participant. We show the resultant trajectories with overlaid fixation points from one participant for each direction (8000 observations each) (see Figure 6). All results for this participant were statistically significant (again, single Factor ANOVAs were calculated, with $F(7998, 1) > Fcrit$ and $p \approx 0$ for each direction) the same was true for the other four participants in all cases.

We looked at all results for each participant, regardless of direction, and found that the struck ball B was indeed looked at more overall than the striking ball A (see Figure 7). ANOVA results were significant for all participants: $F(63998, 1) > 1000, P \approx 0$ in each case. However, upon closer examination it emerges that the case is not quite as black and white as these results suggest. By examining the results collapsed over each direction, we can see that for certain directions, the participants either did not exhibit a preference for either ball, or actually preferred looking at the *striking* ball (see Figure 8). This is particularly curious in the case of the ball coming from the right and hitting the struck ball on top (i.e., right top), as this was one of the situations depicted in O'Sullivan et al.'s [2003] earlier studies. It cannot be due to prior experience, as none of these participants had previously done this task. Differences in mean responses between participants under all conditions were not significant. Using a different metric, where the distance from the fixation point to the line of each ball (rather than to the ball itself) was measured, did not provide significantly different results.

We investigated whether further clues as to the reasons for these rather unexpected results could be found by looking at the individual results broken down further by anomaly type. For the participant whose data is shown in Figure 6, for example, there is no difference in attention depending on which ball was being distorted in any of the cases where the struck ball was attended more, but far more confusion

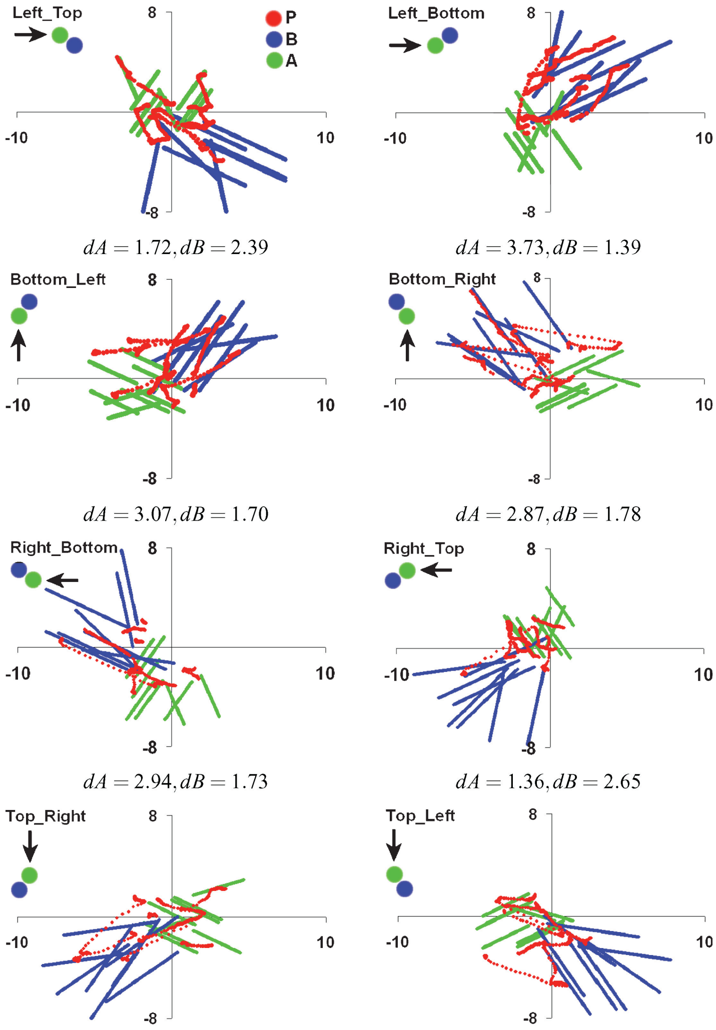

$dA = 1.72, dB = 2.39$ $dA = 3.73, dB = 1.39$

$dA = 3.07, dB = 1.70$ $dA = 2.87, dB = 1.78$

$dA = 2.94, dB = 1.73$ $dA = 1.36, dB = 2.65$

$dA = 2.08, dB = 2.42$ $dA = 2.09, dB = 2.41$

Fig. 6. Detailed fixation graphs for one participant.

in the last three cases (i.e., right_top, top_right and top_left). However, although we thought that it may be possible that the anomaly caught the attention in these cases, this was not consistently so. Clearly, these issues warrant further investigation in the future, with more experiments in which attention is manipulated in various ways. However, as we wished to show, these results do clearly indicate a pattern of fixation behaviour based on the nature of the collision event.

