

# Introduction to 3D

## a Simple Stereography Primer

By Markus Stone

### Markus Stone | Director


*This article originally formed part of the production bible for the movie Bait 3D.*

I'm simplifying things a lot here, but hey, that's the point of this document; to be an easy introduction to 3D so that everyone can get their head around it and so we can all talk about it in practical terms. With that in mind, I will be using the terms 'Stereoscopic' and '3D' interchangeably. Live with it. I'm not going to start an article who's goal is accessibility with a long boring discussion on terminology. I'll explain the difference at the end.


#### The bit you probably know already.

The basic thing you need to know with stereo is that you have two cameras separated from each other in the horizontal plane – same as your eyes.

When you look at something, both your eyes turn in a little bit to look at it (as well as pulling focus). This is called converging. Our cameras do the same thing. Where the different views meet (cross over) is called the plane of convergence.



When you play these images back on a 3D screen, the plane of convergence will correspond to where the screen is.




Objects at this point will NOT appear blurry, but can be watched quite comfortably without glasses. Objects that are placed either further away or closer to camera than the plane of convergence will begin to have separation on the screen, and will become two discrete images (left and right) instead of one. The distance between these images is called parallax.

Separate left and right images



Stereo overlapped image (Anaglyph)



In the Anaglyph image above you can see the difference between the left and right images. Where they overlap, the red and cyan images add together to form a black and white image, which represents the plane of convergence. Where they are different, you will see coloured 'fringes' which represent the amount of background parallax.

### **Limits of Parallax**

As you move things further away from the plane of convergence they display more and more separation (parallax). But there are limits to how far you can go with this.

The most distant object you can look at is a star. When you look at a star, your eyes are not converging at all – your eyes are parallel. There is nothing you can look at in the physical world that will cause your eyes to point outward (away from each other). So unless you want your audience to end up looking like Marty Feldman, you need to control how much parallax there will be on the background.



Marty Feldman, Patron Saint of 3D

You do this by moving the cameras closer or further apart. If the cameras are too close, you don't get much depth and the image tends towards 2D. Too much, and it causes people serious eyestrain.


For objects that are in front of the plane of convergence, there is also a limit to how close they can come to camera before the eye has trouble converging on them (try looking at your nose! Both eyes can see it, but you can't focus on it). At this point the only projection into the audience you'll get is projectile vomiting as your audience develops a headache en masse and stampedes for the door.

The practical parallax limits for shooting 3D are somewhat different from the limits your eyes have in normal life (for a bunch of boring technical reasons I won't go into), but the important thing to remember is that for shooting 3D there are limits to the amount of parallax we can have for both 'into the screen' (positive) and 'into the audience' (negative) spaces.

When we talk about parallax, we typically talk about it as a percentage of screen width because it applies to any screen, and doesn't depend on the resolution. Parallax at 1% of screen width in 2K works out to 20.48 pixels, or 10cm on a 10m screen.

