A Beginner's Guide to the Human Field of View*

John P. Costella†
School of Physics, The University of Melbourne, Parkville, Victoria 3052, Australia

(20 November 1995)

Abstract

I describe a simple kitchen-table construction that allows one to get a good feel for the field of view of the human eye.

In considering the design and construction of head-mounted, immersive display devices, it is useful to know how much solid angle is actually viewable by the eye—the "field of view". Over the decades, the capabilities of the human visual system have been investigated by researchers in meticulous detail; and we may call on this research to give us precise quantitative answers to almost any question that we might wish to ask. But it can be difficult for designers of virtual worlds to track down this information in the published medical literature; and, even when found, it can be difficult for non-experts to "get a good feeling" for the tables and graphs presented. To do so, one needs an accurate way of portraying *solid angles*. Unfortunately, it is intrinsically impossible to represent solid angles without distortion on a flat sheet of paper, much less by describing it in words. Seeing as that we are still far from having "virtual reality on every desk", it is also not currently possible to use the technology itself to portray this information.

The reader is therefore asked to procure the following pieces of hardware so that a mathematical construction may be carried out: a mathematical compass—preferably one with a device that locks the arms after being set in position; one red and one blue ballpoint pen, that both fit in the compass; a texta (thick-tipped fibre marker), or a whiteboard marker; a ruler, marked in millimetres (or, in the US, down to at least sixteenths of an inch); a simple calculator; a pair of scissors, or a knife; an orange, or a bald tennis ball; a piece of string, long enough to circumnavigate the orange or tennis ball; and a roll of sticky tape.

The first task is to wrap a few turns of sticky tape around the sharp point of the compass's "pivot" arm, slightly extending past the point. This is to avoid puncturing the orange or tennis ball (as appropriate); orange-equipped readers that like orange juice, and do not have any objections to licking their apparatus, may omit this step.

The next task is to verify that the orange or tennis ball is as close to spherical as possible for objects of their type, and suitable for being written on by the ballpoint pens. If this proves to be the case, pick up the piece of string and wrap it around the orange or ball; do not worry if the string does not yet follow a great circle. Place your right thumb on top of the string near where it overlaps itself (but not on top of that point). With the fingers of your left hand, roll the far side of the string so that the string become more taut; let it slip under your right thumb only gradually; but make sure that no parts of the string "grab" the surface of the orange or ball (except where your right thumb is holding it!). After the rolled string passes through a great circle, it will become loose again (or may even roll right off). Without letting go with the right thumb, mark the point on the string where it crosses its other end with the texta. Now put the orange or ball down, cut the string at the mark, and dispose of the part that did not participate in the circumnavigation. Fold the string in half, pulling it taut to align the ends. At the half-way fold, mark it with the texta. Then fold it in half again, and mark the two unmarked folds with the texta. On unravelling the string, there should be three texta marks, indicating the quarter, half and three-quarter points along its length. Now pull the string taut along the ruler and measure its length. This is the circumference of the orange or ball; store its value in the calculator's memory (if it has one), or else write it down: we will use it later.

We now define some geographical names for places on our sphere, by analogy with the surface of the Earth. For orange-equipped readers, the North Pole of the orange will be defined as the point where the stem was attached. For tennis-ball-equipped readers, a mark should be made arbitrarily to denote the North Pole. Mark this pole with a small 'N' with the blue ballpoint pen. Similarly, the small mark 180° from the North Pole on an orange's surface will be termed the South Pole. Tennis-ballers, however, will need to use the piece of string to find this point, as follows: Wind the string around the ball, placing the two ends at the North Pole, and, as before, roll it carefully until it is taut; inspect it from the side to ensure that it appears to "dissect" the ball into equal halves. The half-way texta mark on the string is now over the South Pole; mark this spot on the ball with the blue ballpoint pen.

