Welcome

nordic

td forum

Seminar on Stereoscopic Production for Computer Animation

Assumptions

- We are going to limit our discussions to pure computer graphic production. We will not be discussing any live action techniques or how to deal with photographic plates cameras or any other live action issues.
- We are assuming that all of you have some sort of production process so we are approaching this topic from the perspective of the best ways to integrate stereoscopic production into your existing pipelines.
- Everyone has an understanding of 3D production techniques.
- No matter what software you are using or how your pipeline works, the basic principles of stereoscopic production are applicable.
- Our software selection for the purposes of this seminar is not a recommendation or an endorsement, but rather a choice we made for our own convenience with regard to illustration and demonstration.

The Illusion of Depth

There are two main categories of visual depth cues.

Monocular Depth Cues

- Interposition
- Relative Height
- Relative Size (Size Consistency)
- Linear Perspective
- Texture Gradient
- Lighting Effects (Shadows)
- Linear Perspective
- Atmosphere (Aeriel Perspective)
- Motion Parallax
- Retinal Image Size (Requires a priori knowledge)
- Depth of Field Effects (This is a synthetic depth cue.)

Binocular Depth Cue

- Stereopsis

Monocular Depth Cues

Interposition

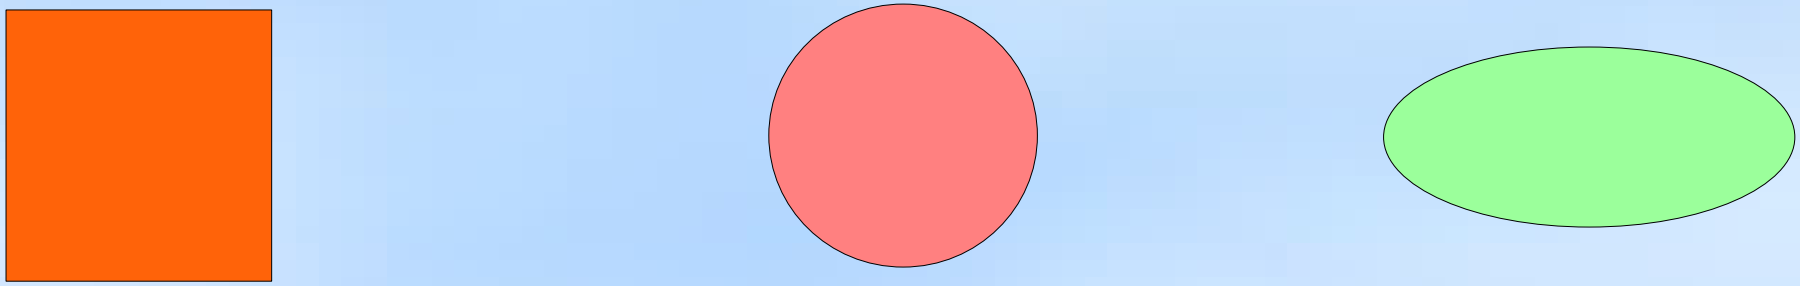

If one object occludes another, the object that is occluded is likely more distant from the observer.

Monocular Depth Cues

Interposition

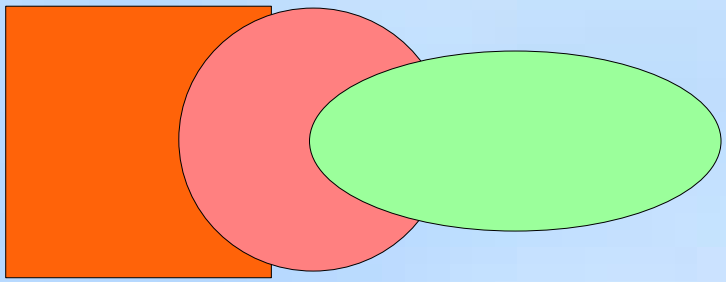

If one object occludes another, the object that is occluded is likely more distant from the observer.

Monocular Depth Cues

Relative Height

If one object is higher on the screen than other, it is likely more distant from the observer.

Monocular Depth Cues

Relative Size (Size Consistency)

Sometimes called Retinal Size

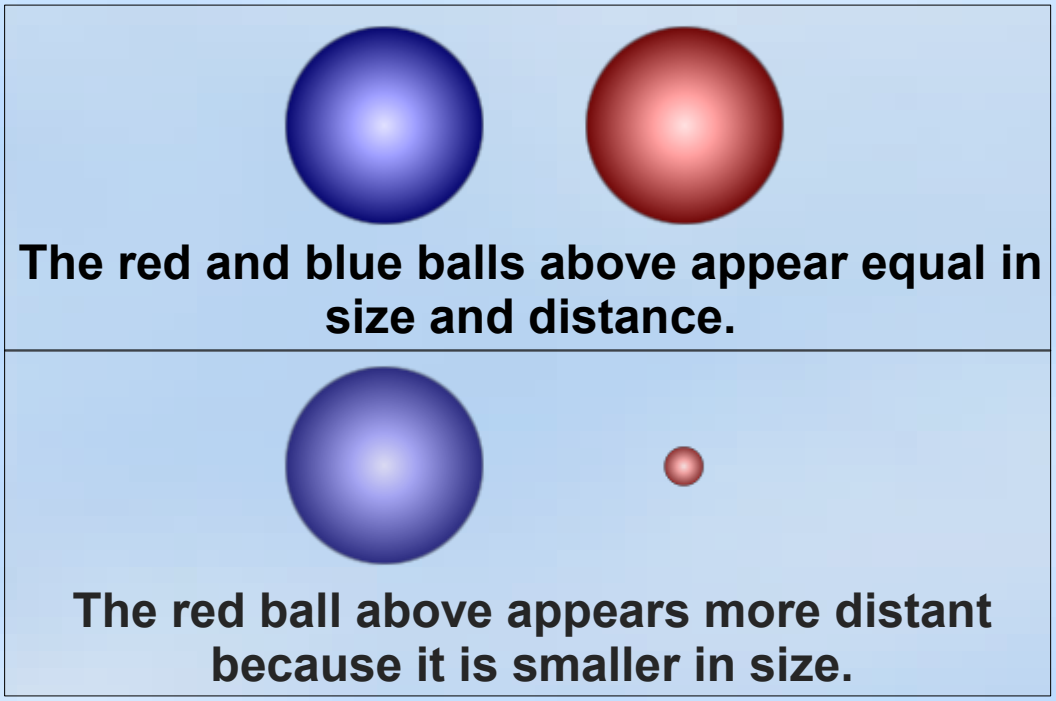

The red and blue balls above appear equal in size and distance.

The red ball above appears more distant because it is smaller in size.

An object is more likely to change position than to change size. If an object gets smaller, it is likely to be receding from the observer and if it gets larger, it is more likely to be approaching the observer.

The corollary is that smaller objects appear more distant than larger objects, especially if they are similar in other respects.

