A Model of Collision Perception for Real-Time Animation

Carol O'Sullivan[1], Ralph Radach[2], Steven Collins[1]

[1]Image Synthesis Group, Computer Science Department, Trinity College Dublin, Ireland
{Carol.OSullivan, Steven.Collins}@cs.tcd.ie
[2]Institute of Psychology, Technical University of Aachen, Jägerstr. 17, 52056 Aachen, Germany
ralph@ifp.psycho.rwth-aachen.de

Abstract. A model of human visual perception of collisions is presented, based on two-dimensional measures of eccentricity and separation. The model is validated by performing psychophysical experiments. We demonstrate the feasibility of using this model as the basis for perceptual scheduling of interruptible collision detection in a real-time animation of large numbers of visually homogeneous objects. The user's point of fixation may be either tracked or estimated. By using a priority queue scheduling algorithm, perceived collision inaccuracy was approximately halved. The ideas presented are applicable to other tasks where the processing of fine detail leads to a computational bottleneck.

1 Introduction

The aim of interactive animation systems is to create an exciting and real experience for viewers. The tendency in the past has been to attempt to achieve this by matching as closely as possible the physics of the real world. Of course, what a person perceives is strongly affected by the physical behavior of the world around them, but it is the human visual system that receives and interprets the visual cues from the surrounding environment, and it ultimately determines what we perceive. Therefore, we must look beyond the laws of physics to find the secret of reproducing visual reality. In interactive animation applications such as VR or games, it cannot be predicted in advance how a user or the entities in a virtual world will behave, so the animation must be created in real-time. As the number of independently moving objects in the scene increases, the computational load also increases. Possible scenarios are crowd simulations or rockfalls, where large numbers of visually homogeneous entities move around a virtual world in real-time. There are many bottlenecks in such systems, collision detection being a major one. A trade-off between detection accuracy and speed is necessary to achieve a high and constant frame-rate. It is possible to reduce perceived inaccuracy by taking perceptual factors into account, and also by estimating where on the screen a viewer is looking, either using an eye-tracking device, or by attaching more importance to certain objects or regions in a scene.

In Section 2 of this paper, we present the background to our work. In Section 3 a model of human visual perception of collisions is presented, based on two-dimensional measures of *eccentricity* and *separation*. We demonstrate how such a model can be used in a real-time, adaptive collision detection algorithm to reduce the perception of collision-handling inaccuracies when animating large numbers of similar objects. The model is validated by performing psychophysical experiments, described in Section 4. Finally, Section 5 presents conclusions and plans for future research.

2 Background

Traditional collision detection algorithms have required a large amount of geometrical intersection tests. To improve the efficiency of such algorithms, hierarchical representations of entities were developed to localise the areas where the actual collision occurred. These include sphere-trees [14][21][22], OBB-trees (Oriented bounding boxes) [11], and hierarchies of k-DOPs (Discrete Orientation Polytopes) [17]. While speed and efficiency has been the main focus of such research, the issue of a constant frame rate is also paramount. This problem has been addressed in part by exploiting coherence [5], and by using an interruptible collision detection algorithm [14]. The advantage of an interruptible algorithm is that the application has full control over the length of time that collision processing may take. However, inaccuracies in the handling of collisions may cause the viewer to perceive unrealistic behaviour of colliding entities. The following scenarios are possible when considering the problem of collision detection:

- Fully accurate collision detection: This will give frame-rate problems.
- Interruptible collision detection: Produces good frame-rates, but bad perception.
- Track or estimate the viewer's point of fixation; Use a perceptual model to schedule collision processing; Validate and refine the model with data from psychophysical experiments: This should reduce perceived inaccuracy within a target frame time.

