A Perceptual Framework for Optimising Visual Detail in Virtual Environments

Martin Reddy [1]
Department of Computer Science,
University of Edinburgh,
JCMB 1408, King's Buildings
Mayfield Road,
Edinburgh, EH9 3JZ

Abstract. In most virtual reality (VR) graphics engines, there is a correlation between the degree of detail used to represent a scene and the time it takes to render an image of that scene. This can be loosely stated by the observation that a more visually complex scene will take longer to display than a less visually complex one. The rate at which a VR system can display images is of paramount importance; particularly when one considers that lags in an immersive VR system can cause effects of motion sickness and degrade user performance. It is therefore desirable to investigate the minimum amount of visual detail required in a scene and thus gain the optimal performance from the graphics system. This paper investigates limitations of the human visual system in order to formulate an integrated, perceptually-oriented framework for automatically degrading the level of detail (LOD) of objects in a scene (with the aim of increasing the overall performance of the system). This will be achieved by applying knowledge and theories from the domain of visual perception to the field of VR: thus basing the LOD framework on solid metrics of vision which accurately and efficiently optimise the visual information presented to the user.

Keywords. level of detail, spatial frequency, contrast sensitivity, visual perception.

1. Introduction

In an immersive virtual reality (VR) system, there are three major sources of lag: *sensor lag*, incurred by monitoring the real environment and communicating this data to the host computer; *processing lag*, incurred by implementing the dynamics and behavioural characteristics of the virtual environment (VE); and *rendering lag*, incurred by displaying a representation of the VE onto the display device [2].

The latter two components depend heavily on the content and complexity of the particular VE being simulated. For example, a VE which is composed of a large number of polygons will take longer to process and render than a VE with fewer polygons. This is not a simple relationship because, for example, the size and geometry of polygons can affect performance; also, the use of texture maps can

---

[1] e-mail: `M.Reddy@ed.ac.uk`, URL: `http://www.dcs.ed.ac.uk/~mxr/`

add visual detail to a VE which would otherwise require a large number of small polygons. However in general, we can state that the more visual complexity which we include in a VE, the greater the lag which will be subsequently induced.

The magnitude of these lags can have very real consequences for the user of a VR system. Visual delays in immersive applications have been reported to cause effects of motion sickness [18]; the symptoms of which include nausea, pallor and cold sweating [14, 20]. Systems which suffer from a noticeable lag can also affect the performance of the user; particularly for coordination and navigational tasks. For example, a lag of around 500ms can seriously degrade hand–eye coordination tasks such as drawing and writing [9]. In addition, it has been reported that reduced frame rates (less than 15Hz) can diminish a user's sense of *presence* within a VE [3]—the feeling that a user has of being in an environment other than where their body is [19].

It is therefore clearly undesirable to produce VEs which are unnecessarily complex. As a result, a number of techniques are commonly employed to reduce the complexity of a VE at any given instant. These include the process of world subdivision, clipping non-visible objects, using various levels of detail for certain objects [1] and supporting texture mapped polygons.

The technique which this paper will focus upon is level of detail (LOD). This involves holding a number of representations of an object, each varying in complexity (e.g. polygon count) and alternating between these during the simulation based upon certain selection criteria—thus allowing the complexity of the VE to be modulated in real-time. The purpose of this paper is to describe possible LOD selection criteria and subsequently develop an integrated framework for implementing these in an immersive VR system. As Bricken states, "Psychology is the Physics of VR" [4]. However little empirical work has been done to investigate the foundations for developing a VR graphics system on models of visual perception. This framework will therefore be founded upon evidence and theories from the domain of visual psychology in order to base the selection criteria on limitations of the human visual system, as opposed to arbitrary judgements and *ad hoc* heuristics.

This topic represents the subject of an on-going research programme within the department of Computer Science at the University of Edinburgh, in association with the department of Psychology. It has been inspired by [7] who appealed for a perceptually-oriented formalisation of LOD techniques.