Monocular Depth Cues

Linear Perspective

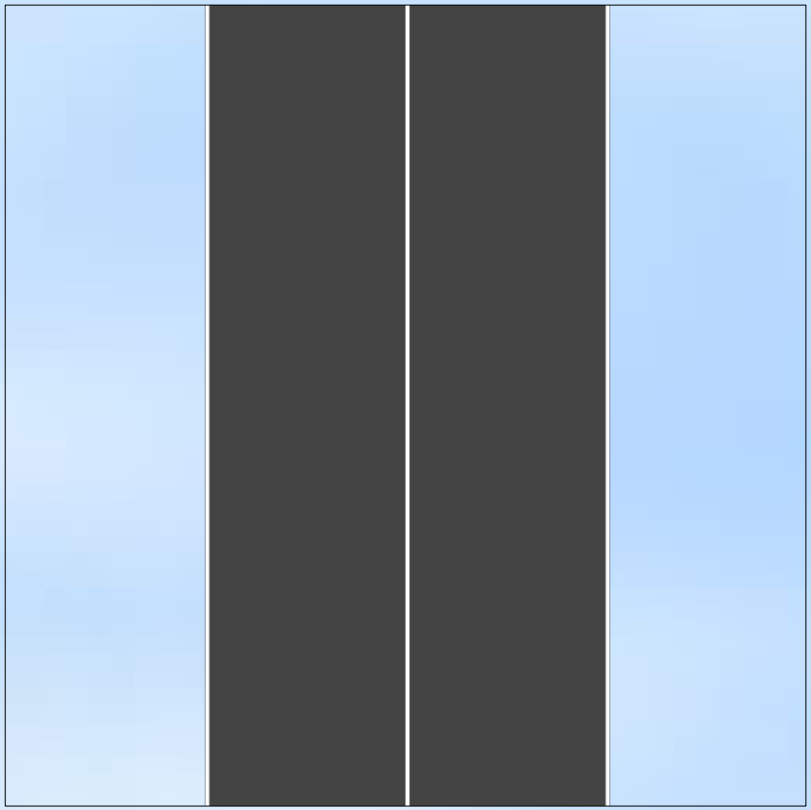

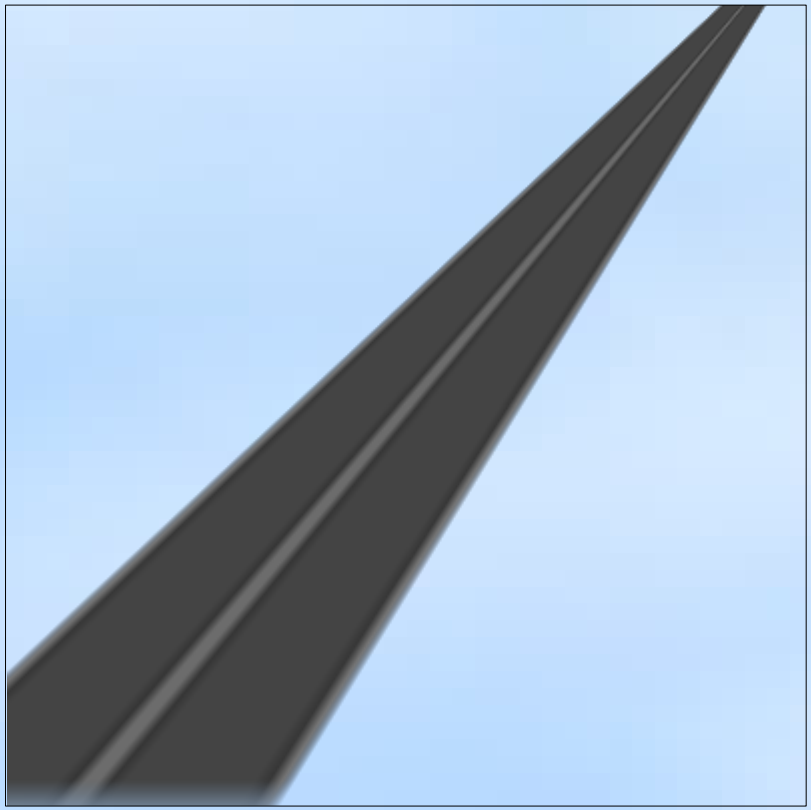

As lines converge, the apparent distance from the observer is inversely proportional to the amount of convergence.

Monocular Depth Cues
Texture Gradient

As textures get less detailed and more even and regular in tone, the greater the perceived distance from the observer.

Monocular Depth Cues
Lighting Effects

Shape and relief can be inferred from shading.