4. PAPER-BASED EXPERIMENTS

As a follow-up to the eye-tracking experiments, we now describe a series of paper-based experiments which we ran to determine if similar errors in judgment were evident when people attempted to *predict* the trajectories of objects, as opposed to simply *detecting* if an anomalous collision response had occurred.

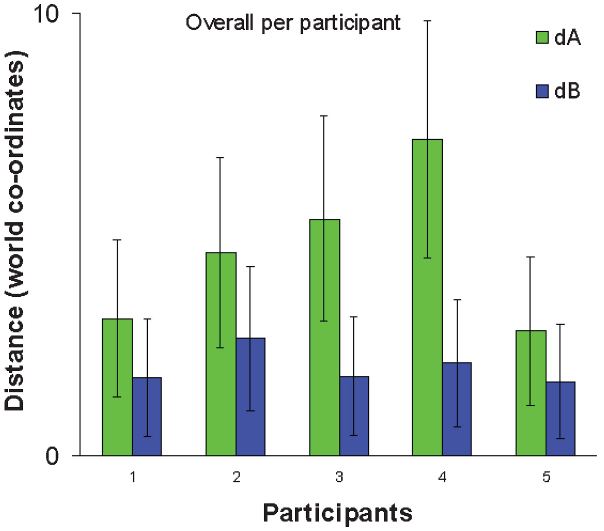

Fig. 7. Results for each participant, collapsed over direction, showing that, overall, they observed the struck ball B significantly more than the striking ball A. Error bars show standard deviations.

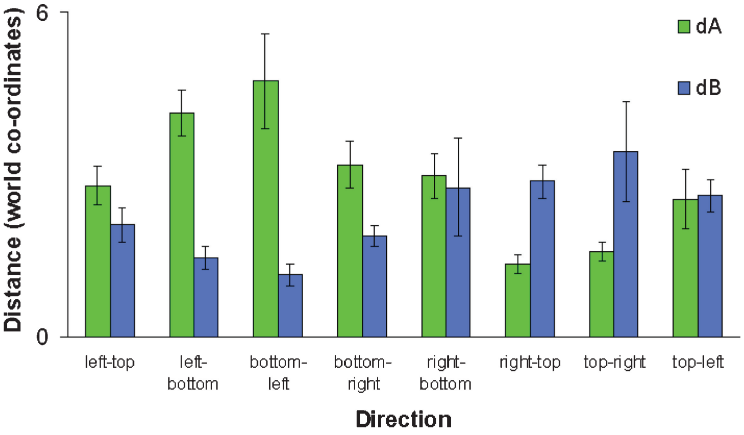

Fig. 8. Results for each direction, showing average distances to each ball over participants. Error bars show standard errors of the means. Results were not significant for right bottom and top left, marginally significant for left top and top right, and significant for the remainder.

Kaiser et al. [1992] showed that animation had a strong effect on people's ability to notice anomalous physical events. Earlier work on naive physics usually involved people sketching predicted paths of objects on paper, such as the trajectory that an object would take after being shot through a curved tube. A significant number of people erroneously predicted, on paper, that the object would continue to follow a curvilinear path, whereas when shown an animation of the event, they were better at recognizing anomalies. We were also interested in determining whether collision anomalies were also more easily detected in the animated scenario.