In recent years the realisation has been growing within the computer graphics community of the advantages to be gained by using knowledge of human perception. Perceptual factors such as size and speed of objects have been used to choose the levels of detail (LOD) at which to render objects in a scene [10]. The advantages of simulating plausible motion, as opposed to physically accurate motion, have been investigated [3]. The results of psychophysical research have been used in realistic image synthesis [8][9][13]. It has long been established that many visual processing tasks deteriorate at increasing eccentricities from the fixation point [2] [28]. Therefore, an eye-tracking device that locates where on the screen a viewer is looking could be an important tool in such real-time systems [6]. In the past, the most common use of eye-trackers has been in medical and scientific research. These types of trackers are very accurate, but also very invasive, involving the use of head restraints, bite bars, or scleral coils which are inserted directly into the eye. More recently, more mobile, non-intrusive, and low-cost trackers have been developed, and their use has been gaining increasing support in the fields of Human Computer Interaction (HCI) and Virtual Reality (VR) [16][25]. The problem with such low-cost solutions is loss of accuracy, both temporal and spatial. We found that using such a tracker was infeasible without excessive restraints being imposed, so we simply attach the point of fixation to the centre of the screen, or attach it to a prominent object in the animation. We maintain that an analysis of human visual perception of collision events could enable a prioritisation of potential collisions to process within a given frame of an animation, hence reducing the negative impact of interruption.

3 The Application

To test the feasibility of the ideas presented herein, a three-dimensional animation system was developed, where non-convex, star-shaped entities move around and interact in real-time within a volume (see [20] for a detailed description).

3.1 Collision Detection

The “Sweep and Prune” algorithm from [5] is used for the broad-phase of our collision detection algorithm, i.e to detect overlaps of the fixed-size bounding boxes of interacting entities. When the bounding boxes of two entities overlap, a collision object is created. For the narrow phase of our algorithm, where more accurate collision detection occurs, an interruptible algorithm based on sphere trees is used, adapted from [21] and [14]. We approximate each entity with a sphere tree during a pre-processing phase, using an octree generation algorithm (see Figure 1). The collision detection algorithm is fully interruptible, thus enabling a fast, albeit approximate, response when necessary. At some point the application deems that collision processing should stop, e.g. when a pre-defined target time has been exceeded. This leaves a list of real collisions and a list of unresolved collisions. We have chosen to treat these collisions as real collisions, thus causing objects to occasionally bounce off each other at a distance.

The key to controlling the collision inaccuracy perceived by a viewer in a given frame of an animation lies in the scheduling method adopted. In [14] *round-robin scheduling* is implemented, where active collisions are resolved by descending one level in the hierarchy of every sphere tree at each iteration of the algorithm, until interruption. However, no account is taken of the perceptual importance of each collision. *Sequential scheduling* starts at the first collision and fully resolves each collision in turn, until completion or interruption, but again perception is ignored. In *perceptually sorted sequential scheduling* a perceptual importance is attached to each collision, and the active collision list is sorted based on this priority using a version of quick-sort adapted for linked lists. The list is then processed sequentially, meaning that collisions which are most important perceptually will be resolved first, leaving the more unimportant collisions to be resolved only if there is time left. In *priority queue scheduling*, we generate not one active collision list, but a set of priority queues, and round robin within them. A higher priority queue is resolved first, and only when all collisions on that queue have been resolved is the next highest queue processed. It is also possible to sort one or more of the priority lists, in *priority sorted queue scheduling*. Our application supports any of these scheduling strategies.

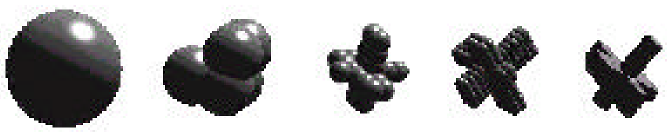

Figure 1: An entity and 4 levels of its sphere tree

3.2 Collision Perception

There are many potential factors that affect a human's ability to notice whether two objects have collided realistically or not. To consider using the factor in a prioritisation algorithm, the effects observed should be significant, and occur in most, if not all, of the subjects tested. In addition, they must be robust. We are looking for factors that can be generalised over a wide range of conditions, because we wish to apply them in a real-world scenario. In the applications being considered, three *types of collisions* may occur: "True" collisions, where entities touch, the collision is detected, and fully accurate collision response occurs (the control situation); Interpenetrations, where the entities also touch, but the collision is not detected or is ignored by the application, allowing them to merge into each other; Repulsions, where the application accepts two objects which are close but not touching as a true collision, causing a repulsion effect. There are certain points to be made in favour of allowing only repulsions to occur. The effect of one entity piercing through another is very noticeable and observations strongly suggest that this effect is more disturbing than the effect of repulsion, especially if the entities are of different colours. Another problem with interpenetration is that the anomaly lasts longer than repulsion, i.e. if two interpenetrating entities are ignored for several frames, they will continue to interpenetrate further and further, hence increasing the chance that they will be observed by a viewer. In addition, the visual perception of repulsion has well documented parallels in the study of spatial vision, hyperacuity, and brain physiology.