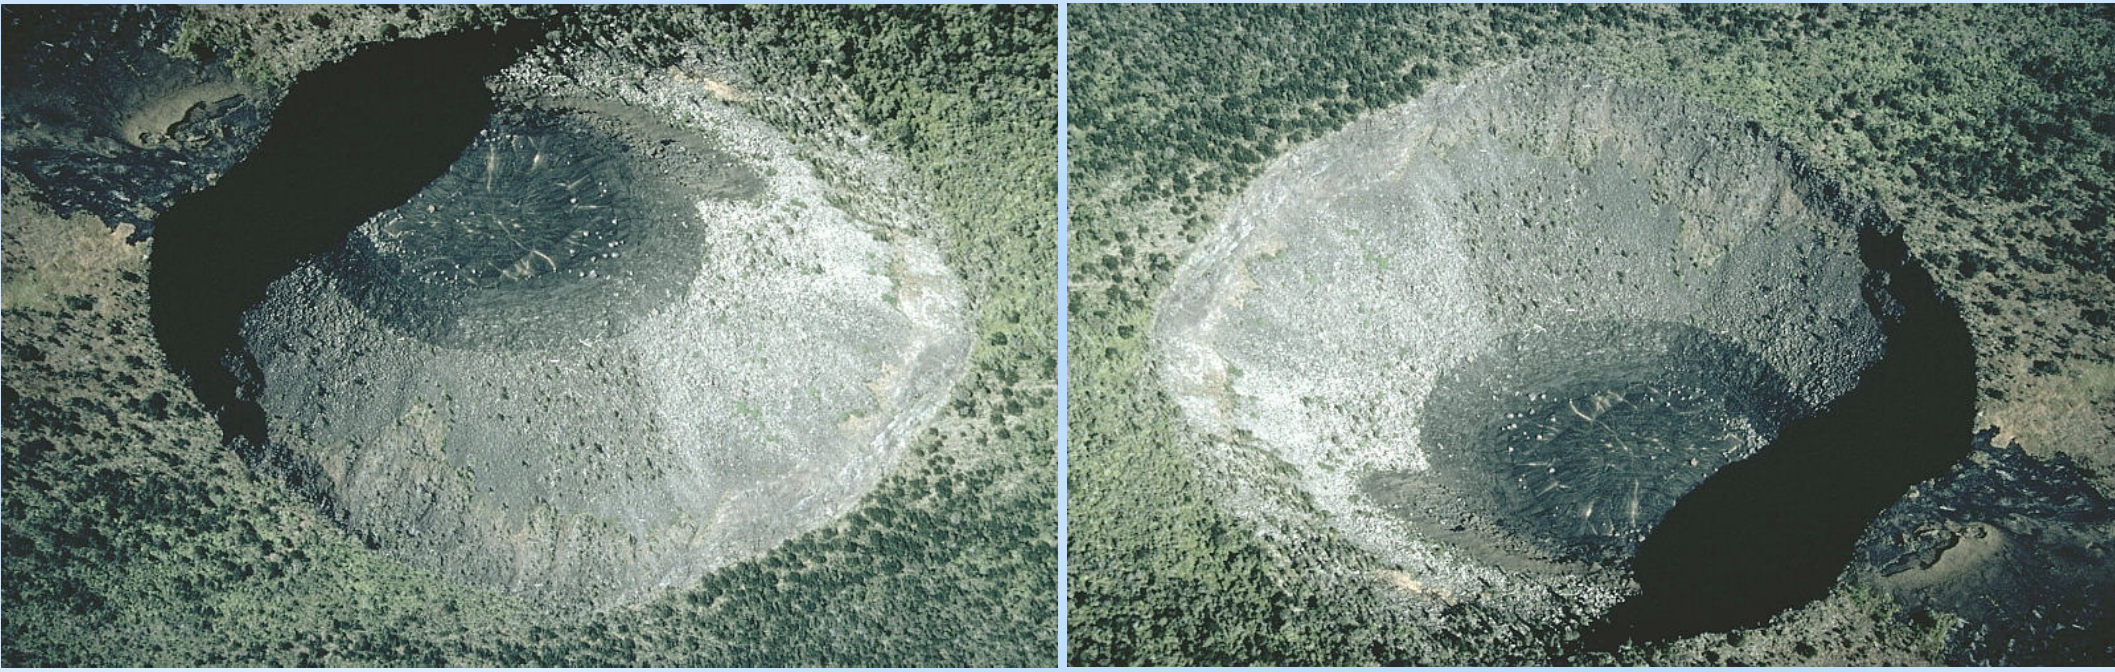

In this image, what appears to be a mound with a plateau at its summit when rotated 180° reveals itself to be a hole with a flat bottom.

Monocular Depth Cues
Atmosphere

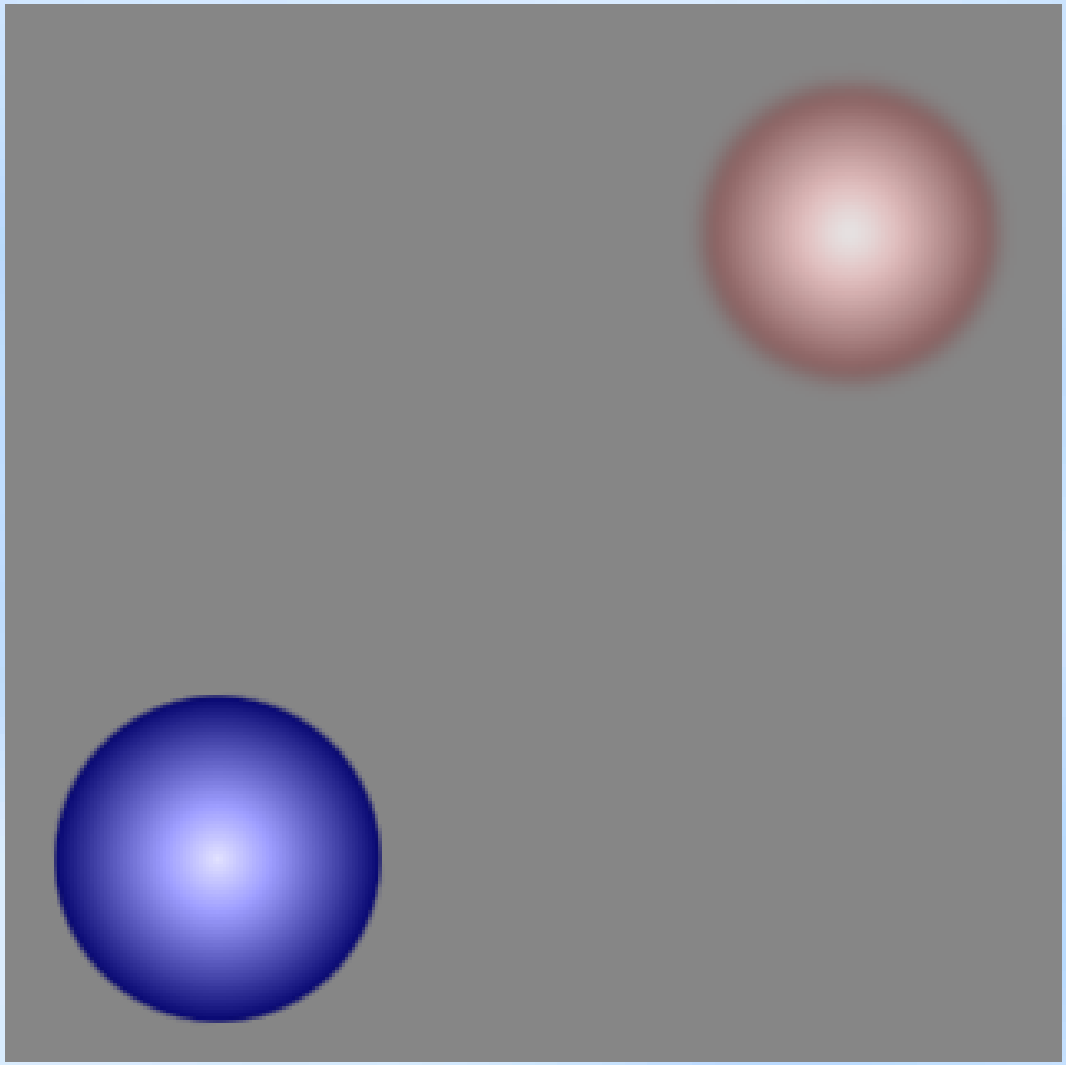

Distance can be inferred from atmospheric effects.

Monocular Depth Cues

Retinal Image Size

A priori knowledge of the known size of familiar objects can be used to infer distance.

Note how the perception of the walls and the lions change when you notice the man working on the scaffold in the background.