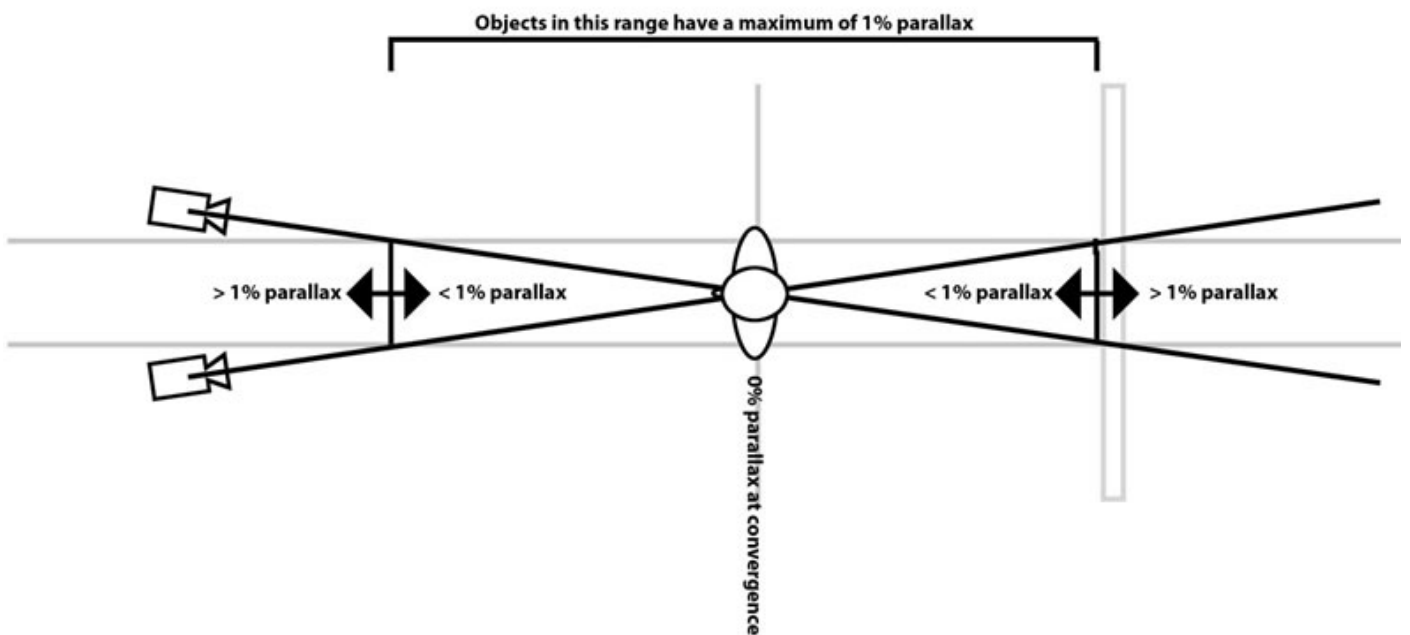
### **The Patented Markus Stone Chopsticks Model for thinking about 3D**

If you imagine two cameras converged on the subject of the scene, and then you move the cameras closer together while staying converged on the subject. You can visualise that as a pair of scissors or chopsticks with a rubber band around the middle opening and closing with the pivot centered on the plane of subject.



If you open up the parallax in front of the convergence point, you end up with more parallax behind as well. You can't change the amount of parallax on one end without applying the same change in parallax to the other.


If we then put in our limits (lets say for simplicity's sake we call it maximum 1% parallax foreground and background) we can make a pair of lines that represent this limit.




It's kind of like opening a pair of scissors inside a cardboard tube – at some point you're going to hit the edges of the tube.

### Where it gets interesting


If you look at the diagram above, you'll notice the optical axes of the cameras cross over at the plane of convergence (the person/subject) and cross the parallax limits at the same point as the flat does. This means that we are achieving our maximum 1% parallax in the background. At this point if we move the flat away from camera, the optical axes of the cameras will be outside of our limits –therefore we are getting too much parallax.




The solution is to reduce the distance between the cameras in order to keep the parallax on the background flat within our limits.



Now, lets say that we want the subject to appear bigger in frame, so we move the cameras forward.




If we reconverge on the subject, we'll find that because it is now closer, our background parallax has widened out to exceed our limits.



If you move the camera closer to the subject, you need to move the cameras closer together .

In fact to maintain the same parallax as we move the cameras closer and further away from the subject, they end up pretty much moving along the same line of optical axis we have already drawn – that way they keep the same angle (the cameras don't have to be within our 1% limit line, just the subjects in the frame).



You could also zoom in to make the subject bigger. Considering that zooming in is the same as magnifying the middle part of the image, we can see that a 2 x zoom will double the parallax in the image. Essentially the 1% line we have been working with for deciding our parallax limits becomes effectively a 2% line, so again, we again need to bring our cameras closer together to control the parallax.