2. Limits of Vision

2.1 Visual Acuity

The resolution of the human visual system must clearly possess a finite threshold which is ultimately determined by the spacing and pooling of photoreceptors in the retina (the cells which detect incoming light). However, when a user dons a Head Mounted Display (HMD), the resolution of that device will normally be far lower than that of their biological vision system[2] and so the resolution of the HMD will define the user's visual acuity while they are immersed in the VR system. As a demonstration of this point, the eye can detect detail down to a size of about 0.5 min of arc [10]; whereas the angular resolution of a modern LCD based HMD is in the order of 10 min of arc (E.g., based upon manufacturers' figures, the Virtual i-O *i-glasses!* offers an angular resolution of 6.8 min of arc, the Forte Technologies *VFX1*: 10.4 min of arc, and the VictorMaxx *CyberMaxx*: 12.6 min of arc).

2.2 Peripheral Vision

The eye's sensitivity to detail is not uniform across the entire visual field. Instead we find that our visual acuity is highest towards the centre of the retina, at a point called the *fovea*. There are a number of physiological reasons for this: for example, the concentration of cone photoreceptors in the retina decreases rapidly with eccentricity (distance from the fovea). Also, the ratio of cortical cells in the brain which are devoted to the foveal region is notably disproportionate: with around 80% of all cortical cells dedicated to the central 10 degrees of the visual field [6]. The result of these characteristics is that our vision is maximally sensitive within a central region of approximately 4 deg of arc, and drops off smoothly towards the periphery [21]. This reduction in visual acuity across the retina is significant: with around a 35-fold difference existing between the fovea and the periphery [16].

2.3 Motion Sensitivity

The human vision system cannot resolve as much detail in an object which is moving across the retina as it can in an object which is stabilised on the fovea. This causes the familiar effect of objects blurring as they move past our point of fixation, or as we pan our head to fixate on another target. The reason for this effect is thought to be due to the eye's inability to track rapidly moving targets accurately; thus causing a slippage in the retinal image [15].

3. Formulating the Framework

Based upon the preceding examination, we can see that there are three principal ways in which visual complexity can be optimised in a VE with respect to limitations of the display device and the human visual system. We can therefore describe these in terms of criteria for the LOD selection mechanism:

1. **Size/Distance** : When a complex object is positioned at a distance from the viewpoint, many of its detailed features will be projected onto an area less than the size of one pixel. As such, these features will make little or no

---

[2] The often quoted example is that when using a modern HMD, the user's visual acuity is reduced to around 20/200 (i.e. effectively, the user can see from 20 feet what a "normal" person can see from 200 feet). This, by definition, renders them legally blind.

contribution to the visual representation of the object at that distance. It is therefore possible to select a lower LOD model when the object exceeds a particular threshold distance from the viewpoint. Doing so will have little impact on the display, but will afford a certain performance advantage. This technique can be simplistically represented in pseudo-code as:

```
getViewpointPosition( viewpoint ) -> position1
getObjectPosition( object ) -> position2
distanceBetween( position1, position2 ) -> objdist
if ( objdist < Threshold1 )
   switchInHighLOD( object )
else if ( objdist > Threshold2)
   switchInLowLOD( object )
endif
```

Note that the values for the constants `Threshold1` and `Threshold2` need not be equivalent. In fact, `Threshold2` is often made slightly larger than `Threshold1` in order to introduce a degree of hysteresis and avoid the distracting flicker of an object continually switching LOD at the threshold distance [1].

An alternative method for implementing distance LOD is employed by the Open Inventor toolkit [23], and the AVRIL graphics library. This approach bases the LOD selection upon the rectangular area occupied by the bounding volume of an object after it is projected into screen coordinates. This is perhaps a more accurate technique because it is the size of a feature in screen space which determines whether it will be visible on the display device; also, it avoids the issue of choosing a single arbitrary point within the object's 3D volume to use for the distance calculation. However, this technique is best implemented by the underlying graphics renderer and cannot easily be implemented at the application level.

2. **Eccentricity** : As described previously, our visual system is limited to detailed resolution in only a very small region of the visual field (around 4 deg of arc). The field of view (FOV) of contemporary HMDs range from around 30 degrees to 120 degrees of horizontal arc (although most low-end HMD units offer a FOV of less than 60 deg of arc). There is therefore substantial opportunity for a highly detailed object to be present within the user's peripheral field. When this is the case, the object's visual complexity may be wasted because the eye cannot resolve features in the periphery to the same degree as in the fovea. A number of graphics systems have been developed to take advantage of this phenomenon by degrading the detail of features in the user's peripheral vision [11, 13]. A method for incorporating this facility into a VR system is proffered by [17].