Monocular Depth Cues

Depth of Field

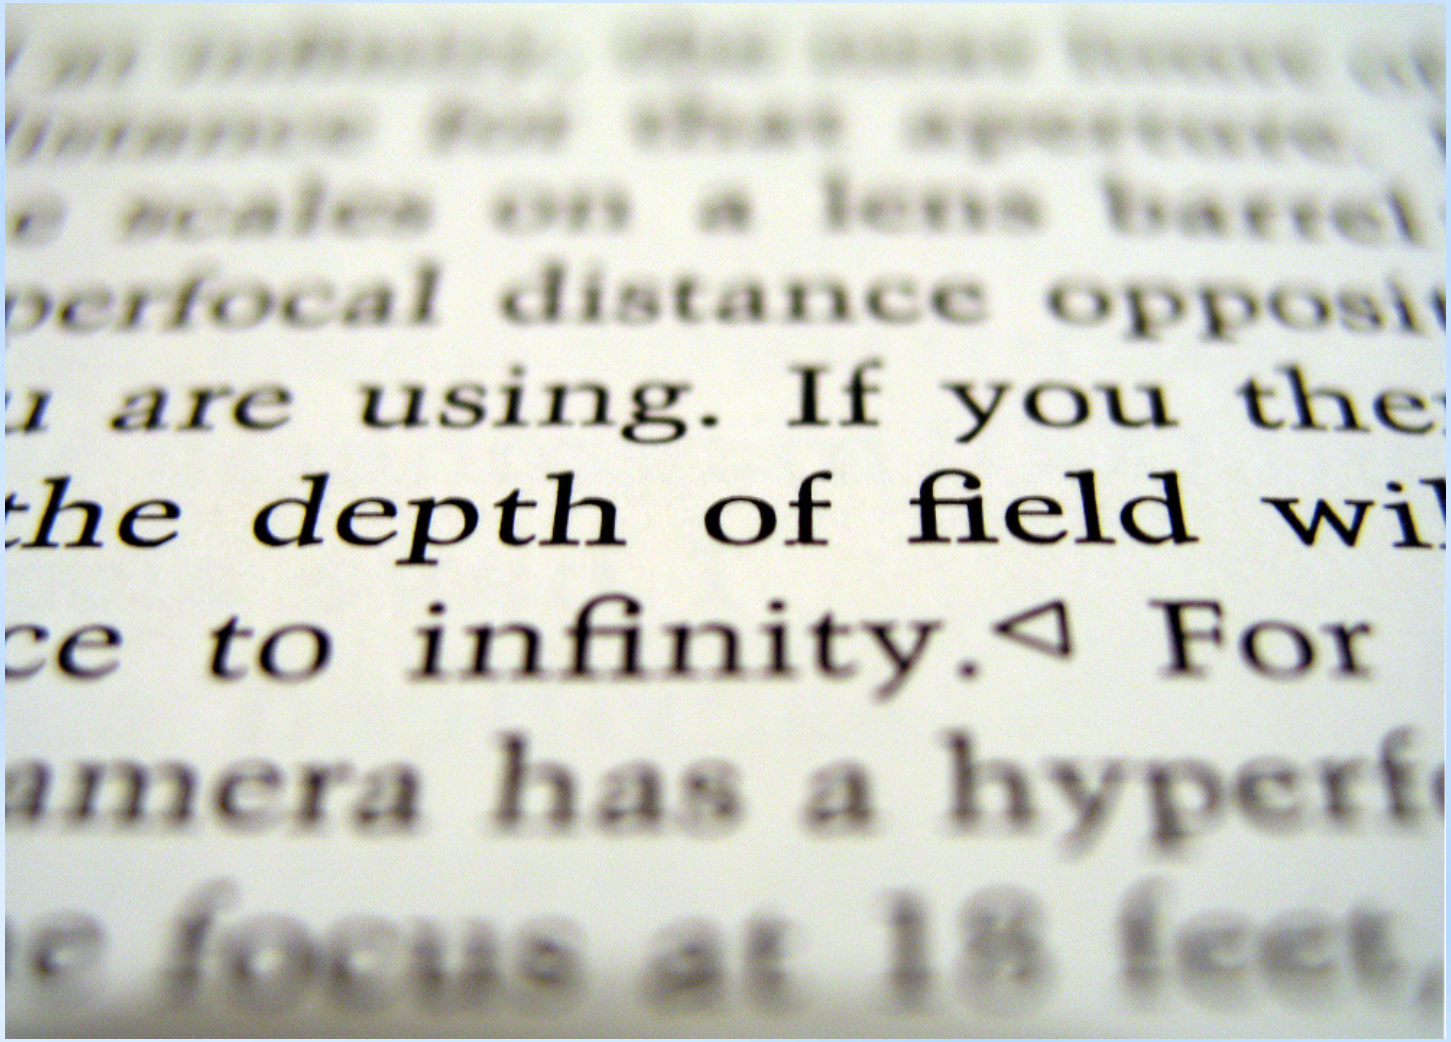

A priori knowledge of the effect of imaging systems, like lenses can be used to infer distance.

Binocular Depth Cues Stereopsis

There are only three parameters and they are:

- ~~Interocular Distance~~
- Interaxial Distance
- Convergence

All depth cues, monocular, binocular and extraocular are integrated to formulate a single, constantly updated, consistent hypothesis and if there is ambiguity among the depth cues, the perception of depth is compromised.

Stereopsis is a binocular system for range finding that relies on the parallax effect produced by two, simultaneous views of the object being ranged that are from different perspectives. The mechanism of the perception of depth in humans relies on the horizontal disparity of the views which is detected as retinal disparity and is synthesized in the brain.

