Building a Stereo Pipeline From the Ground Up:

A Comprehensive Study of Disney's

The Secret of the Wings

Siggraph Asia 2012

Vladimir Sierra

Jason Carter, John Park

DisneyToon Studios

Abstract

The Secret of the Wings is the fourth installment of Disney's successful Disney Fairies franchise, and the first Disney Fairies movie that audiences around the world are able to enjoy as a fully stereoscopic experience. As 3D TV and 3D Cinema continue to enjoy worldwide acceptance, establishing flexible and efficient means for creating 3D content is a growing and essential component of modern movie-making. During this half-day course, we will take a comprehensive look at all aspects of stereoscopic film creation examining lessons learned from Disney's *The Secret of the Wings*. In particular, this course aims to provide a thorough understanding of stereo creation for CG-animated films with an eye towards artistic, technical, and pipeline execution.

Introduction

The primary focus of this talk is to present the topic of stereo creation from the perspective of how it fits into production. Specifically, by the end of this course, we would like our audience to come away with an understanding of what it means to create content in stereo in terms of resources needed, how it a affects production, and hopefully how this new artistic element can be shepherded creatively throughout the making of a film. While no prior experience with stereo is necessary for this talk, exposure to recent stereoscopic titles is certainly helpful. Given that the more technical aspects of stereo design have been covered in previous Siggraph talks, this talk will have very little in the way of theory or delivery mechanisms. We want to present a more practical approach to stereo with most of the emphasis on the artistic and production aspects of it. As a case study, we will describe the work-flow we implemented during production of *The Secret of the Wings.* We feel that providing insight into the process we went through to arrive at our stereo choices for the film has educational value that hopefully will be of use in a wide variety of stereo endeavors. By illustrating the progression of stereo sequences from depth script all the way through final color, our primary goal is to give a complete picture of how stereo can organically evolve alongside the changing needs of a film production.

Course Outline

Excerpt from The Secret of the Wings (10 minutes)

Introduction (10 minutes)

Stereography—The Art of Making Stereo (30 minutes)

- Stereo and Story
- Depth Script Design
- Working with Pixel Offsets

Stereo Workflow Part I: *Pipeline Design at DTS* (30 minutes)

- Pre-Stereo Digital Pipeline at DisneyToon Studios
- Stereo Pipeline for *The Secret of the Wings*

Stereo Workflow Part II: *Workflow Principles* (30 minutes)

- Depth Script
- Stereo Layout
- Stereo Animation
- Stereo Color
- Depth Grading

Intermission (15 minutes)

Overview of Stereo Concepts (20 minutes)

- Retinal Image Size
- Perspective
- Interposition or Overlapping
- Aerial Perspective
- Light and Shade
- Textural Gradient

Motion Parallax

Accommodation

Convergence and Disparity

Stereopsis

Stereo Techniques in CG: Creating Stereo Pairs (20 minutes)

Two-Camera Rendering

Z-Depth Compositing

Stereo Conversion

Compositing via Pixel Offsets

Depth Grading

Stereo Techniques: Design Principles (20 minutes)

Depth Budget

Creating Layered Depth

Layered Depth and Multi-Rigging

Stereo and World Scale

Camera Staging

Stereo Techniques: Advanced Concepts (30 minutes)

Direct/Indirect Stereo

Stereo Windows

Production Lessons Learned (15 minutes)

Lessons from *The Secret of the Wings*

Closing Thoughts

Q/A (15 minutes)

Stereography: The Art of Making Stereo

Stereo and Story

Stereoscopic films have experienced a boom over the course of the last few years in terms of variety of titles being produced for 3D, and the availability of venues in which to view rich 3D content. Whether stereo is here to stay, however, is very much an open question, and the answer may be directly tied to how organically film-makers integrate stereo as part of their visual language. In the same fashion that the addition of sound and color expanded the story-telling possibilities available to film-makers, the challenge that will determine the longevity of stereo as a visual tool is in how organically it can be made a part of the story-telling fabric. It is our belief that, in order for stereo to be truly embraced as an art form, it must become as integral part of the film-making process as possible. In describing the process of incorporating stereo as an integral part of *The Secret of the Wings*, we hope to contribute to the evolution of stereo and stereography at large.

