Assessment of Visual Requirements for Flight Simulation

Keith K. Niall
Defence and Civil Institute of Environmental Medicine, Toronto, Canada

Byron Pierce
Air Force Research Laboratory, Mesa, Arizona

The assessment of visual requirements for flight simulation is a study both sophisticated and fraught with complexity. It is commonplace for authors to declare the immensity or the richness of their subject, but vision is an old topic where much has been forged already, and where much remains to be accomplished (see Warner & Casey, 1995). Some issues in vision are deceptive in their apparent simplicity: a favorite anecdote is that Marvin Minsky once assigned the "problem of vision" as an exercise to a summer graduate school in engineering at M.I.T, fully expecting that his students would solve the problem by term's end. Needless to say, they did not. A lot of effort can be saved by learning to *think* about vision clearly, by learning to gauge the scope of particular problems in vision research as no larger and no smaller than they are. "For in psychology there are experimental methods and *conceptual confusion.* The existence of the experimental method makes us think we have the means of solving the problems which trouble us; though problem and method pass one another by" (Wittgenstein, 1958, p. 232e).

There is a dilemma characteristic to the assessment of visual requirements for flight simulation. One must decide whether to provide either complete conditions or else necessary conditions for visually guided performance on a task. That is, one can decide to reproduce exactly the changing pattern of light that reaches an aviator's eyes during flight, or else one can decide to reproduce only those aspects or properties of the visual scene that can be shown to be crucial to performance. On first consideration this may seem to weigh realism against expediency. There is a great deal of cost and care involved in promoting optical rigor in displays for flight simulation. Indeed, one of our goals is to reproduce the changing pattern of light presented to an aviator - to ensure that the illuminated scene presented by the simulator is indistinguishable from what might be experienced in flight. Yet it is worthwhile to remember this is an ideal. Usually we express the point by saying that if more money were available for engineering development, or if our graphics computers were only much faster, then realism would improve.

This emphasis on complete realism forestalls another legitimate concern: that we don't need to reproduce everything an aviator sees in the cockpit. That is because another of our goals is to determine which visual properties or aspects of the simulation are necessary to the task at hand, for those tasks that are made difficult by their dependence on vision (e.g., Heikens, Gray, O'Neill, & Salisbury, 1999). From this knowledge, we may say in detail how it is the aviator accomplishes the task, that is, in a refueling task by attending to the rivets of the companion plane, or in a search task by the systematic control of eye movements in raster pattern. Our second goal, then, is to appreciate the role of learning and adaptation in aircrew training insofar as it is guided by vision. Incomplete simulations can be used to great training advantage, because aviators adapt with marked elasticity to many novel or impoverished conditions of illumination. Moreover, observers learn to interpret new kinds of displays with alacrity: night vision goggles (Rabin, 1993), infrared cameras, and other devices present visible information that is useful in identification tasks, but where the characteristic reflectances of familiar objects may be different than under ordinary daylight conditions of observation. One goal of flight simulation is to provide complete optical fidelity, but some other goals are to understand which aspects of a simulation are essential for a task, and to extend the sensory facility of the aviator by presenting information from new domains, or by presenting information in new ways.

METRICS

The search for appropriate metrics is a recurring theme in the assessment of visual requirements. A *metric* is loosely characterized as a scheme or a strategy for measurement: mathematics is the source of these schemes or strategies. Then the search for a metric consists in the application of appropriate mathematics to the domain of our concern. The establishment - that is, the choice and the publicly enacted convention - of a metric enables us to answer how much and how well a flight simulation succeeds. How much, in that an appropriate metric taken from optics can be applied to gauge the completeness of optical fidelity that is achieved in the simulation (e.g., Beall & Loomis, 1997). How well, in that enumeration and measurement can be made of the optical properties that are known to play a part in performance of a task. Indices of performance can be compared to the number and measure of these properties, as well: Three temptations should be avoided in the search for a metric; these temptations are easy to describe and difficult to avoid. The first is the temptation to pursue mathematics for mathematics' sake: an application ought to be enlightening at least to those with a general education in mathematics, if not always enlightening in practise to many who lack such an education. The second is the temptation to use a wrench for a hammer. The astute choice of a metric is an act of ingenuity or creativity: no recipe book will help to guide any truly novel application. The third, most dangerous temptation is quietism or cynicism in the search for a relevant metric. To put the matter vividly, trying to specify visual requirements without an appropriate metric is like trying to dig a hole in the earth with one's bare hands: an ill-considered waste of time. Mathematics is the language of science, and understanding is not given to those who refuse even to communicate.