In this experiment, 29 participants were given a sheet of paper with a printed screen-shot of the situation at the time of collision. A line indicated the precollision trajectories of the two balls and participants were asked to draw the predicted postcollision trajectories on each piece of paper see

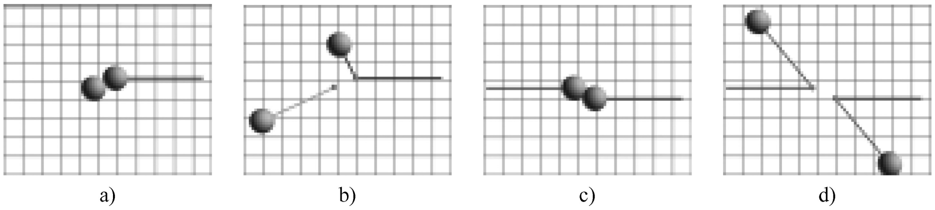

Fig. 9. Participants were asked to sketch the predicted postcollision trajectories of two spheres: (a) One ball moving precollision; (b) correct postcollision trajectories after (a); (c) two balls moving precollision; (d) correct postcollision trajectories after (c).

Figure 9. We showed 14 participants situations where only one ball was moving and showed 16 participants situations with both balls moving (one participant performed both tasks). In the former case, seven had participated in the previous animated study and seven had not—there was no significant effect of experience. In the latter experiment, no participants had performed the animated study. Participants were given the same instructions as in O'Sullivan et al. [2003], i.e., they were told that the balls were of equal mass and that there was no gravity or friction.

The results are shown in Figure 10. A summary of the results from O'Sullivan et al. is shown (a), which clearly demonstrates that people were more sensitive to errors in the struck ball and that they preferred expansion of the struck ball's trajectory, but had no preference for the striking ball's direction. For the one-moving case (b), people clearly exhibited the same biases as in the animated scenario in that they were better at predicting the trajectory of the struck ball and most people preferred expansion of the struck ball's trajectory (11 out of 14). Again, they had no significant preference for the striking ball's direction. By comparing the charts in (a) and (b), we can see that, in fact, error detection performance in the animated scenario was worse than in the paper-based study. This is different from the findings of Kaiser et al. [1992] and is a positive result with respect to masking anomalies in real-time collision handling. No bias in favour of either ball was evident when predicting trajectories in the two-moving case (c), but people generally underestimated the angle of both balls (13 out of 16—"right" and "left" are also used to label the balls, as both balls are moving initially). We can see from (d) that overall error was significantly higher for the striking ball in the 1-moving case. This bias was statistically significant: there were 14 participants, each with two values, struck and striking, hence: $F(1, 26) = 6.1$, $p < 0.02$. No difference in performance was found in the 2-moving case, although people were far less accurate at sketching both moving trajectories.

The results of these paper-based experiments would seem to cast doubt on the hypothesis that people were more sensitive to degradations in the striking ball's trajectory solely due to attention, i.e., because they were looking at it. If that were the sole factor, it should not have any effect when predicting the trajectories on paper and people should be equally able to sketch both trajectories accurately. Clearly, the difficulty of predicting the trajectory of the striking ball in this situation is due to the more complex and uncertain nature of the event itself. Even though participants were told that there was no friction or spin in the system, perhaps they were unable to avoid comparison with their real-world experiences. For example, if a spinning cue ball hits a static object ball in billiards or snooker, the spin has more of an effect on the trajectory of the cue ball. Nevertheless, as shown in our eyetracking experiments, people are still more interested in what happens a struck ball during a causal launching event, even if its behaviour is easier to predict (or perhaps because it is).

5. CONCLUSIONS AND FUTURE WORK

In this article, our main aim was to start an investigation into the feasibility of predicting attention in dynamical simulations, which typically involve colliding entities. As attention is clearly an important