3. **Motion** : We have seen that an object moving rapidly across the retina cannot be perceived in as much detail as a stationary object. It would therefore be sensible to attempt to reduce an object's LOD in relation to its motion. Considering that an immersive VE is an inherently motion-rich experience

(e.g. users' head movements generate motion flows as do autonomously moving objects within the VE), then we would expect that a motion-sensitive LOD mechanism will favourably affect the performance and interactivity of the system. Funkhouser and Séquin incorporated support for this feature into their architectural walkthrough system [7].

4. Formalising the Framework

Of the three criteria presented above, distance modulation is the most prolific LOD mechanism in use today (although surprisingly most VR toolkits do not explicitly support this, and often leave the task of distance LOD selection to the application program). However, one problem which is commonly associated with LOD techniques is the often detectable flicker when different models are swapped in. This is because the selection thresholds are normally arbitrary in nature and not based on any models of visual perception. If the user is not to be distracted by the changes in visual complexity, then the LOD selection thresholds must be based on a solid metric which accurately models what the user can and cannot see.

The remainder of this paper is therefore concerned with introducing such a metric from the domain of visual perception; investigating how this can be applied to an immersive VR system; and analysing the pertinent issues which are involved in the implemention of this process.

4.1 A Metric for Visual Detail

Campbell and Robson devised a means of measuring a subject's visual acuity by using a pattern known as a *contrast grating* [5]. This is simply a pattern where contrast is varied sinusoidally across the display, producing a series of alternating light and dark vertical bars (see Figure 1). The spacing between bars is measured by a quantity called *spatial frequency*, defined in units of contrast cycles per degree of visual field (c/deg). For example, a high spatial frequency implies a small distance between adjacent bars and hence represents a stimulus of high detail.

The visibility of a grating is dependent upon its spatial frequency and its contrast (luminance difference between adjacent bars). A curve known as a Contrast Sensitivity Function (CSF) can be plotted to record a subject's ability to resolve a grating based upon these two factors. For example, the CSF in Figure 2 states that for a contrast grating moving at 3 deg/s, the subject will be unable to resolve any spatial frequencies greater than about 8 c/deg.

The effect of increased velocity and/or eccentricity is to shift the CSF closer towards the y-axis: thus reducing the threshold of visual acuity in both cases. I.e. a subject will be able to resolve fewer high spatial frequencies (regions of high detail) as the stimulus moves faster across the retina [12] or is presented further towards the periphery [22].

Spatial frequency has been used to analyse and describe the limitations of the human

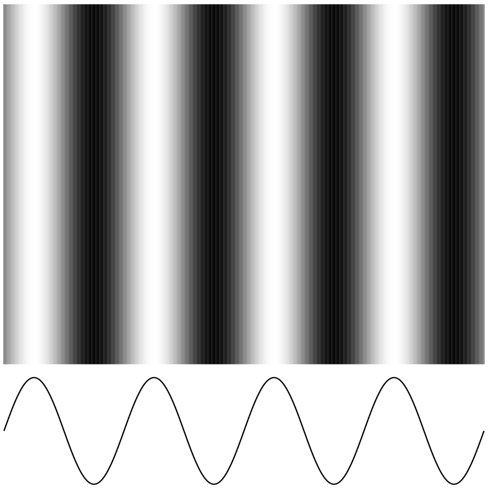

Figure 1: *A sample contrast grating. The curve below the grating illustrates the sinusoidal variance of contrast across the image. If this grating was positioned to occupy 1 deg of visual arc, then it would have a spatial frequency of 4 c/deg.*

visual system for over 25 years. If this metric can be applied to the domain of real-time computer graphics, then we would have a wealth of literature at our disposal to determine the perceptual content of a computer display.

4.2 Applying Spatial Frequency to Computer Graphics

Visual psychologists often adopt a reductionist approach during vision experiments in order to isolate and analyse the variable of interest. As a result, much of the corresponding literature is concerned with simple contrast gratings which vary only over one dimension, and do so in a harmonic fashion. However, a rendered image of a VE is obviously 2D and involves complex changes of intensity. The first question we must therefore ask is whether knowledge of the visibility for a 1D harmonic contrast grating can be applied to a 1D complex contrast grating. Then, whether it can be applied to a 2D complex pattern, such as a computer-generated image.