How can stereo be used to help tell a story? The answer is, of course, very much a matter of interpretation. In the same manner that, say, lighting can be used to drive the emotional tone of a story, so too stereo can be tied to specific aspects of a film to reinforce a particular mood or character development. Similarly, where stereo takes place relative to the audience can be used to heighten some pivotal aspect of story. The possibilities are endless. Every film is as different as the film-makers that create them. In designing the stereo experience for *Avatar*, James Cameron made stereo choices that consistently kept the focus of the action right at screen space for the entire film. This illustrates a point that is important to establish early on. While the options for stereo execution are limitless, whatever stereo decisions are made for a film should be consistent, and should have an awareness of the language that has been established for the telling of it. At minimum, there are three aspects of film language that should be considered during the planning stages for stereo:

- Pacing: The pace in which shots flow throughout a film
- Staging: The language of cinema: close-ups, medium shots, open shots, etc
- Story arc: The thousand and one ways in which film moments are designed to affect an audience.

Interestingly, because of the added dimension that stereo brings into the mix, the geometry of how a scene appears on screen stretches from a two-dimensional rectangle to a three-dimensional cone that coincides with the camera viewing frustum.

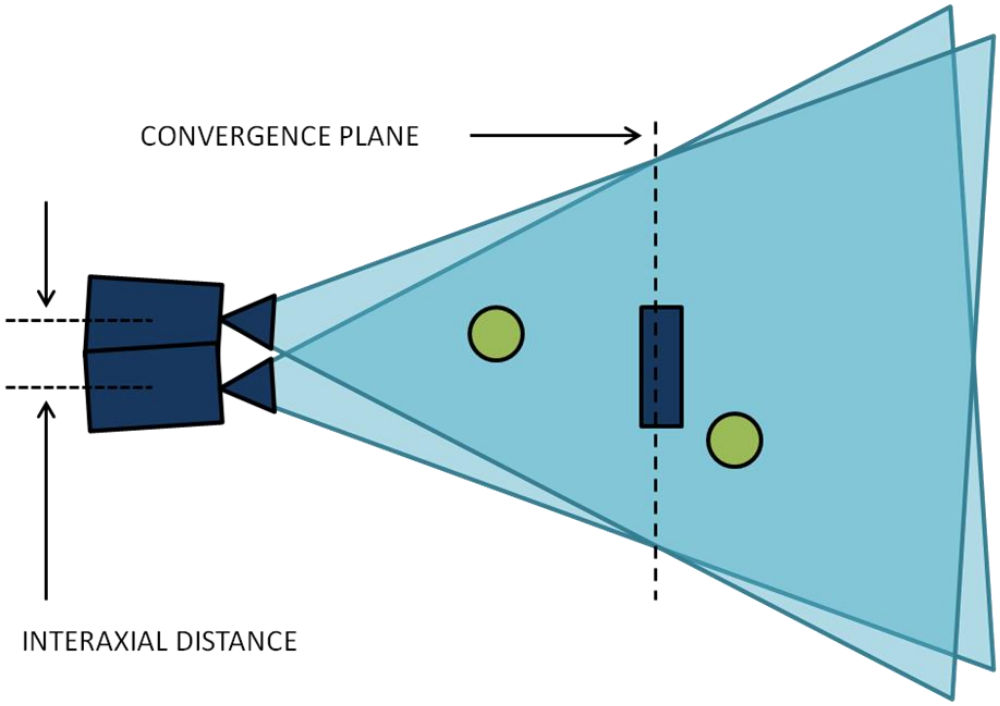

Figure 1: *A basic two-camera stereo set-up.* © Disney Enterprises, Inc.