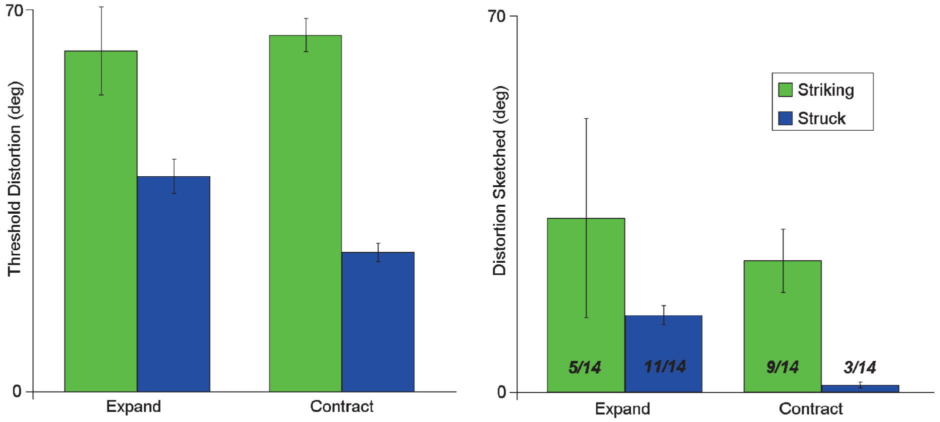

(a) Average Points of Subjective Equality (PSEs) over 15 participants, recorded in O'Sullivan et al.'s [2003] dynamic scenario

(b) One ball moving initially: average errors in paper-based sketched trajectories for 14 participants

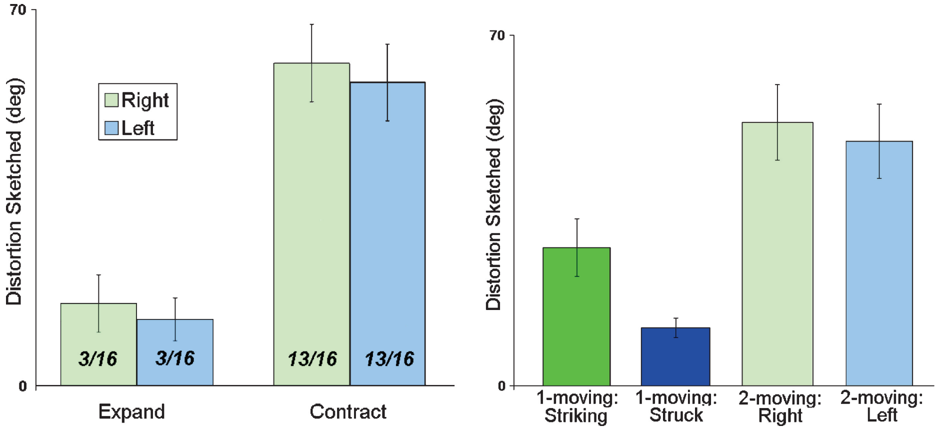

(c) Two balls moving initially: average errors in paper-based sketched trajectories for 16 participants

(d) Overall errors in paper-based sketched trajectories

Fig. 10. Paper-based experiments.

factor in the perception of causality [Choi and Scholl 2004], the ability to predict and exploit this fact could be very useful for real-time animation and physically-based simulation. The main motivation for this work comes from the success of similar approaches used in rendering and user interface design, where knowledge of the focus of attention allows resolution and computational power to be concentrated on areas or objects in a display that are most important. Similar savings could be made in the computationally expensive talk of collision detection and physical response determination. For example, low level of detail, proxy objects could be used for collisions between objects that are not currently being attended by the viewer, or simpler response models could be applied to the objects that are less likely to be noticed. Our second aim was to determine whether there is something about such dynamical events themselves i.e., their names, that attracts attention is some way. If this was the case, even in the simple examples depicted, then we could say that it may be possible in the future to derive a model that could predict attention automatically for physical animations.

Our investigation began with a series of eye-tracking experiments in which we attempted to establish whether a stationary ball that is struck by a moving one receives more attention post-collision than the striking one. We found, for the simple cases depicted, that this was indeed the case in general, although there were some noteworthy exceptions. However, this allocation of attention did not fully explain people's better error detection performance with the struck ball, which leads to the conclusion that the nature of the event and it's ability to attract attention must both be taken into account in any future model of dynamic event attention.