In the first instance, Campbell and Robson performed experiments with compound waveform gratings in order to investigate how the visibility of these is related to that of their component harmonic gratings. They found that the appearance of a compound grating is characterised by the independent contributions from each of the harmonic components. Their results showed that if a compound grating is displayed such that some of its high frequency components are below threshold, then these features will not be visible in the compound grating and can be removed without any perceivable change being made to the grating. These results form the basis of the *multi-channel* model of visual perception which exists to the current day. Essentially, this theory proposes that differently sized stimuli are detected and

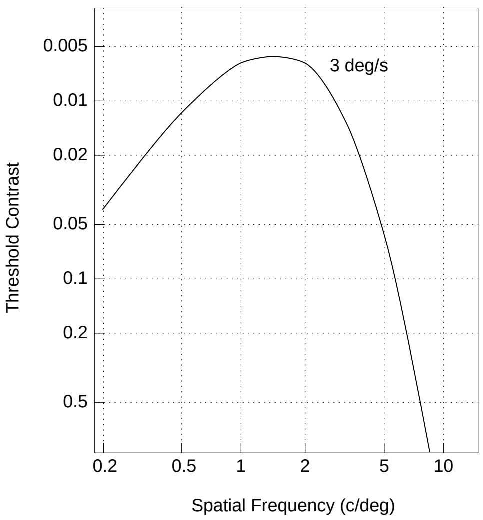

Figure 2: *A Contrast Sensitivity Curve illustrating detection thresholds for contrast gratings moving at 3 deg/s. All gratings represented by the region below the curve are considered to be detectable. Adapted from [12].*

processed independently by the visual system.

Secondly, we must address the problem of how to describe a 2D image in terms of spatial frequency. This is done by incorporating an orientation variable, i.e. the grating in Figure 1 is oriented horizontally, however it could just as easily be oriented vertically or at any other arbitrary angle. Therefore to describe a 2D computer-generated image in terms of spatial frequency we must supply a 2D array containing values for all relevant spatial frequencies at a number of sample orientations.

It should perhaps also be noted that the CSF curves refer to gratings of harmonic contrast variations; whereas a computer display is composed of discrete pixels and so presents a square-wave distribution. However, it has been shown that a square-wave grating has the same visibility characteristics as a sine-wave grating—with a constant scaling factor applied— [5]. So we can assume that the CSF is applicable to computer displays such as monitors and HMDs.

5. Implementing the Framework

5.1 When and Where to Calculate Spatial Frequency

Spatial frequency is a measure of the detail which is presented to the visual system—i.e. in an immersive VR system it is a measure of the detail presented on the HMD. In order to accurately gauge the perceptual content of a VE, we must therefore apply our analysis to the rendered image which is transmitted to the HMD. We cannot accurately predict what will be displayed by simply looking at the geometry of the scene because the geometry can be displayed differently depending upon the shading model being applied, the level of lighting which is in effect, the use of texturing etc. We must therefore base our analysis of the scene upon the rendered image, not the geometry of the objects within the scene.

There are however a number of complications. If we take the isolated case of one object, then the spatial frequency content of that object will change as it rotates or moves away from the viewpoint, i.e. spatial frequency is viewpoint dependent. Therefore, to accurately calculate the spatial frequency content of that object we must apply our analysis to the display every time the object or the viewer moves. This is obviously unacceptable in a time-critical application like VR. However, more to the point, this is impractical because we are required to know the spatial frequency profile *before* the image is displayed—once the image has been rendered we will have already expended the computational resources which we wish to conserve. Therefore we must endeavour to pre-calculate the spatial frequencies in an object off-line. This must be done from several viewpoints around the object in order to capture all of the object's features. We can then interpolate these values during the simulation in order to predict the spatial frequency content of any arbitrarily positioned object in real-time; and subsequently select the most suitable LOD to utilise.

5.2 Calculating Spatial Frequency of a Rendered Image

Given a snapshot of an object from a particular viewpoint, we want to be able to find all of the relevant spatial frequencies which are component in that image. Specifically, we wish to locate each 'feature' in an object and calculate its fundamental (lowest) frequency (where a feature is loosely defined as a region of similar contrast). A feature's fundamental frequency will be inversely related to its cross-sectional length at each orientation (see Figure 3).