A summary of the physiological reasons for decreased spatial resolution in the periphery appears in [7]. Information projected onto the central part of the retina, i.e. the fovea, receives more processing. This means that detection of fine detail is facilitated in the fovea, and deteriorates with *eccentricity*. Therefore, this is a factor that is likely to affect the ability to detect a collision anomaly. *Separation* or gap-size will also most likely be a factor, as there is a one-to-one mapping from the cells in the retina to the cells in the primary visual cortex, called a retinotopic mapping, and it is quite precise, enabling spatial location information to be efficiently processed [26]. However, other factors may also affect collision perception, such as the *location* of the entities in the periphery [19]. The visual cortex contains many cells that are selective for *orientation*, i.e. they perform best when stimuli are oriented at a particular angle [15]. The *direction of offset* of two stimuli has been found to be a significant factor when performing a detection task in the periphery, and more important that the orientation of the stimulus itself [29]. *Image motion* has been observed to have a degrading effect on various types of visual task [4]. The faster a stimulus moves, the less fine detail the retina and hence the visual cortex can determine. The visual cortex also contains cells which are sensitive to *direction of motion* [30]. A "centrifugal directional bias" has been found in the brain of the macaque monkey, where there were more cells responsive to directions away from the fovea than any other direction [1]. Cells have also been found in another area of the brain (MST) which are responsive to certain types of spiral motion [12].

The presence of *distractors* could also affect a human's ability to accurately detect collision or non-collision, as could the nature of the distractors. These issues have been extensively researched in the area of Visual Search [23][27]. A pop-out effect occurs if the distractors are in a clearly distinguishable perceptual grouping from the target to be searched for (e.g. a different colour). If such an obvious grouping is not immediately apparent (i.e. all entities are visually similar), it is necessary to focus attention on each item in turn, i.e. to perform a serial search. In such a task, performance is significantly worse than in the pop-out tasks. In the situations that we are considering, e.g. simulations of large numbers of homogeneous interacting entities, such as crowd scenes or rockfalls, there will not be obvious perceptual groupings of objects. Therefore, the ability of viewers to detect a collision anomaly in such a scenario is of major interest to us. Other important factors which we do not consider here may also affect the perception of collisions, such as size, acceleration, colour, shape, and semantics.

3.3 A Model of Collision Perception

A model of collision perception has been developed, which may be used both to prioritise collisions for perceptually-sorted scheduling, and also to estimate perceived inaccuracy. In this way, we can test the feasibility of our approach, and also focus on the type of psychophysical data we wish to gather. When considering the inaccuracy present in a frame of an animation, we must distinguish between geometrical inaccuracy ∇, and perceived inaccuracy P. The geometrical inaccuracy in a scene is an estimate of the overall three-dimensional error that has been incurred by accepting non-collisions as real, causing entities to repulse without touching. If N is the number of collisions, and g_i is the maximum erroneous separation between the objects in collision i, we estimate this error by summing the potential gaps left during such "non-collisions":

$$\Delta = \sum_{i=1}^{N} g_i$$

In our applications, we cannot calculate the exact size of the gap between two entities which was left due to an incorrectly detected error, as this would take an excessive amount of time, and thus defeat our purpose. Instead, we can use the information available to us to estimate an upper bound on the maximum size of such a gap. We use the three-dimensional distance between the centres of the last two spheres found to be intersecting from the sphere-trees of each colliding pair. Hence, the further down the sphere-tree hierarchy each collision is allowed to progress, the more accurate the estimate will become. Alternatively, we could use the Hausdorff distance, as in [14].

Not all collision inaccuracies contribute equally to the inaccuracy perceived by the user in a single frame of an animation. Hence, the perceived inaccuracy P present in two frames of an animation with identical geometrical inaccuracy may be quite different depending on how the frame is viewed. For now, let us assume that eccentricity e, i.e. distance from the viewer's fixation point, and on-screen separation, estimated by maximum 2-dimensional gap size g' (g-prime, to distinguish it from the 3D separation), are the only two factors which affect perceived inaccuracy. If two spheres are interpenetrating, we find the midpoint on the line segment inside the intersection. We call this the *Centre of Collision*. We estimate the user's fixation point F for each frame, so eccentricity e is the 2-dimensional distance from F to the centre of collision. Collisions closer to the fixation point contribute more to the perceived inaccuracy of a frame than those further away and hence should receive higher weighting. Collisions further from F should receive lower weighting.