PROPERTIES OF LIGHT, LAND, AND SKY

A primary demand for the visual display of a flight simulator is that it provide adequate resolution. What drives this requirement of resolution? There are specific tasks that demand a high-resolution display for their simulation: an example is identification of the aspect angle of an aircraft in the middle distance (in terms of kilometers). One could say that high-resolution displays are required because our eyes are capable of resolving very small subtenses of visual angle, but that begs the question. The important question is what there may be in the environment that requires an observer to resolve small differences.

There are two aspects of the air environment that determine the requirement for resolution, but that are so obvious as to be ignored: land and sky. That is, one determining factor is the difficulty of accurately depicting terrain complexity (see Kleiss, 1996), and another factor is the modeling of scattering and absorption characteristics of the atmosphere. Depiction of the terrain may be categorized as: (a) the depiction of landmarks such as cultural features and gross geographical forms, and (b) the depiction of texture, including local geology, vegetation, and coloring of the soil or surface. But such a distinction is hardly adequate: first, it makes texture a heading on the same footing as *miscellaneous* - texture may then include anything at all. Second, the distinction is blurred: what was texture a moment ago may become a vitally important feature when size scale changes as one continues a flight path. There are several ways in which observers use terrain: in way-finding, in visual search, and in the intuitive judgment of speed, orientation, and heading by the relative motion of elements in the scene due to perspective (i.e., by what has come to be known as optic flow: Cutting, Vishton, Fluckiger, Baumberger, & Gerndt, 1997). Even these several uses of detail in the landscape will mandate that depiction of the terrain be as finely-grained as is possible with current display technology - and we do not yet know how to count all the uses an aviator makes of changes in terrain.

How then do characteristics of the atmosphere determine resolution requirements? The scattering of light by aerosol particles, and absorption contribute to an overall attenuation of light in the atmosphere, that is, they contribute to the reduced transmittance of the air. In a pure, dry atmosphere (this ideal is known as a Rayleigh atmosphere) that consists of a mixture of gases without dust, a black object subtending a small visual angle against a bright sky could have a visible range of more than 350 kilometers along the observer's line of sight. (Middleton, 1952, pp. 18 and 104). However, this capability does not contribute to the requirement of resolution in a visual display as much as does a photometric relation: the cosine law of illumination (Smith & Atchison, 1997, p. 277). As an observer's line of sight draws parallel to a flat surface, as occurs when flaring to land on a runway, the visual range of objects and textures that are seen on the flat surface changes dramatically. The measure of the solid angle they subtend changes with the cosine of the angle the line of sight makes to the flat surface. Features on the runway that appear well spaced at one moment are crowded together by the effect of scene compression in the next moment (the magnitude of this compression is given by the cosine law: Domini & Caudek, 1999; Niall, 1997). In flight simulation, the effect of compression of detail poses a requirement for increased resolution, and this effect becomes crucial during approach and during takeoff. Then the transmittance of light in air and the cosine law of illumination can be considered characteristics of the atmosphere that determine resolution requirements in flight simulation.

How do we quantify the resolution that we require in a flight simulator? There is a hack solution to this problem: that the device be able to resolve a pair of lines that are separated by two arc-minutes of visual angle, for an observer at a comfortable distance from the display screen. That number is meant to represent the separation in visual angle between a pair of lines that can just be resolved by the human eye: it is supposed to represent an "eye-limiting resolution" of film quality. Yet the specific forms that are presented to an observer influences the minimum angular subtense that the observer can resolve: in a task called vernier acuity (Fahle & Harris, 1998), line offsets can be distinguished by eye at angular subtenses of a fraction of an arc-minute - less than the postulated minimum resolution. The measurement of resolution may seem to pose a problem, because in flight we are often concerned with distance and with orientation, and we are, accustomed to measuring in those terms. Measures of distance and orientation do not suffice to gauge resolution, because those geometric properties are not preserved in the propagation of light by which an image is formed. Measure of visual angle and the cross-section of solid angle are somewhat better: in a display image, these measures change in a regular manner as an object of fixed size is depicted to recede in distance from the observer. Another family of measures may render greater service in the future. It has been suggested that measures that do not change with the distance and orientation of objects in the environment can be used to gauge how closely observers resolve a variety of forms. This family of measures is outlined in Mundy and Zisserman (1992; also Niall, 1999).