So to recap;


How far apart the cameras are depends upon;

- Distance to the subject
- Volume of space being captured (distance from the subject to the background)
- Focal Length

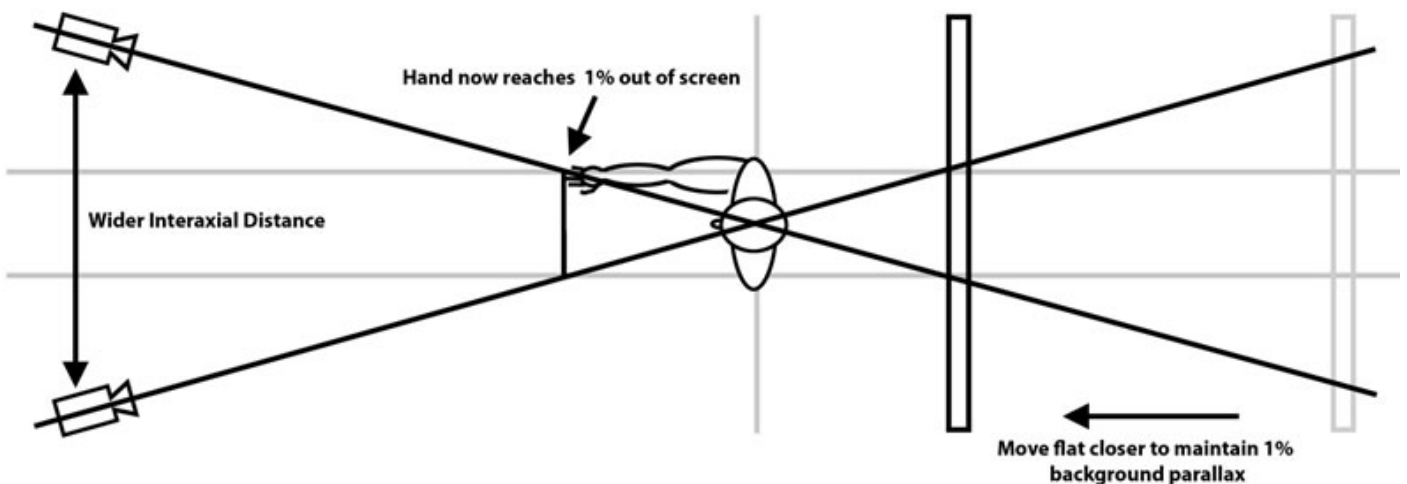


## Where it gets tricky


It gets interesting where you also want something to come off the screen into the audience. Say our actor holds out their arm towards the camera. If you want the actor's hand to have the maximum projection off the screen, it would need to reach exactly to the point where the camera axes cross the 1% parallax barrier. Rarely do these distances coincide.




If the actor's hand is shorter than the optimum distance, you could widen the interaxial (distance between the cameras), but then the background parallax could blow out unless you also move the subject and camera closer to the background to control this.



Alternatively if you're in close and the hand projects out too far, you could reduce the parallax to cope, but then you're getting less than the optimum amount of 3D in the background.



Best bet here is to move the background further away or the camera/subject further from the background as a unit).



In either of these cases, you may find that you can't move the background (ie – if it's the horizon) and if the shot size is critical, you can't move the camera as this changes the size of the object in

frame.

Without being able to change the length of the actor's arm, or the distance to the horizon, your only other option is to position the actor such that they occupy the perfect position according to the length of their arm – but obviously this changes how big the subject appears in frame. Moving back and zooming in doesn't help, unfortunately. As the two effects cancel each other out and you're left with exactly the same problem. Thing to be aware of here is that if the positioning of the subject, background and foreground (protruding) object are fixed, along with the shot size, there is a good chance that a compromise will have to be made – either in terms of having less than the maximum projection from the screen, or less than the optimum amount of background parallax.

It's also worth mentioning that the distances to camera are critical for protruding objects – bringing an object just 10cm closer to camera can easily add another 1% to the parallax and push it over the edge into unwatchable territory.

The other, and perhaps more effective solution is to simply shoot the foreground element against a greenscreen. You can then max out the parallax on the foreground and dial in an appropriate background parallax in post.

## **Breaking Windows**

Watching a 3D presentation is like looking through a window. The window pane corresponds to the edge of the screen and we look through it to the world beyond. Where there is an object on the edge of the frame, we expect it to be occluded by the edge of the window as it would if you were looking out your kitchen window. With the example of an object hugging the left side of the frame, the right eye will see more of the object, just like your right eye would see around the window pane more in real life.