From this point, it will be assumed, for definiteness, that the object is an orange; possessors of tennis balls can project a mental image of the texture of an orange onto the surface of their ball, if so desired. Place the string around the orange, making sure the North and South Poles are aligned. (If possessors of oranges find, at this point, that the mark on the orange is not at the half-way point on the string, then either mark a new South Pole to agree with the string, or get another orange.) Now find the one-quarter and three-quarter points on the string, and use the texta to make marks on the orange at these two points. Put the string down. Choose one of the two points just marked—the one whose surrounding area is most amenable to being written on. This point shall be called *Singapore* (being a recognisable city, near the Equator, close to the centre of conventional planar maps of the world); write a small 'S' on the orange with the pen next to it. (This mark cannot be confused with the South Pole, since it is not diametrically opposite the North!) The marked point diametrically opposite Singapore will be called *Quito*; it may be labelled, but will not be used much in the following. Next, wind the string around the orange, so that it passes through all of the following points: the two Poles, Singapore and Quito. Use the blue ballpoint pen to trace around the circumference, completing a great circle through these points, taking care that the string is accurately aligned; this circle will be referred to henceforth as the *Central Meridian*. (If the pen does not write, wipe the orange's surface dry, get the ink flowing from the pen by scribbling on some paper, and try again. Two sets of ballpoint pens can make this procedure easier.) Now wrap the string around the orange, through the poles, but roughly 90° *around from* Singapore and Quito; *i.e.*, if viewing Singapore head-on, the string should now look like a circle surrounding the globe. This alignment need not be exact; the most important thing is that the string start and end on the North Pole, and pass over the South Pole. Make a small *ballpoint* mark at the one-quarter and three-quarter positions of the string. Now wind the string around the orange so that it passes over both of these two new points, as well as over both Singapore and Quito. Trace this line onto the orange. This is the Equator.

We are now in a position to start to relate this orangeography to our crude mapping of the human visual system. We shall imagine that the surface of the orange represents the solid angle seen by the viewer's eye, by imagining that the viewer's eye is located at the *centre* of the orange; the orange would (ideally) be a transparent sphere, fixed to the viewer's head, on which we would plot the extent of her view. Firstly, we shall consider the situation when the viewer is looking straight ahead at an object at infinity, with her head facing in the same direction. We shall, in this situation, define the direction of view (*i.e.*, the direction of *foveal view*—the most detailed vision in the centre of our vision) as being in the *Singapore* direction (with the North Pole towards the top of the head). One could imagine two oranges, side by side, one centred on each of the viewer's eyes, with both Singapores in the same direction; this would represent the viewer's direction of foveal view from each eye in this situation. Having defined a direction thus, the oranges should now be considered to be *fixed* to the viewer's head for the remainder of this section.

We now concentrate solely on the *left* eye of the viewer, and the corresponding orange surrounding it. We shall, in the following, be using the calculator to compute lengths on the ruler against which we shall set our compass. The compass will then be used to both measure "distances" between points inhabiting the curved surface of the orange, as well as to actually draw circles on the thing.

The first task is to compute how long 0.122 circumferences is. (For example, if the circumference of the orange was 247 mm, punch "$247 \times 0.122 =$" on the calculator, which yields the answer 30.134; ignore fractions of millimetres. Readers using Imperial units will need to convert back and forth between fractions of an inch.) Put the *red* ballpoint pen into the compass, and set its arms, with tips against the ruler, so that the tips are separated by this length of 0.122 circumferences (whatever the calculator gave for that answer). Now centre the *pivot* of the compass on Singapore, and draw a circle on the orange with the red pen (which is often easier said than done, but *is* possible). The solid angle enclosed by this red circle (subtended, as always, at the centre of the orange—where the viewer's eye is assumed to be located) indicates, in rough terms, all of the possible directions that our viewer can "look directly at"; in other words, the muscles of her eye can rotate her eyeball so that any direction in this solid angle is brought into central foveal view.