PHYSIOLOGICAL CONSTRAINTS

Not only the medium of propagation of light, but also the constitution of the visual system itself influences the perception of properties and surfaces we may wish to display in a flight simulator. Often the many contributions of the visual system are construed as a succession of effects, a causal chain that ends with sight; yet such a picture is misleading. Those contributions are better conceived as posing constraints or bounding conditions on sight: that is, as being impediments to sight. Some of these bounding conditions are familiar in terms of visual requirements for acuity: near-sightedness and far-sightedness can be understood as contributions of refraction in the resting state of the lens, whereas presbyopia can be understood as the loss of ability to change the accommodative state of the lens. It may also be remembered that the greater contribution to refraction in the eye is made by the cornea, not the lens (such components do interact: Valluri et al., 1999). Astigmatism is the condition in which the shape of the cornea departs appreciably from spherical. Other bounding conditions on sight are more subtle. Aircrew selection standards in many countries mention both visual acuity (usually in terms of discrimination on the standard Snellen or Landolt charts) and ability to discriminate colors in terms of hue. The utility of both these standards for aviation is in dispute. Changes in visual acuity with age are widely familiar (Fahle & Daum, 1997; Kikukawa, Yagura, & Akamatsu, 1999) and a source of chagrin to pilots in their regularly scheduled medical examinations. If the influence of the visual system on the perception of form and color is misconstrued as a chain of causes, then these constraints that lens and cornea impose on vision will be understood as only the first links in an indefinitely long series of causes, each of which may be considered for its contribution to visual requirements.

Yet the story ends sooner than that. Such "neural" factors distinct from the said optical constraints are seldom if ever elaborated independently. All these factors may and should be filed under the heading of *acuity* instead. The term *acuity* has been appropriated as a rubric for a variety of concerns, where finally acuity indicates success or degree of failure in discriminating form or discriminating color. Standards of performance in discrimination are amenable to better treatment, as when they are elevated into the useful construct known as the *standard observer.* Several versions of a normal observer have been specified for the discrimination of color (Trezona, 1998), although similar conventions have not yet been established for the visual discrimination of form (hence our continued reliance on eye charts).

DISPLAY PROPERTIES

The optical, mechanical, and electronic characteristics of the components of flight simulators may also constrain the performance of vision-based tasks. Naturally some issues concerning specific devices quickly become yesterday's news, as one technology or corporation is superseded by another. Some issues do recur: the angular placement of a seam or an occluder between adjacent panels of a large display can interfere with the performance of a specific task (Bentz, 1980: e.g., in air-to-air refueling, or in a formation manoeuver). Eternal compromise seems to mark other issues: as much as practical, the field of view presented by a flight simulator should be unrestricted. An unrestricted field of view has never been either convenient or practical, if unrestricted is taken to mean the presentation of a panoramic view, a complete wrap-around scene that includes a significant vertical extent and that enables the observer to turn around and view the scene behind. A large field of view has been a difficult requirement to satisfy, partly because it conflicts with another requirement that all parts of a visual display should be presented at maximum resolution. But, from the time of this writing, it was not terribly long ago that the rendering of figures in perspective and the development of a hidden-surface algorithm were big news in graphics.

The display systems that are used in contemporary flight simulators have two main components: the image generator, and the associated image source, the latter including software, database, and display. An image generator can be any process or system by which a computer produces an image. The cost of these devices ranges from several thousands of dollars for a personal computer to well over a million dollars for the latest and greatest of high-end devices. The choice of an image generator involves two concerns: the update rate (the speed with which an image can be computed) and its database (the complexity of the scene to be computed). Several studies have been made of the requirements for image generator update rate and scene complexity in flight simulation. A basic result is that an image generator used in flight simulation must compute and display imagery at maximal complexity and' at high update rate. In the simulation of flight at high speeds and low altitudes, a minimum update rate of 60 frames per second is indicated; this generally rules out the use of low-end image generator devices. Most image generator devices that meet this requirement combine a general computing processor with a specialized, high-speed graphics processor. In the simulation of other flight tasks, a lower update rate can sometimes be used.