Any feature may contain a range of spatial frequencies but we are only interested in the fundamental frequency of each because we can conclude that if this is not visible, then none of the higher frequencies will be visible either. At a higher level, every object will be composed of a number of features, but we are only concerned with the smallest of these in each object (because if the smallest feature is visible, then so will all of the larger features).

From the rendered snapshot of an object we can calculate the values for each feature's fundamental spatial frequency in units of cycles per pixel (c/pixel)—essentially a measure of how many pixels a feature extends over at each orientation. Then once we know the FOV for the HMD, we know how much of the visual field is subtended by a single pixel. We can therefore apply a suitable scaling factor to the c/pixel values to gain results in units of c/deg.

Therefore, for each object, we can find the highest spatial frequencies (and their orientations) at any instant. The orientation of the spatial frequencies is of little importance when considering a stationary object (except perhaps to compensate for the aspect ratio of the HMD). However, because we wish to include motion

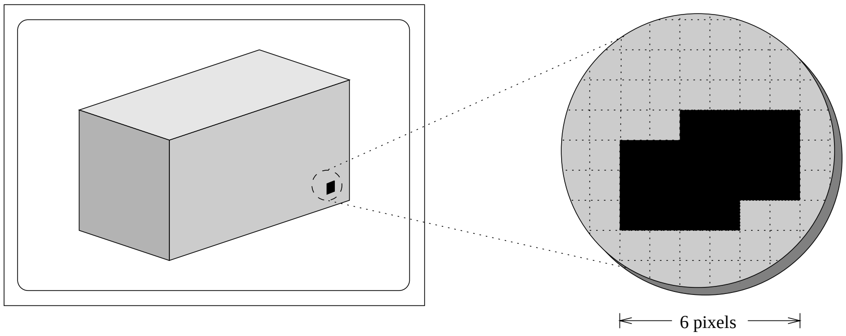

Figure 3: *A snapshot of a simple object showing a single small, black feature. The fundamental spatial frequency of the feature is inversely related to its cross-sectional length at each orientation, e.g. the horizontal fundamental frequency* $\propto \frac{1}{6}$ *c/pixels.*

optimisations in our framework, we must record the orientation of spatial frequencies because the visibility of a moving spatial frequency is dependent upon its alignment with the direction of motion.

For example, consider a long, thin, vertical feature. Its horizontal spatial frequency (derived from its length) will be very high, but its vertical spatial frequency (derived from its height) will be very low. If this feature moves horizontally, then at a certain velocity it will become invisible to the naked eye. If this feature then moves at the same velocity but in a vertical direction, it will remain visible because the vertical spatial frequency is far higher than its horizontal counterpart. This effect can be observed in the natural world: e.g. when you look fixedly out of the window of a fast moving train towards a meshed fence of horizontal and vertical wires, then you cannot see the vertical wires of the fence, only the horizontal wires which are aligned with your direction of motion.

5.3 Developing an Integrated Implementation

Once we have a method for estimating the perceptual content of a display in terms of spatial frequency, we can then begin to look at integrating this into the LOD selection criteria. Obviously for this to be effective, there must be some correlation between spatial frequency and LOD. Specifically, we might expect that because a lower LOD model is more coarse and has less detailed features in it, it should contain fewer high spatial frequencies. This correlation has been borne out by informal observations made by the author.

Based upon the preceding discourse, a perceptually-based LOD system could be implemented as follows:

Off-line : Before the simulation of the VE is initiated, each LOD for every degradable object is analysed to discover its spatial frequency profile. This involves taking a number of snapshots of an object from different viewpoints. For each snapshot, all of the features within that image are located and their size calculated in units of c/pixel for a number of orientations. These values can then be scaled to units of c/deg based upon the FOV of the HMD being used.

On-line : During the simulation, the LOD Scheduler will analyse the distance, eccentricity and velocity of each degradable object (in units of: m, deg and deg/s, respectively). By applying the results from the various CSF curves it becomes possible to calculate the highest spatial frequency which the average person will be able to resolve in this situation. Then with the results from the off-line process, we can estimate the spatial frequencies which will be contained in each LOD if it were displayed. This will enable us to select the lowest available LOD which contains frequencies entirely above the visibility threshold. I.e. if two different LOD models are expected to be perceived equally, then the less complex model will be chosen.