This graphic is instrumental in understanding stereo design for two reasons. One, it demonstrates the concept of stereo as how visual imagery is distributed inside this viewing cone. Secondly, it illustrates that a stereographic experience, whether at home or in a theater, actually is distributed in three regions: in front of the viewing screen, (audience-space), in the vicinity of the presentation screen, (convergence plane), and lastly behind the screen (behind-screen-space). Thus, when we look at these three regions together, stereo design boils down to the way in which an artist chooses to distribute the contents of a scene inside this viewing cone. Although there are a multitude of stereo rig designs for both software and hardware needs, changing the parameters of this stereo cone always depends on the same two parameters:

- Interaxial Distance: The distance between the cameras
- Convergence Plane: The distance to the plane where LE and RE coincide

So, now that we have a basic understanding of what stereo means in terms of camera placement on a set (digital or live), and of how we can control the appearance of stereo relative

to a viewing screen, we can take a step back, and begin to think about how we can go about designing a stereo experience for an entire film. If we have little to no experience with stereo, a good place to start is by watching a variety of stereo titles across different genres. This experience will little by little shape our own opinions as to what kinds of stereo experiences most nearly reflect the sensibilities of the films we want to make. From this process, we can begin thinking about our own film in terms of stereo. We don't need fully concrete ideas at this stage necessarily, but at the very least, we can ask ourselves where in our film we would want to maximize the stereo experience, and conversely, what moments we would opt for a more subtle treatment. Ideally, stereo design should begin to take shape as early as storyboarding. Certainly if the production affords time for a camera script, the stereo department should be aware of what kinds of lensing decisions the director wants to make. For instance, if a film makes use of long lenses, and has a very fast-paced tempo, that is going to lead to a very different stereo treatment than for a film that has a calm cinematic quality to it.

Depth Script Design

One of the most interesting things about 3D design, and 3D creation for that matter, is the fact that the tools that we have available to create stereo content make its execution very easy to master. This is particularly the case when we compare stereo to other aspects of digital film-making such as modeling, lighting, and even compositing. Compared to those disciplines, which can take years to master, a new stereo artists can pick up the notion of stereo creation fairly quickly and be able to design shots in a short amount of time. However, in the same manner as the other disciplines mentioned, mastering the art of stereo is something that takes developing a good eye, and years of training. The reason why the early efforts in the resurgence of stereoscopic films we've experienced have delivered mixed results is that their use of stereo was not consistently disciplined or complementary of the stories they were trying to tell. We are only now beginning to find out that stereoscopic imagery, when used as a gimmick to be taken advantage of, actually does a disservice to filmic story-telling. Fortunately, along with this realization, we've come to understand that stereo needs a vision, and this vision is best executed through a depth script. Similar to the color script, the depth script is the very first step toward incorporating stereo design as a more organic part of the story-telling process.

However, as essential as this role is, the importance of the depth script goes beyond helping to create memorable film experiences. Given that the art of film-making is inherently a collaborative process, its first and most important role is that of communicating the stereographer's intended contribution to a team of artists and producers. In this sense,

presenting a thorough, well-thought-out explanation of the depth script is nearly as important as the depth script itself. Additionally, the presentation of the depth script may be the one opportunity a stereographer has to communicate how a project may resonate on a personal level, and establish a working dynamic that may last throughout the length of a project.

During the planning stages of *The Secret of The Wings*, stereographer Vladimir Sierra had the dual task of exposing director Peggy Holmes to the language of stereo, while at the same time implementing a vision of the film that would be complementary to her style of story-telling. Creating a depth script that would educate while at the same time communicate a deep familiarity with the film became paramount. It was in this manner that, using his background in data visualization for the aerospace industry, stereographer Vladimir Sierra set out to create a depth script presentation that, in addition to indicating the ebb and flow of depth throughout the film, would provide a 3D representation of what those depth choices would look like at any point in the story. The resulting visualization was a kind of 3D room containing a depth graph along with simple iconography that could be moved in space according to how much depth budget corresponding to each segment of the film. With additional design input from CG Supervisor John Park, this virtual room was designed in such a way that, when viewed in stereo, one could see a clear representation of the complete range of stereo available, and what portion of that budget was allocated at any given point. This budget was represented in terms of depth percentages.