Things get interesting when you have an object that is in negative parallax (audience) space that breaks the edge of the frame. Let's say we take that same object and move it towards us so that it is still breaking the edge of the frame, but is protruding into audience space. Now the opposite happens. Now we can see more of the object in the left eye, and less in the right eye. As well as this, the object is being obscured by the edge of the window that is behind it. Neither of these situations occur in the real world, and our brains lack the ability to process it, resulting in a line of brain shear down the edge of these images wherever an object forward of the plane of convergence breaks the edge of frame. The size of the problem depends directly upon how much parallax there is in the foreground object.

## **Floating Windows**

You can fix this by cutting off the left hand side of the left eye (literally cropping off the left side of the frame) until the object in the foreground is cut off at the same point in both shots. Without your glasses on you will see a black line down one side of the screen, but with your glasses on you will never pick it. This is what's known as floating the window and is a common technique (was

used all the way through 'G-Force'). A side effect of this technique is that it sucks everything back into the screen so the object that was trying to project out of the screen now feels as if it is no closer than the plane of convergence, so it's mostly used as a fix for objects that give us a case of the edge uglies, although you can use this technique to borrow FG parallax to give greater depth to the shot.

## **Violating the Old School**

The old school way of shooting was not to let anything violate the edge of frame in the foreground. So if you had a shot of a man standing in a field, the closest object in the frame would be the ground, not the man, therefore convergence would be set to the closest part of the foreground so as to avoid edge violations. In the final edited program the viewer's eyes would be converging back and forth between the subject successive shots that could be at different distances from the viewer. If you hold your finger up close to your nose, then try to focus on a distant object, you will notice it takes a moment for the eye to reconverge on the distant object. The ramifications for editing are obvious - no quick editing.

The more recent method is to always converge on the subject of the scene so that as you cut from shot to shot your eyes are already converged on the subject of the incoming shot which means you can cut quicker.