It will be noted that, all in all, this solid angle of foveal view is not too badly "curved", when looked at in three dimensions. Place a *flat plane* (such as a book) against the orange, so that it touches the orange at Singapore. One could imagine cutting out the peel of the orange around this red circle, and "flattening it" onto the plane of the book without too much trouble; the outer areas would be stretched (or, if dry and brittle, would fracture), but overall the correspondence between the section of peel and the planar surface is not too bad. (Readers possessing multiple oranges, who do not mind going through the above procedure a second time, might actually like to try this peel-cutting and -flattening exercise.) The surface of the plane corresponding to the flattened peel corresponds, roughly, to the maximum (apparent) size that a traditional, desk-mounted graphics screen can use: any larger than this and the user would need to *move her head* to be able to focus on all parts of the screen—a somewhat tiring requirement for everyday computing tasks. Thus, it can be seen that it is *the very structure of our eyes* that allows flat display devices to be so successful: any device subtending a small enough solid angle that all points can be "read" (*i.e.*, viewed in fine detail) without gross movements of one's head cannot have problems of "wrap-around" anyway.

Virtual reality systems, of course, have completely different goals: the display device is not considered to be a "readable" object, as it is in traditional computing environments—rather, it is meant to be a convincing, *completely immersive* stimulation of our visual senses. In such an environment, *peripheral vision* is of vital importance, for two reasons. Firstly, and most obviously, the convinceability of the virtual reality session will suffer if the participant "has blinkers on" (except, of course, for the singular case in which one is trying to give normally-sighted people an idea of what various sight disabilities look like from the sufferer's point of view). Secondly, and quite possibly more importantly, is the fact that, although our peripheral vision is no good for *detailed* work, it is especially good at detecting *motion*. Such feats are not of very much use in traditional computing environments, but are vital for a virtual reality participant to (quite literally) "get one's bearings" in the spatially-extended immersive environment. We must therefore get some sort of feeling—using our orange-mapped model—for the range of human peripheral vision, so that we might cater for it satisfactorily in our hardware and software implementations.

The following construction will be a little more complicated than the earlier ones; a "check-list" will be presented at the end of it so that the reader can verify that it has been carried out correctly. Firstly, set the compass tip-separation to a distance (measured, as always, on the ruler) equal to 0.048 circumferences (a relatively small distance). Place the pivot on Singapore, and mark off the position to the *east* (right) of Singapore where the pen arm intersects the Equator. We are now somewhere near the Makassar Strait. Now put the *blue* ballpoint pen into the compass, and set its tip-to-tip distance to 0.2 circumferences; this is quite a large span. Now, *with the pivot on the new point marked in the Makassar Strait*, carefully draw a circle on the orange. The large portion of solid angle enclosed by this circle represents, in rough terms, the range of our peripheral vision. To check that the construction has been carried out correctly, measure the following distances, by placing the compass tips on the two points mentioned, and then measuring the tip-to-tip distance on the ruler: From Singapore to the point where this freshly-drawn circle cuts the Central Meridian (either to the north or south): about 0.18 circumferences; from Singapore to the point where the circle cuts the Equator to the *west*: about 0.16 circumferences; from Singapore to the point where the circle cuts the Equator to the *east*: about 0.23 circumferences. If these are roughly correct, one can, in fact, also check that the *earlier*, foveal construction is correct, by measuring the tip-to-tip distance between the red and blue circles where they cross the Equator and Central Meridian. To the west, this distance should be about 0.04 circumferences; to the east, about 0.13 circumferences; to the north and south, about 0.08 circumferences.

There are several features of the range of our peripheral vision that we can now note. Firstly, it can be seen that our eyes can actually look a little *behind* us—the left eye to the left, and the right eye to the right. (Wrap the string around the orange, through the poles, 90° east of Singapore; the peripheral field of view cuts behind the forward hemisphere.) To verify that this is, indeed, correct, it suffices to stand with one's face against a tall picket fence; with one's nose between the pickets, and looking straight ahead, you can still see things going on on *your* side of the fence! That our field of vision is configured in this way is is not too surprising when it is noted that one's *nose* blocks one's field of view in the other direction (but which is, of course, picked up by the other eye); hence, our eyes and faces have been appropriately angled so that our eyes' fields of vision are put to most use. But this also means that it is *not* possible to project the view from both eyes onto *one single* plane directly in front of our faces. Of course, we would intend to use two planar displays anyway, to provide stereopsis; but this observation is important when it comes to *clipping* objects for display (for which a single planar clip, to cover *both* eyes, is thus impossible).