Binocular Depth Cues

Interaxial Distance

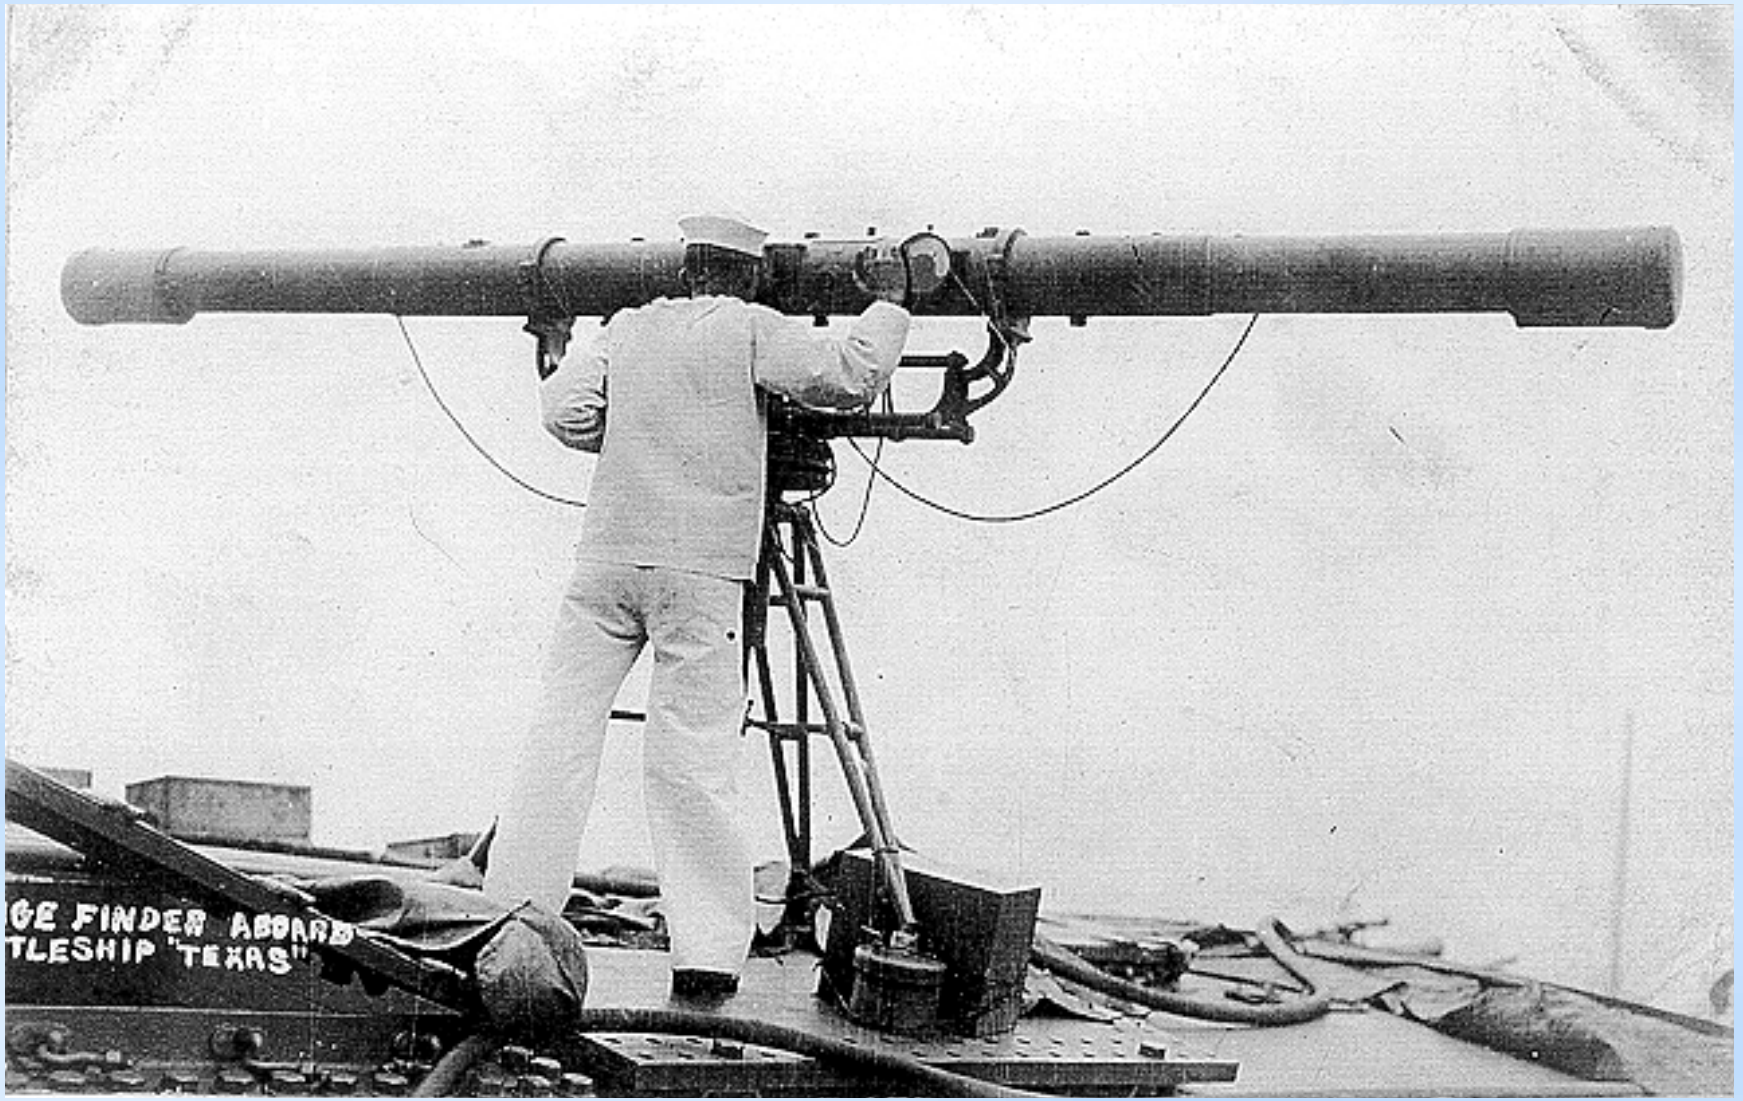

The distance between the prime axis of the stereoscopic imaging system is the interaxial while the interocular is the distance between human eyes. You can change the interaxial, but the interocular distance is immutable.

Binocular Depth Cues
Convergence

The Big Convergence Controversy:

Should we converge our stereo cameras?

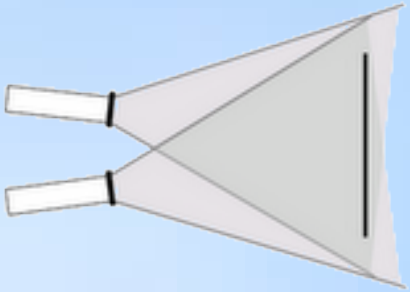

- Guaranteed to keystone
- Potentially impossible to fuse the entire scene simultaneously
- Exacerbates the convergence vs. accommodation conflict
- Can cause eyestrain, headaches and nausea if not done perfectly

Should we shoot (render) everything in parallel?

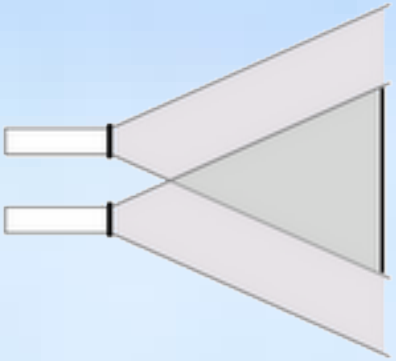

- No keystoning or linear distortions so the chance of eyestrain, headaches and nausea is reduced.
- If the interaxial distance is correct, fusion is possible throughout the entire scene.
- Minimizes the convergence vs. accommodation conflict
- Requires “overshooting” or “overrendering”

Convergence in a pair of stereoscopic images is the point at which there is no visible parallax.

Binocular Depth Cues
Convergence

With pure CG production, there is another option:

Off-Axis rendering provides all of the advantages of parallel rendering and the efficiency of On-Axis rendering without the distortions.