There are a number of factors that mandate image generator capacity. Some considerations are:

1. the number of polygons that the graphics system will compute each second
2. the number of point-light sources to be displayed in a second
3. the number of picture elements (pixels) the system generates
4. the interlacing of the display (interlaced systems compute every other horizontal line during an update; noninterlaced systems compute every line)
5. database management considerations
6. the scheme by which the level of textural detail is altered as an observer approaches a surface
7. contingencies for processor overload when more polygons are specified to be written than can be written in real-time
8. the hidden-line or occultation algorithm that is used
9. smoothing or anti-aliasing techniques
10. the application of computer-generated or photorealistic texture to single polygons
11. the requirement for the database to be correlated with other databases in networked simulations or sensor simulations.

The image computed by the image generator should be presented to the observer with sufficient brightness, contrast, resolution, field of view, and with minimal geometric distortion. This is the job of the image source. Differences in image quality also depend on their characteristics in terms of timing accuracy and positional accuracy. Primary technologies include cathode-ray tubes (CRTs), back-lit liquid crystal displays (LCDs), light-valve liquid crystal displays, and laser-based technologies. For simulator displays with a large field of view, these technologies are paired with projection optics, as opposed to being viewed directly.

CRTs have a single electron gun that scans an image sequentially on a screen, in either an interlaced or noninterlaced manner. Current CRTs can address an image of up to 2500 x 2000 pixels. However, greater addressability does not always mean improved image quality. Back-lit liquid crystal displays project light through a liquid crystal panel. Like CRTs, typical liquid crystal displays execute a sequential scan, although this is no constraint: it is the characteristic of many image generators that are used with liquid crystal displays. The persistence of liquid crystal display elements in time is far greater than that of most CRT phosphors: this greater persistence induces a blurring or smearing of moving images. Advances are being made in liquid crystal display switching speed, and in the number of elements that can be addressed at once in a liquid crystal display. A negative characteristic of most liquid crystal display panels is the visible grid that is formed by the seams or interstices of adjacent pixels. The effect has been described as similar to looking through a mesh of wires, like a screen door. Liquid crystal *light valves* use a CRT coupled to a liquid crystal layer. A high-intensity light source is sent to one side of the liquid crystal layer. There the light is modulated, then reflected by a mirror through a projection lens and onto a screen. Since a CRT does the image scanning for a light-valve, the image is scanned in raster pattern.

Liquid crystal light valves have a high light output and exhibit no "screen-door" effect. Laser-based technologies incorporate new solid-state eye-safe laser light sources: no display phosphors or liquid crystal elements are involved. In laser-based systems, laser light is shaped into many beamlets, then modulated, scanned, and projected onto a viewing screen. Such devices may be able to provide very high resolution, increased contrast, greater brightness, and improved color gamut in dynamic imagery.

There is such a variety of large-format flight simulator displays that they are difficult to classify. A start might be to distinguish screen displays from refractive or reflective displays. (*Refractive* displays bend light a little, or transmit it directly; *reflective* displays reverse the direction of light transmission.) Refractive or reflective displays can be further distinguished as pupil-forming or nonpupil-forming. With pupil-forming displays, there is a definite volume of air within which an observer's eyes must be positioned. This volume lies about an axis of projection. Nonpupil-forming displays allow the observer to move a reasonable distance off-axis and still observe the projected image in focus. Decisions about displays are made on the basis of eye relief, field-of-view, the range of an observer's head movements, light transmission, optical distortion, contrast rendition, and resolution.