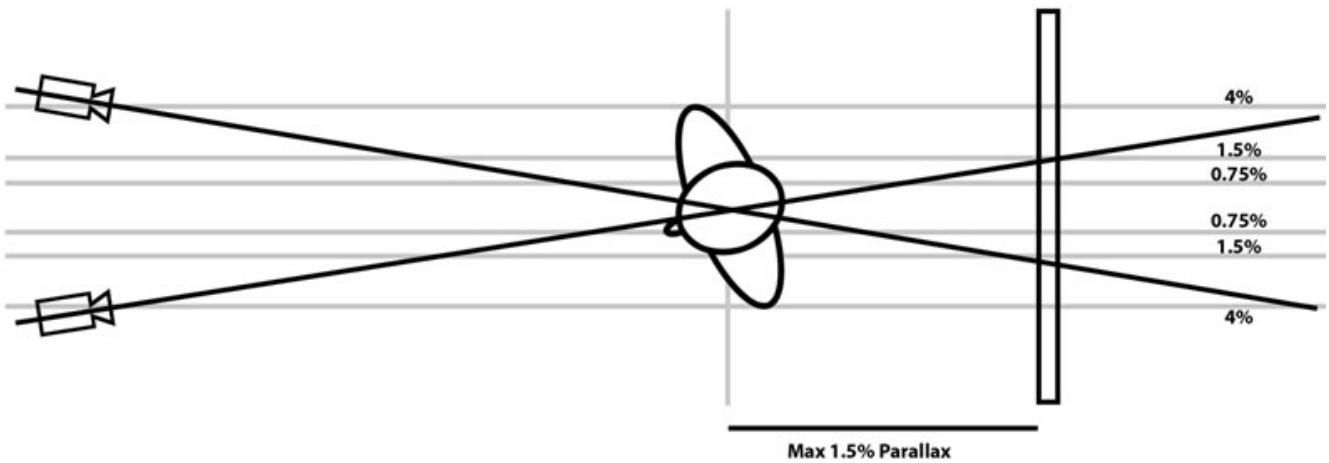
## **The New School**

The newer method of always converging on the subject raises some obvious issues. What do you with an over the shoulder shot where the shoulder is obviously going to be projecting into the audience and violating the edge of the screen? You could float the window, but doing this for a large percentage of the shots could be time consuming in post. A common way of handling this is to keep the foreground parallax below a certain threshold limit, making the edge violation relatively minor. By lighting the foreground subject a little darker so it doesn't draw the eye, when combined with a well-paced edit the viewer's eyes never see the edge violation (anyone notice them all the way through Avatatar?). If no-one 's looking at it, why fix it?

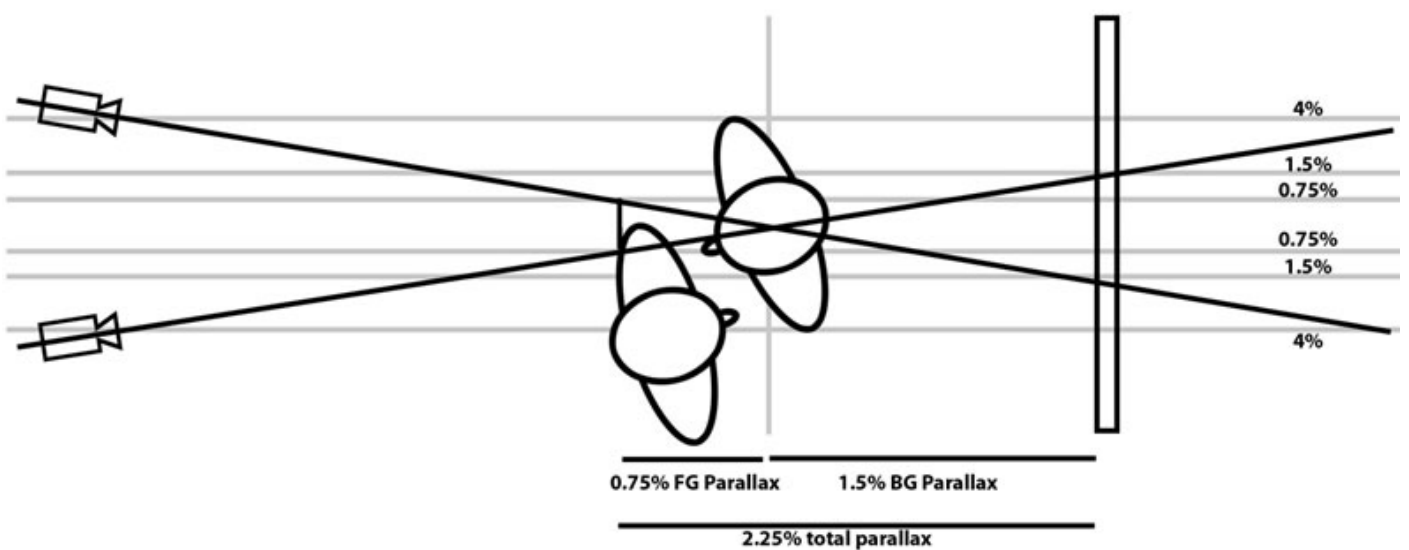
## **Problems at school**

In our earlier examples I spoke about the limits of parallax as 1% of screen for the foreground and background. In actual fact for a feature film you'd normally be talking maximums of 1.5-2% background parallax, 4-5% foreground parallax for a foreground object that doesn't break the frame and 0.5 - 1% foreground parallax for one that does. Lets say we're working with BG=1.55, FG (frame breaking) = 0.75% and FG gag shots 4%.


The subject of the frame can be either in the foreground, mid ground, or background of any given shot. With the subject in the foreground, we are converged on the subject, which means there are no objects in audience space. Background parallax can be maximised out to 1.5%.



If the subject is in the mid-ground, the BG can still be maximised to 1.5%, while adding depth from the foreground elements \*as long as their parallax is less than 0.75% off the screen.\*



If the subject is in the background, essentially everything in the shot is coming off the screen – therefore you only have the 0.75% of foreground parallax to play with.