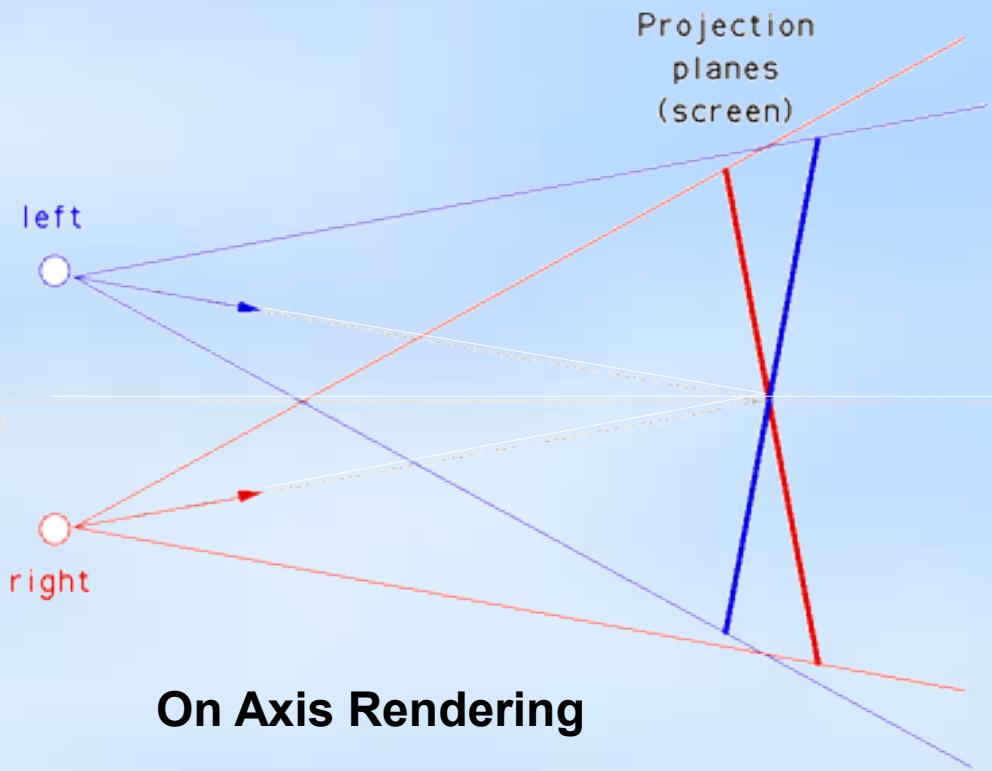

On Axis Rendering

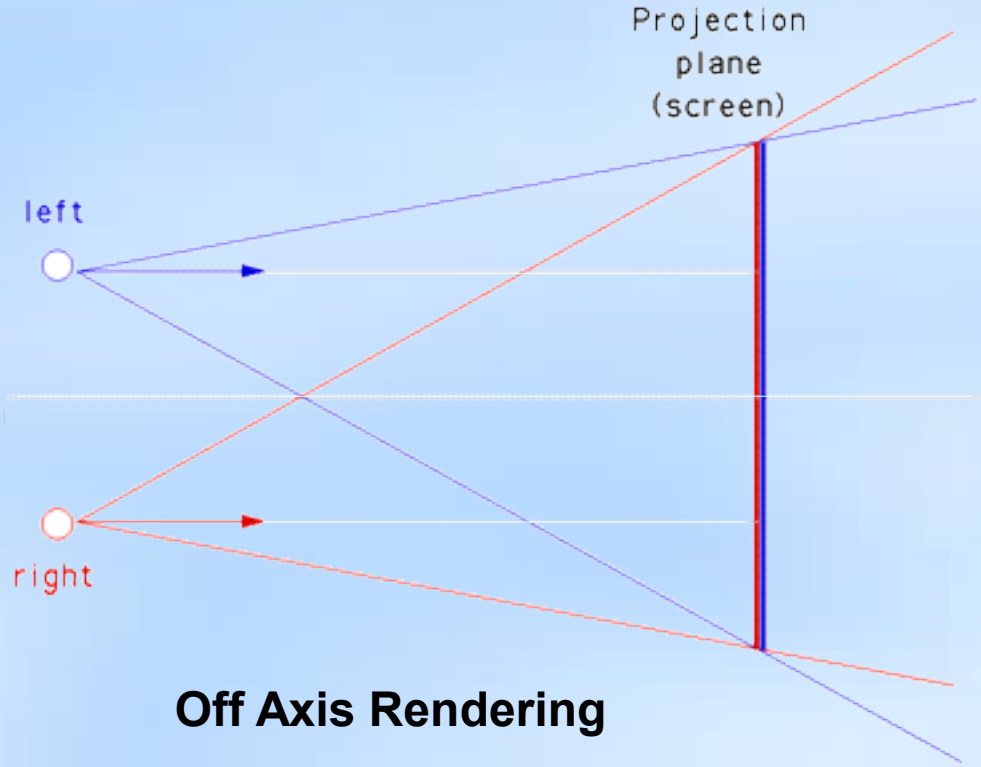

Off Axis Rendering

Binocular Depth Cues Convergence

Fixing the Problem with On Axis Rendering:

Adjusting on axis renders (or photography) is sometimes necessary to reduce eye strain and improve the stereoscopic viewing experience.

It can be done as a 2D process by skewing or corner pinning. The problem with fixing the problem with that solution is that there is significant image degradation.

Occula, a set of Nuke plug-ins make this sort of correction relatively painless, although they still take time and labor.

Occula also corrects color disparity between cameras as well as other alignment and vertical parallax issues.

Binocular Depth Cues Convergence

Reconverging parallel material is trivial:

Adjusting on axis renders (or photography) is sometimes necessary to reduce eye strain and improve the stereoscopic viewing experience.

It can be done as a 2D process by skewing or corner pinning. The problem with fixing the problem with that solution is that there is significant image degradation. (Note: It would be better to get the camera parameters right in the first place.)

Occula, a set of Nuke plug-ins make this sort of correction relatively painless, although they still take time and labor.

Occula also corrects color disparity between cameras as well as other alignment and vertical parallax issues.

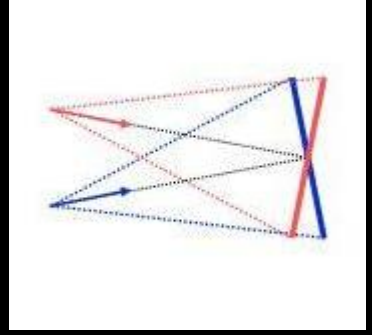

Toe In

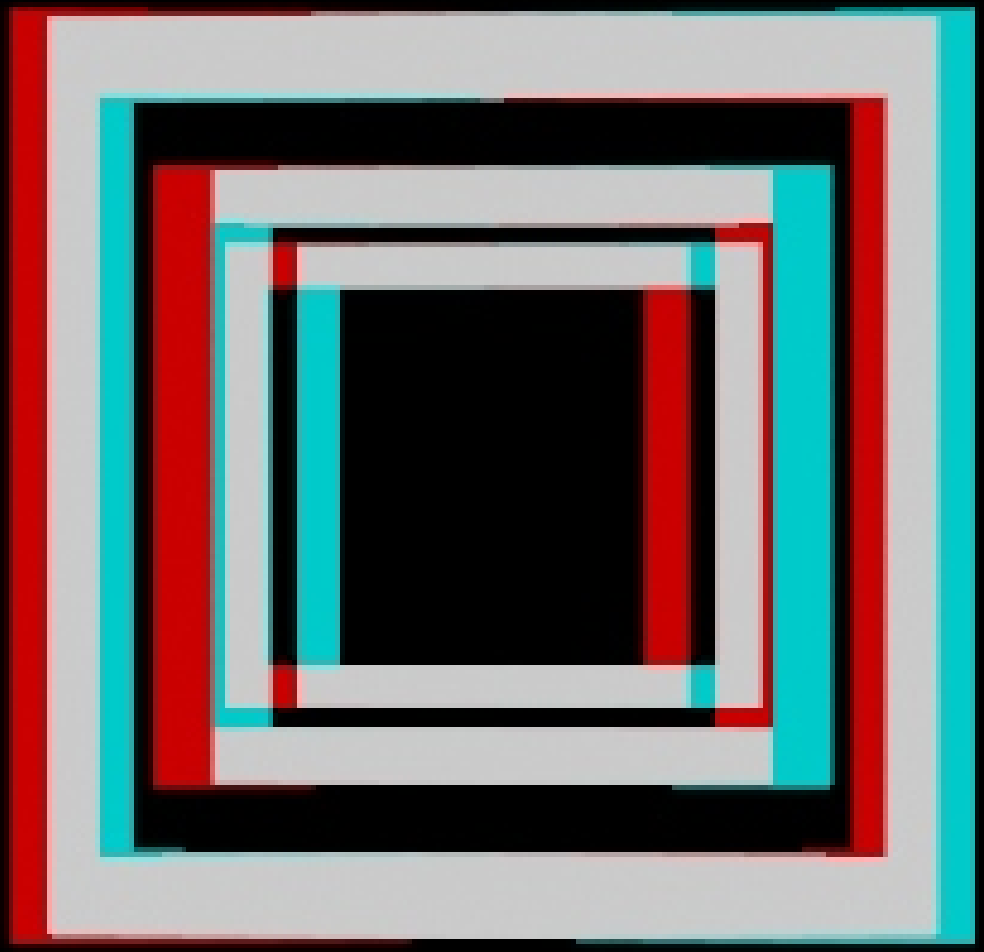

Images rendered by Dan Wexler, Principal Senior Software Engineer at Nvidia.