Similarly, the size of the maximum erroneous on-screen gap, g', may also be used to weight each collision, with larger gaps contributing more to inaccuracy than smaller ones. To calculate an upper bound on the 2-dimensional gap size, we take the centres of the last two spheres found to be intersecting, and calculate the 2-dimensional distance between their projections. If N is the number of collisions in a scene, P may be defined as follows:

$$P = \sum_{i=1}^{N} f(g_i^{'}, e_i)$$

For constant C, two possible estimates of the above function $f : Z \times Z \rightarrow R$ are:

$$f_1\left(g_i^{'}, e_i\right) = \begin{cases} \dfrac{g_i^{'}}{e_i} \text{ if } e_i \geq 1 \\ g_i^{'} \text{ if } e_i = 0_1 \end{cases} \quad \text{and} \quad f_2\left(g_i^{'}, e_i\right) = \frac{g_i^{'}}{\exp\left(\dfrac{e_i}{C}\right)}$$

These functions provide plausible models of how the visual system might work. They all assume a fall-off in the ability to detect collision anomalies with increasing eccentricity and decreasing gap-size. They differ only in the rate of fall-off, and the relationship between gap-size and eccentricity. Function f_1 decreases very rapidly, whereas function f_2 allows a more gradual fall-off (see Figure 5.b). It is not yet apparent which, if any, of the above models is the most appropriate. We have seen that other factors may also effect collision detection. However, it is now possible to test our application with respect to a plausible model of perception, allowing us to manipulate frame times and scheduling mechanisms, and test their effect on the viewer, as measured by our hypothesized perceptual metric.

3.4 Performance

To test performance, we ran the application with 10, 30, 100, 300, and 500 identical objects, each for 5000 frames. The density of objects inside each volume was equal. We first ran all experiments with no interruption, i.e. at full accuracy. In the case of 100, 300 and 500 objects, real-time performance was not achieved. In the 10 and 30 object tests, the frame-rate standard deviation was quite high. In order to estimate the time needed for the non-detection activities of our collision handling functions, we repeated all tests, this time interrupting at 0, i.e. only performing broad-phase tests and collision response. Based on these results, we chose a target time X for each number of objects, and interrupted collision detection at this time. We found that optimal results were achieved by interrupting the animations after 5, 20, 44, 30 and 10 milliseconds, for 10, 30, 100, 300 and 500 objects respectively. Figure 2 shows the improvement in mean time and standard deviation achieved by interrupting collision detection at this time, compared to fully accurate detection. This improvement is up to a factor of 16 in the 500-object case.

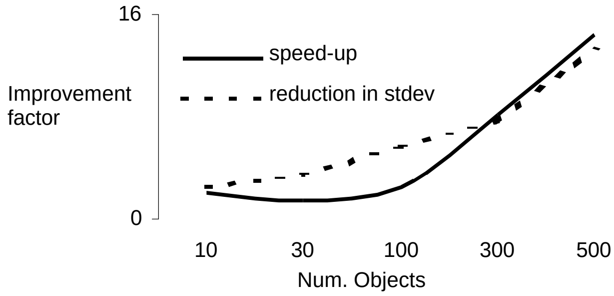

Figure 2: Improvements in average collision handling times and time deviation when using interruptible collision detection

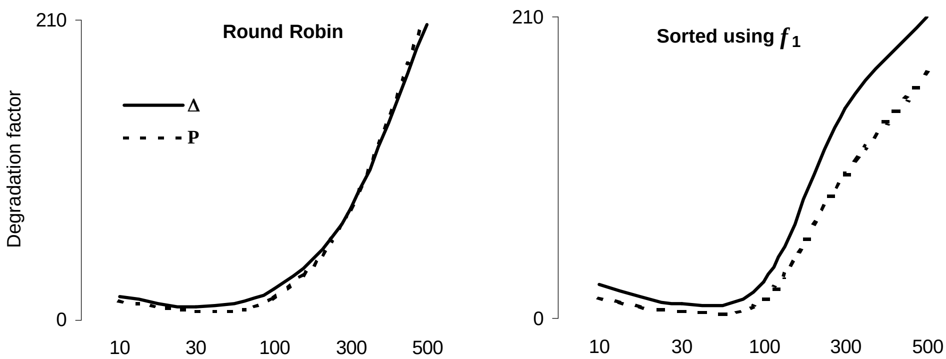

Figure 3: Disimprovement of overall and perceived inaccuracy with interruptible collision detection.