Secondly, we note that a reasonable approximation is to consider the point near the Makassar Strait to be the "centre" of a circle of field of view for the left eye. (This unproved assertion by the author may be verified as approximately correct by an examination of ocular data.) The field of view extends about 75°–80° in each direction from this point, or about a 150°–160° "span" along any great circle. This is, of course, a little shy of the full 180° that a hemispherical view provides, but not by much (as an examination of the orange reveals). Now imagine that one was to cut the peel of the orange out around the large (blue) circle of peripheral vision, and were then to try and "flatten it" onto a plane. It would be a mess! An almost-hemisphere is something that does not map well to a flat plane.

Let us reconsider, nevertheless, our earlier idea of placing a planar display device in front of each eye (with appropriate optics for focusing). How would such a device perform? Let us ignore, for the moment, the actual size of the device, and merely consider the *density of pixels* in any direction of sight. For this purpose, let us assume that the pixels are laid out on a regular rectangular grid (as is the case for real-life display devices). Let the distance between the eye and the plane (or the effective plane, when employing optics) be R pixel widths (*i.e.*, we are using the pixel width as a unit of distance, not only in the display plane, but also in the orthogonal direction). Let us measure first along the x axis of the display, which we shall take to have its origin at the point on the display pierced by a ray going through the eye and the Makassar Strait point (which may be visualised by placing a plane, representing the display, against the orange at the Makassar Strait). Let us imagine that we have a head-mountable display of resolution (say) 512×512, which we want to cover the entire field of view (with a little wastage in the corners). Taking the maximum field of view span from the Makassar Strait to be about 80°, we can now compute R from simple trigonometry. Drawing a triangle connecting the eye, the Makassar Strait and the extremum pixel on the axis, and insisting that this extrumum pixel be just at the edge of the field of view (*viz.* 80° away), we obtain

$$\tan 80^\circ = \frac{256}{R},$$

where we have noted that half of the display will extend to the other side of the origin, and hence the centre-to-extremum distance in the plane is 256 pixels. Computing the solution to this equation, we have

$$R = \frac{256}{\tan 80^\circ} \approx 45.1 \text{ pixels.}$$

For a planar device of resolution 512×512 to just cover the entire field of view, this eye–plane distance is absolute; virtual reality displays, unlike traditional computing environments, by their nature specify *exactly* how far away the display plane must be; this is no longer a freedom of choice. The most important question then arises: how good a resolution does this represent, if we simply look straight ahead? Well, in this direction, we have the approximation $dx \approx R\,d\theta$, where $d\theta$ is the linear angle (in radians) subtended by the pixel of width dx. Since the pixel width dx is, by definition, 1 pixel-width (our unit of distance), we therefore find that

$$d\theta \approx \frac{dx}{R} \approx \frac{1}{45.1} \approx 0.0222 \text{ radians} \approx 1.27^\circ.$$

So how big is a 1.27°-wide pixel? Let us compare this with familiar quantities. The pixel at the centre of the 640×480 VGA display mode, viewed on a standard IBM 8513 monitor from a distance of half a metre, subtends an angle of about 2.2 arc-*minutes*, which is approaching the limit of our foveal vision. Thus, our planar-display wrap-around central pixel looks about the same as a 35×35 square does on a VGA 640×480 display—it is huge! But *why* is it so huge? After all, 512 pixels of a VGA display subtend about 18.6° of arc at the same half-metre viewing distance; stretch this out to cover 160° instead and each pixel should get bigger in each direction by a factor of about $160/18.6 = 8.6$. So why are we getting an answer that is four times worse in linear dimensions (or sixteen times worse in terms of area) than this argument suggests?