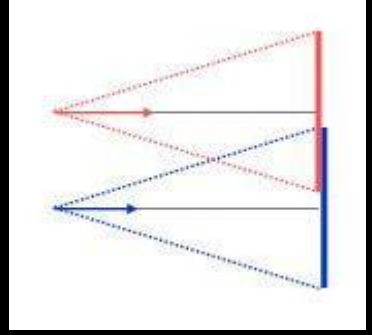

Parallel

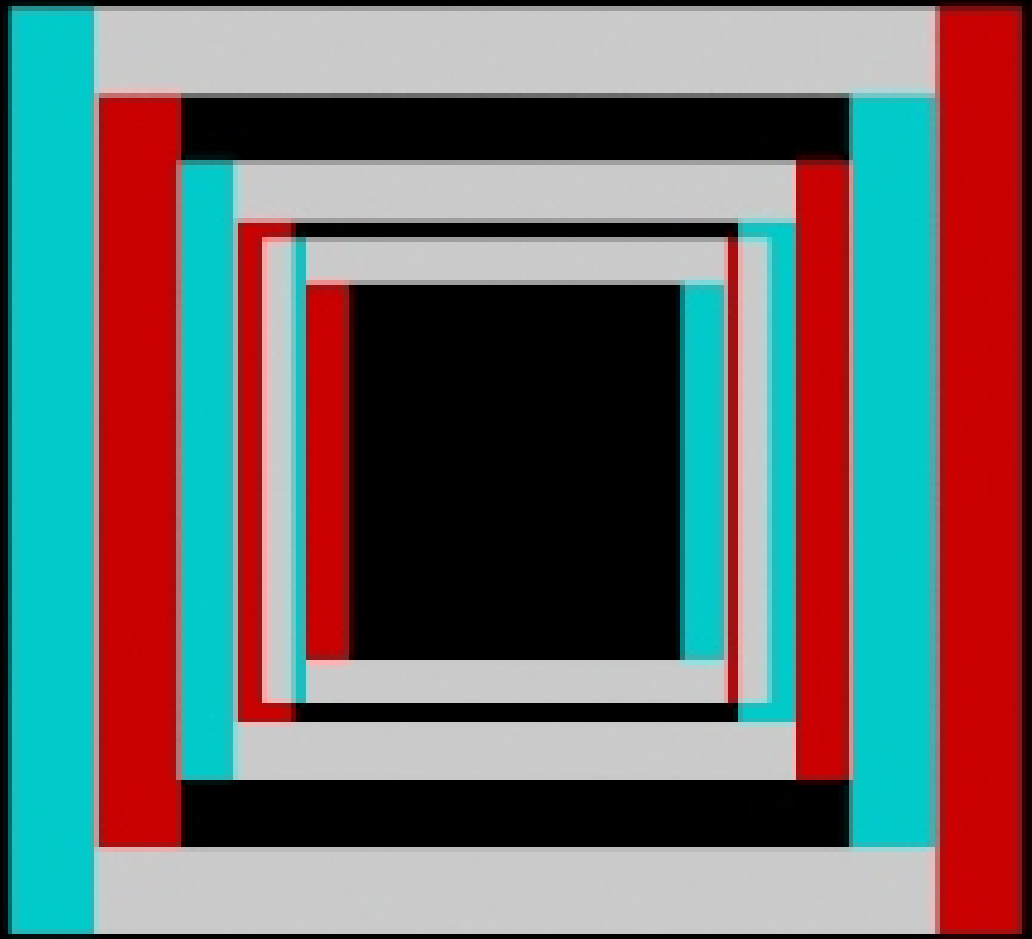

Images rendered by Dan Wexler, Principal Senior Software Engineer at Nvidia.

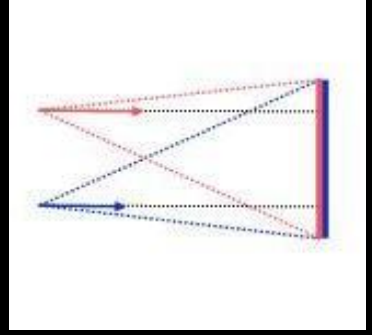

Off Axis

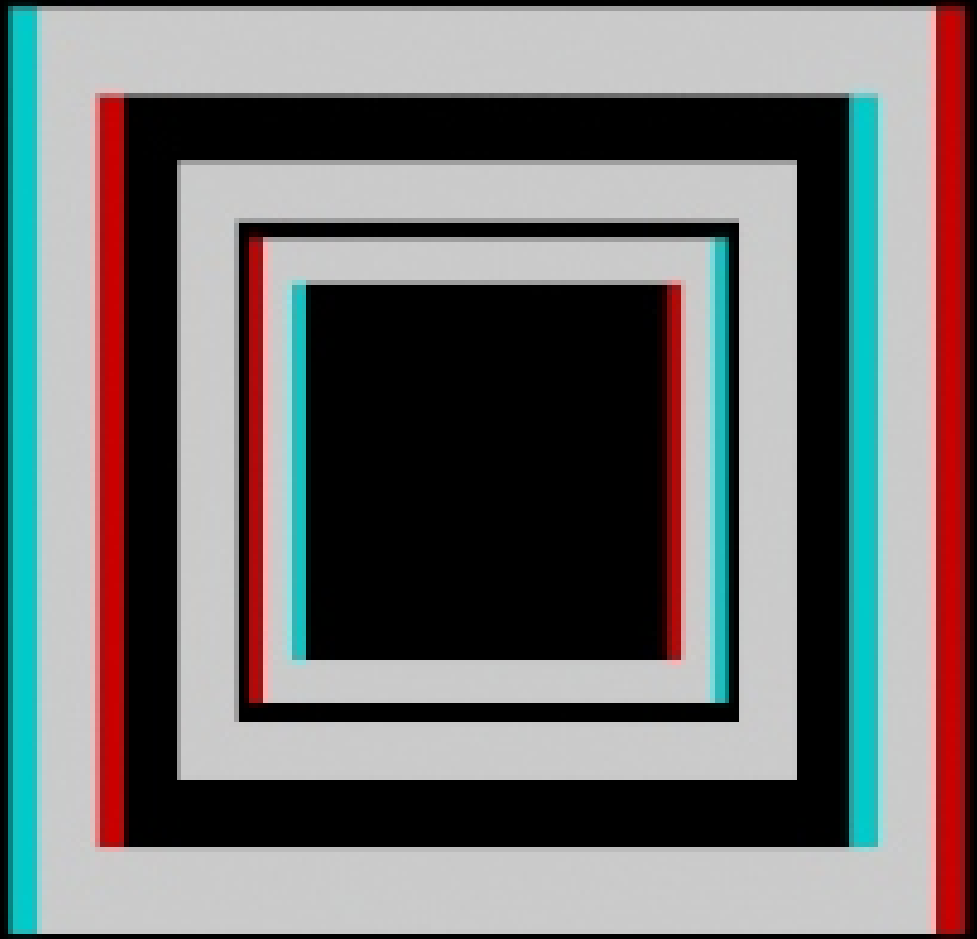

Images rendered by Dan Wexler, Principal Senior Software Engineer at Nvidia.

Binocular Depth Cues
Convergence

Once you decide if to converge, the next question is where?