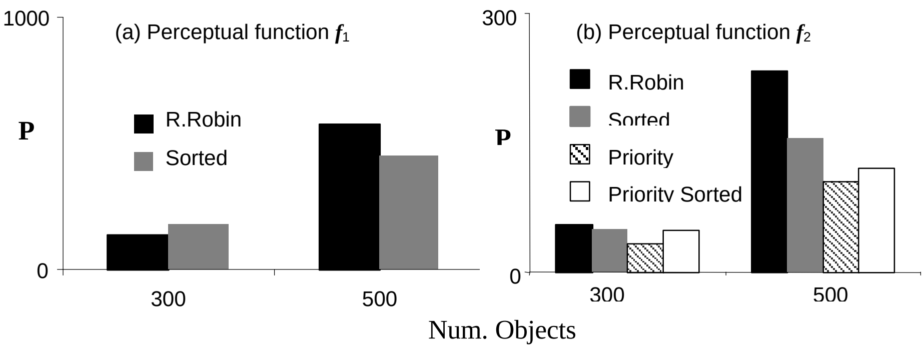

Figure 4: Comparison of perceived inaccuracy with various scheduling mechanisms and perceptual functions

Figure 3 shows the increase in both geometrical and perceived inaccuracy for the same tests. As the number of objects increases, both the perceived and total inaccuracy increases significantly, up to 200 times worse. We ran the same tests using perceptually sorted sequential scheduling, using function f_1 to measure inaccuracy, and to prioritise collisions. We found that the added overhead of performing a quick-sort at each frame reduced the time available for collision detection, causing an increase in overall inaccuracy Δ. Figure 4(a) shows the absolute values of P for both scheduling mechanisms. In the 500 object case the perceived inaccuracy is slightly improved, but is higher for 300 objects. It is also possible that function f_2 is more appropriate to model the human visual system. Therefore, we repeated the above tests for both 300 and 500 objects, but this time using f_2 to measure inaccuracy and sort collisions. We can see from figure 4(b) that based on this model, there is an improvement in perceived inaccuracy when sorted scheduling is used, and this is most noticeable in the 500-object case. We then implemented both priority queue scheduling, where the collisions on each active list are resolved in simple round-robin fashion, and priority sorted queue scheduling, where the first priority list is sorted and processed sequentially, and the second list is processed in round-robin order, thus reducing the sorting overhead incurred in the fully sequential case. The results are shown in Figure 4(b). It is clear from these results that the simple priority queue scheduling produced the best results for both 300 and 500 objects, approximately halving inaccuracy in both cases.

4 Psychophysical Experiments

To determine the functional field of view for the perception of collisions, we carried out 3 experiments with 12 participants. Filled white circles of 1 deg diameter presented on a black background served as stimuli. Subjects were asked to detect collisions (no gap between two circles) versus repulsions with minimal gap sizes of 0.1° or 0.4°. In experiment 1, we examined detection performance in a static display as a function of 3 levels of eccentricity, 4 directions of offset and 8 locations (up-left, right-down etc.). Experiments 2 and 3 used a dynamic situation with targets moving at a speed of 2.9° per second. In experiment 2 we added distractors that were different in colour from the colliding entities, and in experiment 3 we added distractors which were identical to the colliding entities (the "real-world" scenario). Motion was 2-dimensional, which does not impact upon the generality of our results, since it has been shown that humans use only two-dimensional visual information to make decisions about collision events, such as time to collision [24].