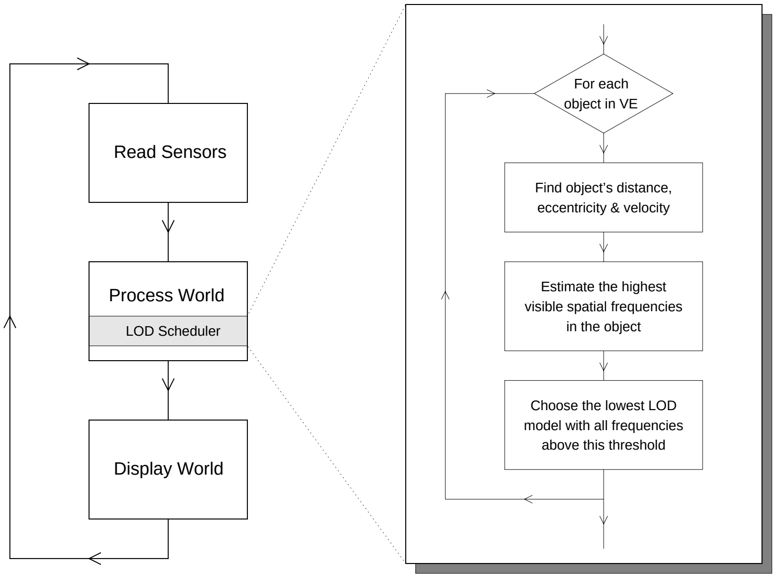

Figure 4: *Illustrating how the proposed LOD Scheduler (right-hand flow diagram) fits into the typical main loop of an immersive VR system (left-hand flow diagram).*

Figure 4 presents the standard *sense–process–display* loop of a typical VR application and illustrates how the LOD Scheduler fits into this scheme. The position of the Scheduler within the 'process' stage is important in order to gain the optimal performance. It must follow any processing which updates the location and velocity of each object (including the viewpoint); but it should precede any subsequent processing (e.g. collision detection, dynamics equations etc.) so that this will be applied to the actual model about to be displayed.

6. Conclusions

In this paper we have looked at a paradigm for reducing the visual complexity of objects in a VE based upon various characteristics of the human visual system. This has been done with the expressed aim of improving the performance and interactivity of the VR system, whilst doing so based upon proven metrics and principled selection criteria.

One point which should be highlighted however, is that any system which attempts to make judgements on the perceptual content of a display must, to be completely accurate, track the user's gaze. This is because the user's perception is based upon the image formed on their retina, which at any instant could be focussed on any region of the display device. It would be acceptable, though less accurate, to assume that the user will always will be looking towards the centre of the display. This was the assumption made by [7] and also by [11].

In this respect, it is the author's opinion that the above approximation can be more strongly advocated in an immersive VR system than in a desktop VR system. This is because whenever the user evokes a large change in their point of fixation, there will normally be an associated head movement. As a result their resting gaze will generally be quite closely related to their head orientation. Thus the differential between a user's point of fixation and the centre of the display will be less pronounced in an immersive, head-tracked system. (N.B. It should be noted that this is merely an intuitive speculation and remains unsubstantiated by empirical data at this time).

Therefore, if an appropriate eye-tracking technology is not employed then one should be aware that the user may be able to perceive slight inconsistencies in the rendered image. This could of course be compensated for by slackening the selection thresholds somewhat, or introducing a degree of hysteresis. However, it must be borne in mind that most researchers in the VR field agree that interactive update rates are more important than display fidelity, and so some visual incongruities can be tolerated. This is aptly expressed by Wloka when he states that "Presentation quality is expendable: to be wrong and on time is more valuable than to be right and late" [24]. These same sentiments are encapsulated by Krueger when he makes the statement: "Geometry is not reality. Interactivity is reality" [8].

In conclusion, the product of this paper has been the formulation of a measure to analyse the perceived quality of a computer-generated image; and the development of a framework to utilise this metric to optimise the performance of immersive VR systems.

References

[1] P. Astheimer and M-L. Pöche. Level-Of-Detail Generation and its Application in Virtual Reality. In *Proceedings of the VRST'94 conference*, pages 299–309, Singapore, August 23–26 1994.