Many further questions remain, which will form the basis for future investigations. As Harrison et al. [2004] showed, task is very important in the allocation of attention. The task we gave our participants were very contrived—typically people do not actively watch out for anomalies. It is very likely that their ability to notice anomalies will be greatly lowered by engaging in a distracting task. To investigate these issues, we have developed a more natural scenario involving a snooker system that implements more accurate physics in a terrestrial environment. A particularly interesting question is whether the outcome of the task will have an effect on the perception of anomalies, for example if an erroneously calculated trajectory leads to a ball being potted that would have otherwise missed, (or vice versa)! Furthermore, we showed in earlier work that distracting items in a simulation can also degrade performance [O'Sullivan and Dingliana 2001], perhaps to such an extent that most physical anomalies would be completely imperceptible in simulations involving many colliding entities.

REFERENCES

BAUDISCH, P., DECARLO, D., DUCHOWSKI, A. T., AND GEISLER, W. S. 2003. Focusing on the essential: considering attention in display design. *Commun. ACM 46*, 3, 60–66.

CARLSON, D. AND HODGINS, J. 1997. Simulation levels of detail for real-time animation. In *Proceedings Graphics Interface*. 1–8.

CATER, K., CHALMERS, A., AND WARD, G. 2003. Detail to attention: exploiting visual tasks for selective rendering. In *Proceedings of the 14th Eurographics workshop on Rendering*. Eurographics Association. 270–280.

CHENNEY, S. AND FORSYTH, D. 1997. View-dependent culling of dynamic systems in virtual environments. In *Proceedings of the ACM Symposium on Interactive 3D Graphics*. 55–58.

CHOI, H. AND SCHOLL, B. J. 2004. Effects of grouping and attention on the perception of causality. *Perception and Psychophysics 66*, 6, 926–942.

DUCHOWSKI, A. T. 2003. *Eye Tracking Methodology: Theory and Practice*. Springer-Verlag, London.

GILDEN, D. AND PROFFITT, D. 1989. Understanding collision dynamics. *Journal of Experimental Psychology: Human Perception and Performance 15*, 2, 372–383.

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

HUNT, A. R. AND KINGSTONE, A. 2003. Covert and overt voluntary attention: linked or independent? *Cognitive Brain Research 18*, 102–105.

ITTI, L., KOCH, C., AND NIEBUR, E. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell. 20*, 11, 1254–1259.

KAISER, M., PROFFITT, D., WHELAN, S., AND HECHT, H. 1992. Influence of animation on dynamical judgments. *Journal of Experimental Psychology: Human Perception and Performance 18*, 669–690.

KAISER, M. K. AND PROFFITT, D. R. 1987. Observers' sensitivity to dynamic anomalies in collisions. *Perception and Psychophysics 42*, 3, 275–280.

LUEBKE, D., REDDY, M., COHEN, J., VARSHNEY, A., WATSON, B., AND HUEBNER, R. 2002. *Level of Detail for 3D Graphics*. Morgan-Kaufmann.

MACK, A. AND ROCK, I. 1998. *Inattentional Blindness*. MIT Press, Cambridge, MA.

MARMITT, G. AND DUCHOWSKI, A. T. 2003. Modeling visual attention in vr: Measuring the accuracy of predicted scanpaths. *Eurographics 2003, Short Presentations.* Saarbrücken, Germany, 217–226.

MICHOTTE, A. 1946. *La perception de la causalit. Louvain: Institut Suprieure de Philosophie. Michotte, A. (1963). The perception of causality T. Miles and E. Miles, Trans.).* Basic Books, New York.

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

O'Sullivan, C., Dingliana, J., Giang, T., and Kaiser, M. K. 2003. Evaluating the visual fidelity of physically based animation. *ACM Trans. Graph. 22*, 3, 527–536.

Scholl, B. J. and Nakayama, K. 2002. Causal capture: Contextual effects on the perception of collision events. *Psychological Science 13*, 493–498.

Scholl, B. J. and Tremoulet, P. 2000. Perceptual causality and animacy. *Trends in Cognitive Sciences 4*, 8, 299–309.

Simons, D. J. and Chabris, C. F. 1999. Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception 28*, 1059–1074.

Yee, H., Pattanaik, S., and Greenberg, D. P. 2001. Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. *ACM Trans. Graph. 20*, 1, 39–65.

Received December 2004; revised March 2005; accepted May 2005