In all experiments, detection of the larger gap was significantly better than the smaller one, justifying the inclusion of the separation factor in our model. The effects of location and direction of orientation and motion were found to have a weak influence in experiment 1, and a very strong influence on performance in experiments 2 and 3, thus weakening the overall eccentricity effect. The number of distractors significantly affected performance when the distractors were visually similar to the colliding entities (experiment 4), but not when the distractors were different (experiment 3). Experiment 4 represents the task which most resembles the real-life situation which we are considering, and in this case a strong eccentricity effect was evident (Figure 5.a). The observed behaviour is most closely matched by a function of the form f_2, (Figure 5.b).

We found in Section 3 that the best reduction in perceptual inaccuracy was achieved by using priority queue scheduling, when inaccuracy was measured using f_2. This has now been shown to be a valid approximation to the behaviour of humans when viewing animations of large numbers of visually similar objects.

5 Conclusions and Future Work

We have presented a model of human visual perception of collisions, and have validated it psychophysically. We have shown the feasibility of using this model as the basis for perceptual scheduling of collisions in a real-time animation of large numbers of homogeneous objects. It has been demonstrated that by using a priority queue scheduling algorithm, perceived inaccuracy can be approximately halved when animating 300 or 500 objects. However, this model now needs to be refined to make it applicable in more general cases. The results of the psychophysical experiments have demonstrated that other factors, such as location and direction of motion, can have very strong effects under certain circumstances. Further psychophysical experiments into other factors which we have not yet addressed, such as velocity, acceleration, colour and luminance must also be conducted, if the model is to be truly representative of human behaviour.

Much work remains to be done to improve the application also. At the moment, the collision time-step is equal to the rendering time-step, which can lead to objects interpenetrating or tunnelling through each other if the time-step is too large. We are working on adapting the time-step for high-priority collisions, also using the perceptual model. However, it may be that some level of interpenetration must be accepted as a trade-off, so we must also study the perceptual response of the human visual system to this anomaly also. In addition, a more realistic response must be generated, using the laws of physics. The effects on this process of reduced information about points of contact is also being investigated. Finally, the use of a low-cost, mobile eye-tracker results in a certain amount of spatial and temporal inaccuracy. We need to measure this inaccuracy and incorporate some fault-tolerance into our system. Another approach could be to develop a model that predicts where the next likely fixation will be in each frame. Salience maps [18] could be used to achieve this.

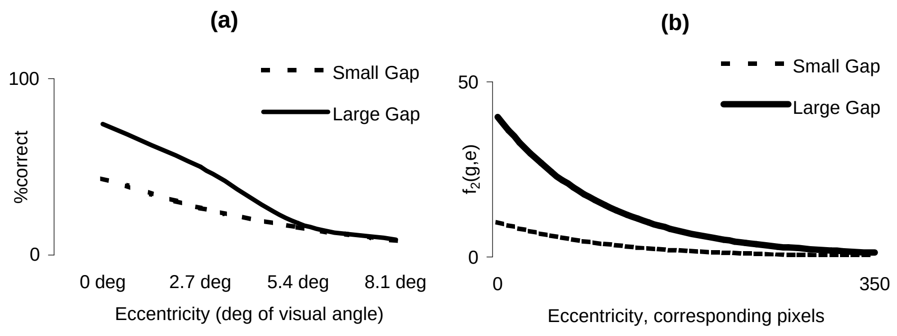

Figure 5: Performance by eccentricity and gap size: observed (a), modelled (b)

References

1 Albright, T.D. Centrifugal directional bias in the middle temporal visual area (MT) of the macaque. Vis. Neurosci. 2(17), 7-188. 1989.

2 Aubert, H. Foerster, 0. Beiträge zur Kentniss des indirekten Sehens. Arch. Ophth 3, 1-37 1857.

3 Barzel, R. Hughes, J.F. Wood, D.N. Plausible Motion Simulation for Computer Graphics Animation. Computer Animation and Simulation '96. 183-197. 1996

4 Chung, S.T.L. Levi, D.M. Bedell, H.E. Vernier in Motion: What Accounts for the Threshold Elevation? Vision Research, 36(16) 2395-2410. 1996

5 Cohen, J.D. Lin,M.C. Manocha, D. Ponamgi,M.K. I-COLLIDE: An Interactive and Exact Collision Detection System for Large-Scaled Environments. ACM Int. 3D Graphics Conference. 189-196. 1995.