[2] P. J. Atkin. Parallel Processing for Virtual Reality. In *Parallel Processing for Graphics and Scientific Visualization*. University of Edinburgh, 25 May 1993.

[3] W. Barfield and C. Hendrix. The Effect of Update Rate on the Sense of Presence within Virtual Environments. *Virtual Reality: Research, Development, Applications*, 1:3–16, 1995.

[4] W. Bricken. Virtual reality: Directions of growth. Notes from the SIGGRAPH '90 Panel. Available on-line from `ftp://ftp.u.washington.edu/public/virtual-worlds/papers/`, September 1990.

[5] F.W. Campbell and J.G. Robson. Application of Fourier Analysis to the Visibility of Gratings. *Journal of Physiology*, 197:551–566, 1968.

[6] N. Drasdo. The Neural Representation of Visual Space. *Nature*, 266:554–556, 1977.

[7] T. A. Funkhouser and C. H. Séquin. Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments. *Computer Graphics Proceedings, Annual Conference Series*, pages 247–254, 1993.

[8] S. Garassini. The Ultimate High of Myron W. Krueger, Father of Artificial Reality. *Tech Images/Paris-Cit*, 18:41–42, 1991.

[9] R.L. Gregory. *Eye and Brain: the Psychology of Seeing.* Weidenfeld and Nicolson, London, fourth edition, 1990.

[10] G. W. Humphreys and V. Bruce. *Visual Cognition: Computational, Experimental & Neuropsychological Perspectives.* Hove Lawrence Erlbaum Associates, 1991. ISBN 0-863-77125-4.

[11] F. Hurault. A Head Slaved Visual System for Flight Simulators. In *Proceedings of the International Training Equipment Conference and Exhibition*, pages 37–42, London, UK, 4–6 May 1993.

[12] D.H. Kelly. Motion and Vision. II. Stabilized Spatio-Temporal Threshold Surface. *Journal of the Optical Society of America*, 69(10):1340–1349, October 1979.

[13] M. Levoy and R. Whitaker. Gaze–Directed Volume Rendering. *ACM SIGGRAPH Special Issue on 1990 Symposium on Interactive 3D Graphics*, 24(2):217–223, 1990.

[14] K.E. Money. Motion Sickness. *Physiological Reviews*, 50:1–39, January 1970.

[15] B.J. Murphy. Pattern Thresholds for Moving and Stationary Gratings During Smooth Eye Movement. *Vision Research*, 18:521–530, 1978.

[16] K. Nakayama. Properties of early motion processing: Implications for the sensing of egomotion. In R.Warren and A.H.Wertheim, editors, *The Perception and Control of Self Motion*, pages 69–80. Lawrence Erlbaum, Hillsdale, NJ, 1990.

[17] M. Reddy. Musings on Volumetric Level of Detail for Virtual Environments. *Virtual Reality: Research, Development, Applications*, 1:49–56, 1995.

[18] C. Regan. An Investigation into Nausea and Other Side-effects of Head-coupled Immersive Virtual Reality. *Virtual Reality: Research, Development, Applications*, 1:17–32, 1995.

[19] M. Slater and M. Usoh. Presence in Immersive Virtual Environments. In *Proceedings of VRAIS'93 (IEEE Virtual Reality International Symposium)*, pages 90–96, Seattle, Washington, September 18–22 1993.

[20] K.C. Uliano, R.S. Kennedy, and E.Y. Lambert. Asynchronous Visual Delays and the Development of Simulator Sickness. In *Proceedings of the Human Factors Society (30th Annual Meeting)*, pages 422–426, Dayton, OH, 1986. Human Factors Society.

[21] W. R. Uttal. *The Psychobiology of Sensory Coding.* Harper & Row, NY, 1973.

[22] V. Virsu and J. Rovamo. Visual Resolution, Contrast Sensitivity, and the Cortical Magnification Factor. *Experimental Brain Research*, 37:475–494, 1979.

[23] J. Wernecke. *The Inventor Mentor: Programming Object-Oriented 3D Graphics with Open Inventor(TM), Release 2.* Addison-Wesley, Reading, MA, 1993.

[24] M. M. Wloka. Thesis Proposal: Time-Critical Graphics. Technical Report CS-93-50, Dept. of Computer Science, Brown University, Providence, RI, November 1993.