One solution to the continuing compromise about resolution is the development of area-of-interest simulations. There, a small portion of the display is presented at fine resolution, while the graphics for the remainder of the screen are rendered more closely. This may be thought to mimic the changes in acuity that occur from foveal vision to peripheral vision (Anstis, 1998), though there are other better descriptions of the difference between central and peripheral vision (Nasanen & O'Leary, 1998). Area-of-interest displays may be presented in different ways, as well. Certain aspects of the display may have high intrinsic value and so be chosen to be rendered in greater detail (one may value incoming aircraft over more cirrus clouds) or the area of enhanced resolution may be slaved to the direction in which an observer is looking. This may be as simple as turning off the parts of the display behind the observer's head, or it may represent a sophisticated and seamless blending of a central area with the larger background, where the central area glides over the display following the direction of gaze. In that case, the area to be rendered in detail must be calculated from eye- and head-tracking devices, in real-time.

Other contemporary issues emerge from the massive computation used to present visual imagery in flight simulators. (Of course, "massive" is a relative term, if you are used to modeling the evolution of galactic clusters, or suchlike.) The image quality of a flight simulator display can be affected by the manner and sequence in which pixels are written (that is: refreshed) on the display screen. The standard method of pixel-writing follows a raster pattern: pixels are written left to right in the top row of the screen, then left to right in the row next to top, and so on until the entire screen has been refreshed. It is possible to engineer multiple-line displays as well: one of these follows a pattern in which all rows of the display are written at the same time, beginning with the leftmost column of pixels and continuing to the rightmost column.

In all these methods, each pixel of the display screen is refreshed completely before any pixel is overwritten: each refresh cycle of an entire screen is called a frame. The frame rate of a display is known to affect its image quality. A frame rate of 60 Hz (60 Hertz is 60 frames per second) is generally accepted and used in North America for both simulation and broadcast. This frame rate has one advantage: it is well above the average critical flicker fusion frequency of the human eye, so that successive frames appear as a continuously changing picture rather than as a flashing series of still pictures. This effect has distinct consequences for the perception of raster-scanned displays. (Thomson & Saunders, 1997). Further increase in frame rate - from 60 Hz at least to 120 Hz - has been shown to enhance image quality.

It may not be clear how else the pattern in which pixels are written can affect the perceived image quality of a display until one considers temporal changes on small parts of the display screen (Gorea & Hammett, 1998; Lindholm, 1992). Events that occur simultaneously should appear simultaneously on the display screen, but this common-sense relation is difficult to maintain when the pixels of the display are refreshed at different times, albeit very quickly. Suppose that a disk is pictured to move quickly across a screen, and not all the pixels of a frame are being refreshed at once. Because the disk moves, it takes on a different position in successive frames. Most often only part of any frame is present on a screen, since the pixels on the screen are being refreshed continuously between one frame and the next. The difference in position of the disk between frames may then appear as a "tearing" or jaggedness at the edges of the disk. Similarly, point sources of light may appear to be doubled. In this interaction between spatial position and motion in time, the pattern in which pixels are written can affect image quality greatly. Effects of asynchrony also appear in what are known as transport delays in simulation. A simulation is not simply a movie, to be viewed passively; observers in the simulator are also active participants. Yet there may be an unnatural lag between the actuation of a device in the simulator cockpit and the depicted effect of that control on the screen image. Such lags can be on the order of hundreds of milliseconds, and these lags have been cited as a primary cause of simulator-induced nausea and headaches (called *simulator sickness* on analogy with motion sickness, but see van der Steen, 1998).

How can all these effects on image quality be assessed? There does exist a variety of methods with which to begin. Given a standard set of patterns, we can measure how well those patterns have been displayed by the simulator (one can also think of variations of the display over time in this context). Or better, we can ask how well that standard set of patterns can be perceived by an observer in the simulator. There are a number of standard sets of patterns (called *basis sets)* with which to begin. A basis set should have the property that an arbitrary pattern on the screen can be reproduced as a combination of the patterns from the set. Several basis sets are in common use: fuzzy bars of different frequencies, whose luminance changes in sinusoidal fashion from left to right (one-dimensional Fourier gratings); black-and-white checkerboards of different sizes (two-dimensional Hadamard or Haar patterns); and fuzzy checkerboards of different sizes, whose luminance changes in sinusoidal fashion from left to right, and up to down (two-dimensional Fourier patterns). Any of these sets may be used (see Papoulis, 1968), though a two-dimensional set is to be preferred where two-dimensional images are in question.