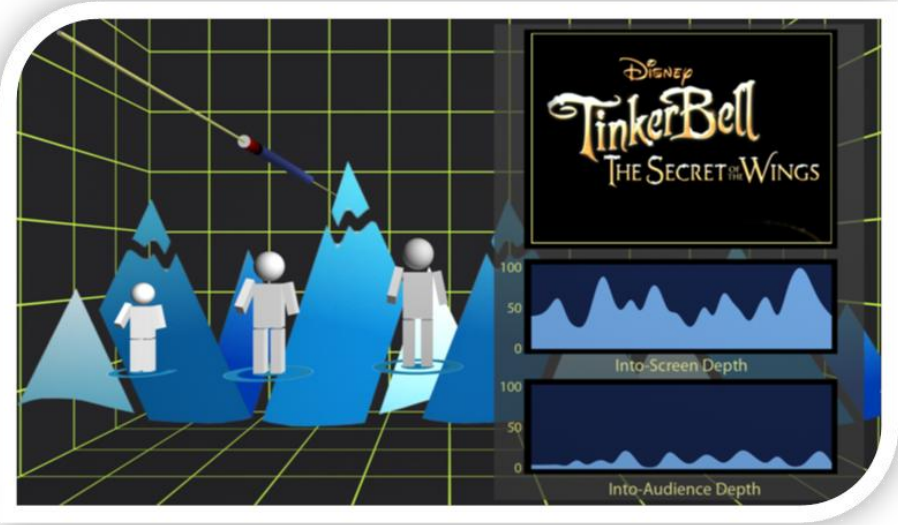

Figure2: *A normal depth moment from* ***The Secret of the Wings.*** © Disney Enterprises, Inc.

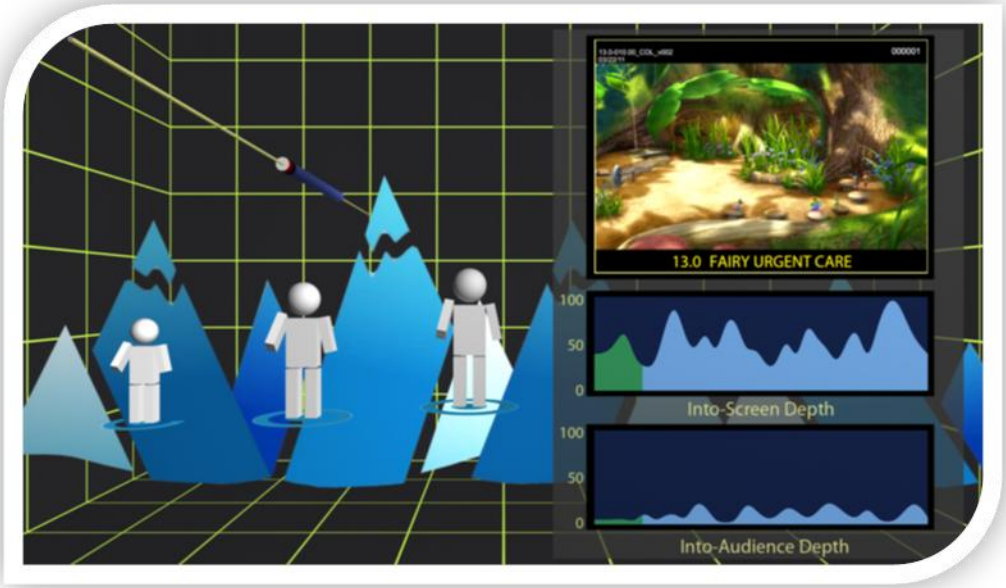

Figure 3: *A shallow depth moment from* ***The Secret of the Wings.*** © Disney Enterprises, Inc.

Figure 4: *An aggressive depth moment from* ***The Secret of the Wings.*** © Disney Enterprises, Inc.