Common Convergence Mistakes:

- Confusing convergence with the focus of the scene
- Animating convergence within a scene
- Changing convergence randomly during a sequence.

A Convergence Strategy

- Come up with one that works, verify it and stay with it.

  Example: The Window Strategy!

Binocular Depth Cues
Convergence

The Window Strategy

- The screen is a window.
- Most action should take place outside of the window.
- Action can extend from the window inside towards the observer, but be careful of occlusion paradox.
- For a normal interaxial distance (~65mm) and roughly human scale scenes, the window should be roughly 2M in front of the camera. (Note: The 1/30th rule)

The closest regular action in your scene should appear just outside of the perceptual “Window” and the interaxial distance should be 1/30th of the distance between the window and the camera. 1/30th of 2 meters is approximately 65mm. If your scene is different from normal human scale, you should adjust your interaxial distance to match using the 1/30th rule.

More Rules

- The objective isn't to duplicate how the eye works, but rather to create a set of images that the typical human visual system can fuse and integrate into an enjoyable stereoscopic viewing experience.
- You can use depth of field if you want, but with other more natural and real depth cues, it is possible to produce compelling synthetic images where everything in the frame is sharp.
- There should be no vertical parallax. Corresponding points in both single views should always lie on the same horizontal axis.
- The interaxial distance should always be about one thirtieth of the distance from the camera to the nearest object.
- You should avoid rolling the camera. This may cause ambiguous depth cues and degrade the stereoscopic experience.
- Be consistent with your stereoscopic camera parameters.

So, How Do I Make a Stereoscopic Movie?

- Let's assume that you all have a 3D digital production pipeline already functioning in your facility and you want to leverage it.
- You want to add stereoscopic capability to that pipeline in order to produce the highest quality stereoscopic material possible at the least cost.
- Change costs money and so does doing more work, so stereoscopic production is going to cost more.
- What is the most efficient and effective way to add stereoscopic production capability to my existing production process?

So, How Do I Make a Stereoscopic Movie?

- To really answer this question, it would be necessary to survey and analyze your existing process.
- Instead of doing that, we are going to demonstrate one way to integrate stereoscopic production capability into a digital pipeline.
- It is certainly not the only way and it may not be the best way for you.
- Again, the criteria we are using and the principles we employ are relevant, even if the particular solution we are demonstrating may not be appropriate for everyone.

So, How Do I Make a Stereoscopic Movie?

- Who really needs to see stereoscopic data and when do they need to see it?

  Most of the time, most people in the digital production don't need to see stereoscopic material to do their work!

  The stereoscopic Layout and Scene Set-up teams definitely need to see stereoscopic material!

  The director and producer will need to see stereoscopic material periodically.

Conclusion: Provide stereoscopic viewing capability only were absolutely needed, but allow for the capability to see it anywhere in the pipeline!

So, How Do I Make a Stereoscopic Movie?

Your production Process:

So, How Do I Make a Stereoscopic Movie?

Your production Process:

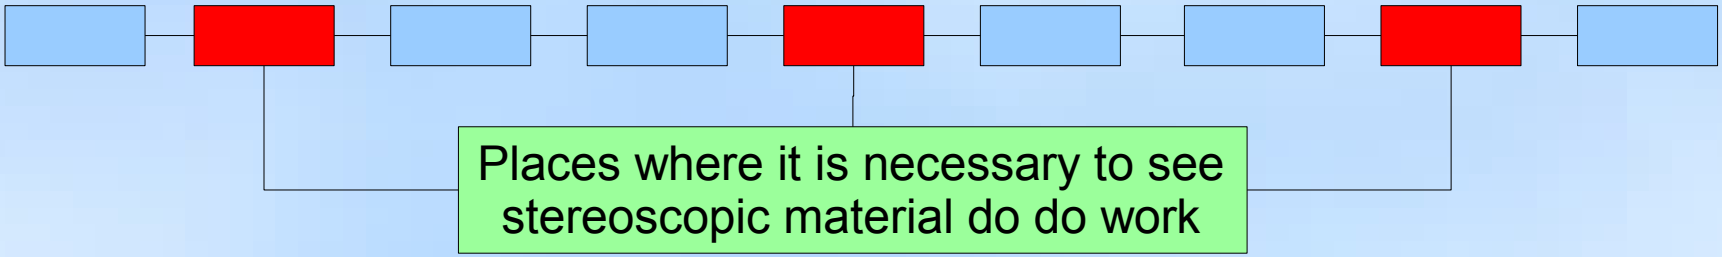

So, How Do I Make a Stereoscopic Movie?

Your production Process:

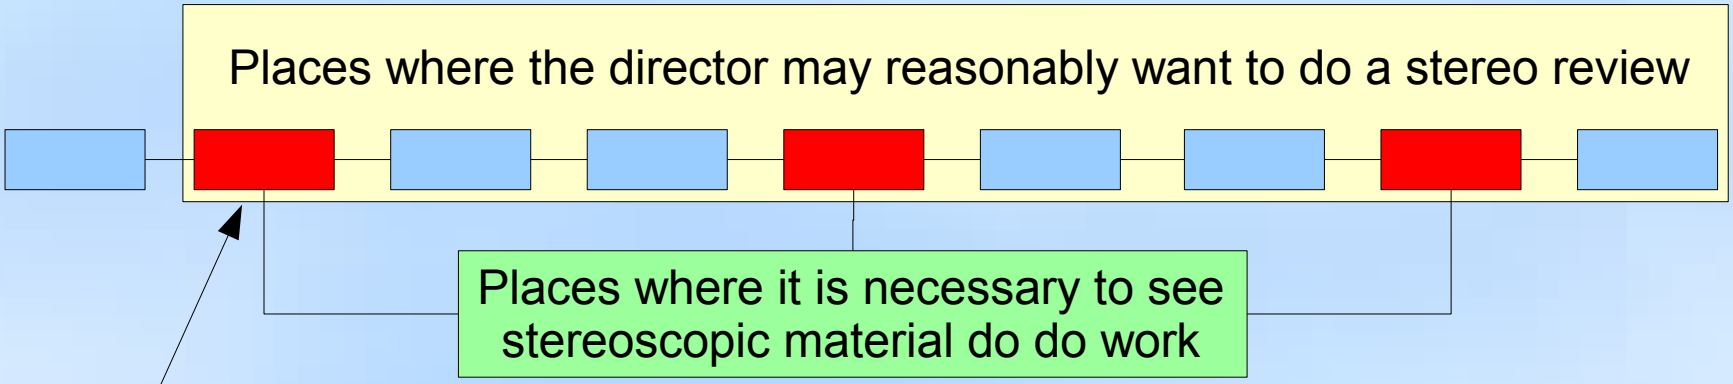

Stereo camera data is defined here and is stored as a separate, shot based asset.

Stereo camera data is not stored with the scene, but as a separate shot based asset. That asset can be used in a render script to generate a second camera at render time so stereo review is possible at any time after the stereo camera is defined.

The process of sending a shot to stereo review can automatically generate the stereo camera!

- This limits access to stereo camera to people with the proper responsibility and training.
- This minimizes the amount of additional viewing equipment necessary.
- This minimizes stereo camera errors that may result from an inexperienced team.

Wake Up!
I'm Done!

Any Questions?