When such a basis set of patterns is displayed, how well are the patterns transmitted at a given luminance, or how well can those patterns be detected? Each pattern in the set is transmitted or detected more or less well at one luminance: very high frequency patterns may not be detected at all in some cases. The function described by this quantitative assessment of transmission or detection is known as a *modulation transfer function* (Jennings & Chapman, 1997; Lopez-Gil, Iglesias, & Artal, 1998). It provides a descriptive and graphical illustration of changes in image quality. Image statistics can be calculated on the modulation transfer function or directly on changes in the patterns themselves. There is a particular statistic or figure of merit that is suited to the comparison of-two-dimensional images: the crosscorrelation coefficient (Collet & Quinquis, 1994; Thomas, Hanson, & Oliensis, 1994). It provides a measure of the similarity of two-dimensional images. Together, these methods are a first step toward a larger field of inquiry (Young & Calvert, 1974)-the application of image-processing techniques to problems of image quality in flight simulation.

TENTATIVE CONCLUSIONS

Slowly we are learning
We at least know this much,
That we have to unlearn
Much that we were taught,
And are growing chary
Of emphatic dogmas;
[Sight] like Matter is much
Odder than we thought.
Auden (1976, p. 207)

One function of a flight simulator is to present the same pattern of light to an observer's eyes as would be present to those of an aviator during flight. Yet an important study of training effectiveness has a narrower focus. There we seek to know which parts of or variations in the visual display affect performance. Sometimes "the experimenter should not hope, and does not need, to display all the information in an ambient array. He is not trying to simulate reality .... What is required is only that the essential invariant be isolated and set forth" (Gibson, 1979, p. 305; cf. Startchik, Milanese, & Pun, 1998). The first goal is one of exact reproduction, whereas the second goal is one of analysis. We are far from realizing either goal. It would be futile to make predictions for the direction of future research. One of the characteristics of research, distinct from the applications of technology, is that in research we are widening the circle of our knowledge, and we simply do not know what lies beyond the circumference. According to one philosopher, the best way of identifying someone who has real genius is that we never discover exactly where it was they lost their way. However, there are some things we can say about current directions of research: that the notion of resolution or acuity has been used as a grab-bag notion that it would be worthwhile to unpack, and explicate in detail; that a close and astute application of metrics is a welcome contribution; and that there has not been sufficient consideration of the "environmental" properties of light, including the structure of landscape terrain and the influence of the atmosphere on aerial scenes. These problems are recommended to the reader for a solution.

REFERENCES

Anstis, S. (1998). Picturing peripheral acuity. *Perception,* 27, 817-825.

Auden, W H. (1976). *Collected poems.* E. Mendelson (Ed.), New York: Random House.

Beall, A. C., & Loomis, J. M. (1997). Optic flow and visual analysis of the base-to-final turn. *International Journal of Aviation Psychology,* 7(3), 201-223.

Bentz, J. L. (1980). Joining techniques for optically combined visual display systems. *Proceedings of the Second Interservice/Industry Training Equipment Conference and Exhibition* (*pp.* 46-59). Arlington, VA: National Security Industrial Association.

Collet, C., & Quinquis, A. (1994). Motion estimation methods and noisy phenomena. *Machine Vision and Applications,* 7, 247-258.

Cutting, J. E., Vishton, P M., Flűckiger, M., Baumberger, B., & Gerndt, J. D. (1997). Heading and path information from retinal flow in naturalistic environments. *Perception & Psychophysics,* 59(3), 426-441.

Domini, F., & Caudek, C. (1999). Perceiving surface slant from deformation of optic flow. *Journal of Experimental Psychology: Human Perception and Performance,* 25(2), 426-444.

Fahle, M., & Daum, I. (1997). Visual learning and memory as functions of age. *Neuropsychologia,* 35(12), 1583-1589.

Fahle, M., & Harris, J. P (1998). The use of different orientation cues in vernier acuity. *Perception & Psychophysics,* 60(3), 405-426.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin Company.

Gorea, A., & Hammett, S. T (1998). Spatio-temporal vernier acuity. *Spatial Vision,* 11(3), 295-313.