The answer comes from considering, as a simple example, the very *edgemost* pixel in the x-direction on the display—the one that is 80° away from the Makassar Strait point, just visible on the edge of our viewer's peripheral vision. How much angle does *this* pixel subtend? An out-of-hat answer would be "the same as the one in the middle"—after all, the rectangular grid is regularly spaced. But this reasoning relies, for its approximate truth, on the fact that *conventional* planar displays only subtend a small total angle anyway. On the contrary, we are now talking about a planar device subtending a *massive* angle; paraxial approximations are rubbish in this environment. Instead, consider the abovementioned relation,

$$x = R \tan\theta,$$

where $R = 45.1$ pixels in our example, and θ is the angle between the pixel's apparent position and the Makassar Strait. Inverting this relationship we have

$$\theta = \arctan\left(\frac{x}{R}\right),$$

and, on differentiation,

$$d\theta = \frac{dx}{\sqrt{x^2 + R^2}}.$$

For $dx = 1$ at $x = 0$ we find $d\theta \approx 1/45.1 \approx 0.0222$ radians, or 1.27°, as before; this only applies at $x = 0$. For $dx = 1$ at $x = 256$, on the other hand, we find that

$$d\theta = \frac{1}{\sqrt{256^2 + 45.1^2}} \approx 0.00385 \text{ radians} \approx 0.22^\circ.$$

Thus, the pixels at the outermost edges of the visible display are subtending a linear angle nearly six times *smaller* than at the centre of the display—or, in other words, it would take about 33 extremity-pixels to subtend the same solid angle as the *central* pixel! But this is crazy! We already know that the peripheral vision, outside the range of foveal tracking (the red circle on the orange), is sensitive to significantly *less* detail that the fovea! Why on earth are we putting such a high pixel density out there? Who ordered that? Sack him!

The reason for this waste of resolution, of course, is that we have tried to stretch a planar device across our field of view. What is perhaps not so obvious is the fact that no amount of technology, no amount of optical trickery, can remove this problem: *it is an inherent flaw in trying to use a planar display as a flat window on the virtual world.* This point is so important that it shall be repeated in a slightly different form: *Any rendering algorithm that assumes a regularly-spaced planar display device will create a central pixel 33 times the size of a peripheral pixel under virtual reality conditions.* This is not open for discussion; it is a mathematical fact.

Let us, however, consider whether we might not, ultimately, avoid this problem by simply producing a higher-resolution display. Of course, one can always compensate for the factor of four degradation by increasing the linear resolution of the device in each direction by this factor. However, there is a more disturbing psychological property of the planar display: pixels near the centre of view seem chunkier than those further out; it becomes psychologically preferable to *look askew* at objects, using the eye muscles together with the increased outer resolution to get a better view. This property is worrying, and could, potentially, cause eye-strain. To avoid this side-effect, we would need to have the *entire* display of such a high resolution that even the central pixel (the worst one) is below foveal resolution. We earlier showed that a 512-pixel-wide planar display produces a central pixel angular length of about 1.27°. Foveal resolution, however, is about 1 arc-minute at best—about what one can get from a Super-VGA display at normal viewing distance. To achieve this, we could simply "scale down" our pixel size, providing a greater number of total pixels to maintain the full-field-of-vision coverage; performing this calculation tells us that we would need

$$512 \times \frac{1.27}{1/60} \approx 39000$$

pixels in each direction for this purpose. No problems—just give me a 39000×39000 head-mountable display device, and I'll shake your hand (if I can find it, behind the 4.5 GB of display memory necessary for even 24-bit colour, with no z-buffering or other sophistication ...). And all this just to get roughly Super-VGA resolution ... do you think you have a market?

The clear message from this line of thought is the following: *planar-display rendering has no future in virtual reality.* Assumption of a planar view-plane is, in traditional computer graphics applications, convenient: the perspective view of a line is again a line; the perspective view of a polygon is also a polygon. Everything is cosy for algorithm-inventors; procedures can be optimised to a great extent. But its performance is simply unacceptable for wrap-around display devices. We must farewell an old friend.

References

* This paper has been taken directly from Sec. 4.1 of [1].

† jpc@physics.unimelb.edu.au; http://www.ph.unimelb.edu.au/~jpc.

[1] J. P. Costella, "Galilean Antialiasing for Virtual Reality Displays," sci.virtual-worlds, 4 November 1992; University of Melbourne preprint UM–P–92/105 (28 October 1992); http://www.ph.unimelb.edu.au/~jpc/papers/gaa.ps.