What all this means is that a shoot with the subject in the foreground will have a total parallax of 1.5%, midground shots will have a total parallax of 2.25%, and where the subject is in the BG, you will only have 0.75% to play with.

To potentially have a total screen depth that changes from 0.75 to 2.25% total screen depth between shots makes a bit of a mockery of the idea of a depth script (that is, designing some parts of the film to have greater depth than others) because the depth varies so much shot to shot it destroys the concept. (you can achieve this using the old school method of 3D because everything is behind the plane of convergence. The effect of a depth script is subtle in any case as most people can't pick even quite large differences in interaxial anyway. It pretty much needs to be the same shot intercut on itself with different interaxials before most people see the difference. In other words, we perceive things as 3D quite well, but the stereoscopic image is only one of a bunch of other cues we use to determine depth; along with other standard 2D depth cues like diminishing size, converging lines, aerial perspective, etc. ( the human eye can only see actual 3D out to about 120m) By just changing the stereo depth we are only changing one factor that our brains use to determine depth, hence the inability of punters to really gauge accurate depth based on the stereography alone.

What 3D does give is an enhancement to the geography of the scene – a greater sense of space and of the immediacy of the action.

That being said, you can look at 3d as a continuum with a 0 interaxial giving a 2D image on the one hand, and too much interaxial melting people's brains on the other. There is a point where even a punter can see that the 3D effect is pretty subtle – generally around 0.4%. Our 0.7% mentioned above is different enough from the 1.5% that it would be noticeably different in terms of depth.

The solution is to shoot the shots were the subject is in the background with 1.5% foreground parallax and float the window to deal with the corresponding edge violations. A depth script, if required can be superimposed upon these parameters.

## The Difference Between '3D' and 'Stereoscopic'

As promised at the beginning of this article, I will now reveal the difference between 'stereographic' and '3D'. I avoided making the distinction at the beginning because this is primarily a piece of communication. If I started this article with a discussion about 'stereographic' v's '3D' there would be a page long precis on nomenclature and you would be softly snoring right now.

If you think about a stereo, it has two speakers that give a sense of depth to the sound, yet is not surround sound. Stereographic is to 3D what stereo speakers are to surround sound (or, more accurately, 'Holophonic' sound).

With stereography, you have a left and right image, just as you have two eyes. But that's all the information there is. You can't for example, peer behind an actor to see the background that is occluded by them. If you move your head while watching 3D, you have the very strange experience of having the world you are watching rotate to maintain the same orientation with you.

'True' 3D means there are an infinite number of viewing positions, such as in a Cad program or 3D Studio Max (or, as in real life!). With 3D, if you move your viewing position, you will see what was behind the actor's head.

However, much as the stereographers love to argue about the names of things, the fact is that no-one is going to start calling 3D movies 'stereographic movies' anytime soon. It simply doesn't market well. That particular ship has sailed. The most correct expression is 'S3D', or 'Stereographic 3D'. As long as you and I know what we're talking about that's what's important.

While we're talking about names for things, there is a debate about what that particular distance between the cameras is called. Some call it 'inter-ocular' distance, others call it 'inter-axial'. Unfortunately they're both wrong. Interocular refers to the distance between the eyes of the observer, not the cameras themselves. Interaxial refers to the axes of the cameras themselves, which sounds right, until you realise that an axis is a line in space that projects infinitely in either direction – the distance between any two crossed axes can be anything from 0 to infinity, depending upon where you measure it.

'Okay smarty pants, then where you you measure it from?' I hear you ask. Not from the sensor plane as you might think, as this is not the optical center of the camera. Panoramic photographers will know that to pan a camera so that it has absolutely no horizontal translation, you need to center not the sensor plane, but the lens node on the axis of rotation. The correct term then is 'inter-nodal'. This is another term that while correct, rolls off the tongue like a brick. That's why I don't expect everyone to start using it soon. As long as both you and the people you're talking to know what is meant that is the main thing.

That is the essence of 3D that I will cover here. If you have any further questions, do not hesitate to ask.