Heikens, M. F., Gray, G. W, O'Neill, H. J., & Salisbury, D. A. (1999). *1997 Canadian Forces Air Operations Vision Survey. Section 1: Operational Vision Requirements.* Toronto, Canada: Technical Report 1999-019 of the Defence and Civil Institute of Environmental Medicine, Department of National Defence.

Jennings, J. A. M., & Charman, W. N. (1997). Analytic approximation of the off-axis modulation transfer function of the eye. *Vision Research,* 37(6), 697-704.

Kikukawa, A., Yagura, S., & Akamatsu, T (1999). A 25-year prospective study of visual acuity in the Japan Air Self Defense Force personnel. *Aviation, Space, and Environmental Medicine,* 70(5), 447-450.

Kleiss, J. A. (1996). Assessing the influence of operational factors on the perceived structure of real-world scenes viewed during low-altitude, high-speed flight. *International Journal of Aviation Psychology,* 6(4), 335-358.

Lindholm, J. M. (1992). Perceptual effects of spatiotemporal sampling. In M. A. Karim (Ed.), *Electro-optical displays (pp.* 787-807). New York: Marcel Dekker.
Lopez-Gil, N., Iglesias, I., & Artal, P (1998). Retinal image quality in the human eye as a function of accomodation. *Vision Research,* 38, 2897-2907.
Middleton, W E. K. (1952). *Vision through the atmosphere.* Toronto, Ontario: University of Toronto Press.
Mundy, J. L., & Zisserman, A. (1992). *Geometric invariance in computer vision.* Cambridge, MA: MIT Press.
Nasanen, R., & O'Leary, C. (1998). Recognition of band-pass filtered hand-written numerals in foveal and peripheral vision. *Vision Research,* 38, 3691-3701.
Niall, K. K. (1997). "Mental rotation," pictured rotation, and tandem rotation in depth. *Acta Psychologica,* 95, 31-83.
Niall, K. K. (1999). Estimates of shape by eye: or, The little invariant that could. *Acta Psychologica, 100,* 291-320.
Papoulis, A. (1968). *Systems and transforms with applications in optics.* New York: McGraw-Hill.
Rabin, J. (1993). Vernier acuity through night vision goggles. *Optometry and Vision Science,* 70(8), 689-691.
Smith, G., & Atchison, D. A. (1997). *The eye and visual optical instruments.* Cambridge, England: Cambridge University Press.
Startchik, S., Milanese, R., & Pun, T (1998). Projective and illumination invariant representation of disjoint shapes. *Image and Vision Computing, 16,* 713-723.
Thomas, J. L, Hanson, A., & Oliensis, J. (1994). Refining 3D reconstructions: A theoretical and experimental study of the effect of cross-correlations. *Computer Vision, Graphics, and Image Processing: Image Understanding,* 60(3), 359-370.
Thomson, W D., & Saunders, J. (1997). The perception of flicker on raster-scanned displays. *Human Factors,* 39(l), 48-66.
Trezona, P W (1998). Theoretical aspects of mesopic photometry and their implication in data assessment and investigation planning. *Color Research and Application*, 23(5), 264-273.
Valluri, S., Minkovitz, J. B., Budak, K., Essary, L. R., Walker, R. S., Chansue, E., Cabrera, G. M., Koch, D. D., & Pepose. J. S. (1999). Comparative corneal topography and refractive variables in monozygotic and dizygotic twins. *American Journal of Ophthalmology,* 127(2), 158-163.
van der Steen, F. A. M. (1998). An earth-stationary perceived visual scene during roll and yaw motions in a flight simulator. *Journal of Vestibular Research,* 8(6), 411-425.
Warner, H. D., & Casey, E. P (1995). *Flight simulator visual system research and development: Bibliography of support provided by the Aircrew Training Research Division* (Tech. Report AL1HRTR-1995-0107). Mesa, Arizona: Air Force Research Laboratory, Warfighter Training Research Division.
Wittgenstein, L. (1958). *Philosophical investigations* (2nd ed.). Oxford: Basil Blackwell & Mott, Ltd.
Young, T Y., & Calvert, T W . (1974). *Classification, estimation, arid pattern recognition.* New York: American Elsevier.