6 Deering, M. High Resolution Virtual Reality. Computer Graphics, 26, 2, 195-202. 1992.

7 DeValois, R.L, and DeValois, K.K. Spatial Vision. New York: Oxford University. 1988.

8 Ferwerda, J.A. Pattanaik, S.N. Shirley, P. Greenberg, D.P. A Model of Visual Adaptation for Realistic Image Synthesis. SIGGRAPH '96. 249-258. 1996

9 Ferwerda, J.A. Pattanaik, S.N. Shirley, P. Greenberg, D.P. A Model of Visual Masking for Computer Graphics. SIGGRAPH '97. 143-152. 1997.

10 Funkhouser,T.A. Sequin, C.H. Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments. SIGGRAPH '93 247-254. 1993.

11 Gottschalk, S. Lin, M.C. Manocha, D. OBB-Tree: A Hierarchical Structure for Rapid Interference Detection. SIGGRAPH. 1996

12 Graziano, M.S.A. Anderson, R.A., Snowden, R.J. Tuning of MST Neurons to Spiral Motions. Journal of Neuroscience, 14(1): 54-67. 1994

13 Greenberg et al. A Framework for Realistic Image Synthesis. SIGGRAPH. 477-494. 1997.

14 Hubbard, P.M. Collision Detection for Interactive Graphics Applications. IEEE Transactions on Visualization and Computer Graphics. 1(3) 218-230. 1995.

15 Hubel, D.H. Wiesel, T.N. Receptive fields and functional architecture of monkey striate cortex. J.Physiology, 195, 215-243, 1968

16 Jacob, R.J.K. Eye Tracking in Advanced Interface Design. In Barfield,W. and Furness, T.A. (Eds) Virtual Environments and Advanced Interface Design. 258-288. 1995

17 Klosowski, J.T. Held, M. Mitchell, J.S.B. Sowizral, H. Zikan, K. Efficient Collision Detection Using Bounding Volume Hierarchies of k-DOPs. IEEE Tr. Vis. & Comp. Graph. 4(1). 1998.

18 Mahoney,J.V. Ullman,S. Image chunking defining spatial building blocks for scene analysis. Z.Pylyshyn (Ed.) Computational processes in human vision. 1988.

19 Nies, U. Bedenk, B., Heller, D. & Radach, R.. Eye movement patterns during search in a homogeneous background. Becker, Deubel, Mergner (Eds). Current Oculomotor Res. 1998.

20 O'Sullivan, C.A. Dingliana, J. Real-Time Collision Detection and Response using Sphere-Trees. Proc. Spring Conference on Computer Graphics, Slovakia. 83-92. 1999.

21 Palmer, I.J. Grimsdale, R.L. Collision Detection for Animation using Sphere-Trees. Computer Graphics Forum, 14(2) 105-116. 1995.

22 Quinlan, S. Efficient Distance Computation between Non-Convex Object. Proceedings International Conference on Robotics and Automation. 3324-3329. 1994.

23 Saarinen,J. Visual search for global and local stimulus features. Perception 23, 237-243. 1994.

24 Schiff, W. Detwiler, M. Information used in judging impending collision. Perception 8, 647-658. 1994.

25 Skerjanc, R. and Pastoor, S. New generation of 3-D desktop computer interfaces, SPIE: The Engineering Reality of Virtual Reality, 439 - 447, 1997.

26 Tootell, R.B.H, Silverman, M.S. Switkes, E, De Valois, R.L. Deoxyglucose analysis of retinotopic organization in primate striate cortex. Science, 218, 902-904. 1982.

27 Treisman, A. Perceptual Grouping and Attention in Visual Search for Features and for Objects. J. Experimental Psychology: Human Perception and Performance 8, 194-214. 1982.

28 Weymouth, R.W. Visual sensory units and the minimal angle of resolution. American Journal of Ophthamology, 46, 102-113. 1958.

29 Yap, Y.L. Levi, D.M. Klein, S.A. Peripheral hyperacuity: isoeccentric bisection is better than radial bisection. J. Opt. Soc. Am. A, 4(8), 1562-1567. 1987.

30 Zeki, S.M. Functional Organization of a Visual Area in the Posterior Bank of the Superior Temporal Sulcus of the Rhesus Monkey. J. Physiology. 236, 549-573. 1974.