One clear advantage that we gained from being able to present a depth script actually in 3D was that there was no question about the intention that Sierra had set forth for the film. The positioning of the graph, as well as the representation of depth relative to elements in the foreground, midground and background was particularly useful, and became the basis of a dialogue that would continue between film-maker and stereographer throughout the length of the project. In fact, it was due in large part to this meticulous breakdown of spatial allocation that allowed our director to vocalize a preference for a stereo experience that would put the majority of the action behind the screen. While on the one hand, this represented a slight departure from Sierra's initial vision for his proposed stereo treatment, it ultimately led director Peggy Holmes to embrace stereo as a valuable addition to her film.

The format of a depth script does not necessarily need to be as elaborately executed as the one we created for *The Secret of the Wings*. As outlined above, its design was meant to educate, as well as to provide assurance of the stereographer's understanding of the film. In this sense, the format of the depth script that Vladimir created for *The Secret of the Wings* may turn out to be a one-of-a-kind visual document. Other projects may dictate more or less information, or a different breakdown of stereo parameters (such as an overall depth value vs. one for positive and negative spaces.) In shaping its design, however, it's generally a good idea to analyze the film in its entirety in terms of important film elements that might drive depth.

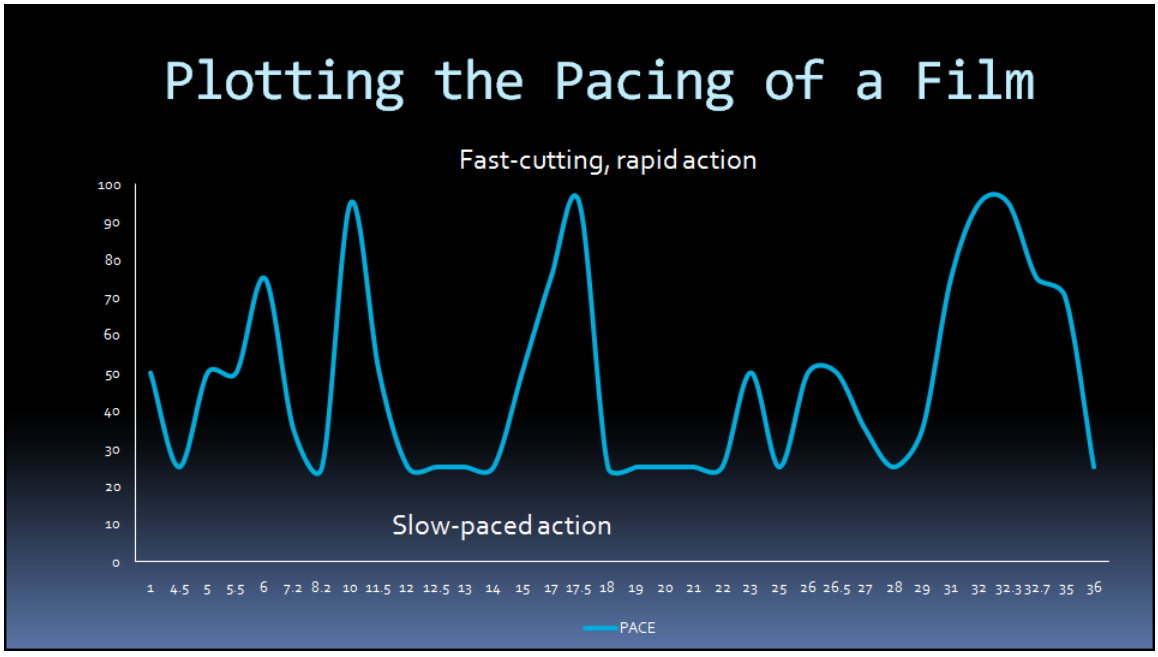

Figure 5: *Plotting the pace of the film, either in terms of how quickly the film flows from shot to shot, or in terms of the dramatic momentum, is a possible beginning for plotting depth.* © Disney Enterprises, Inc.

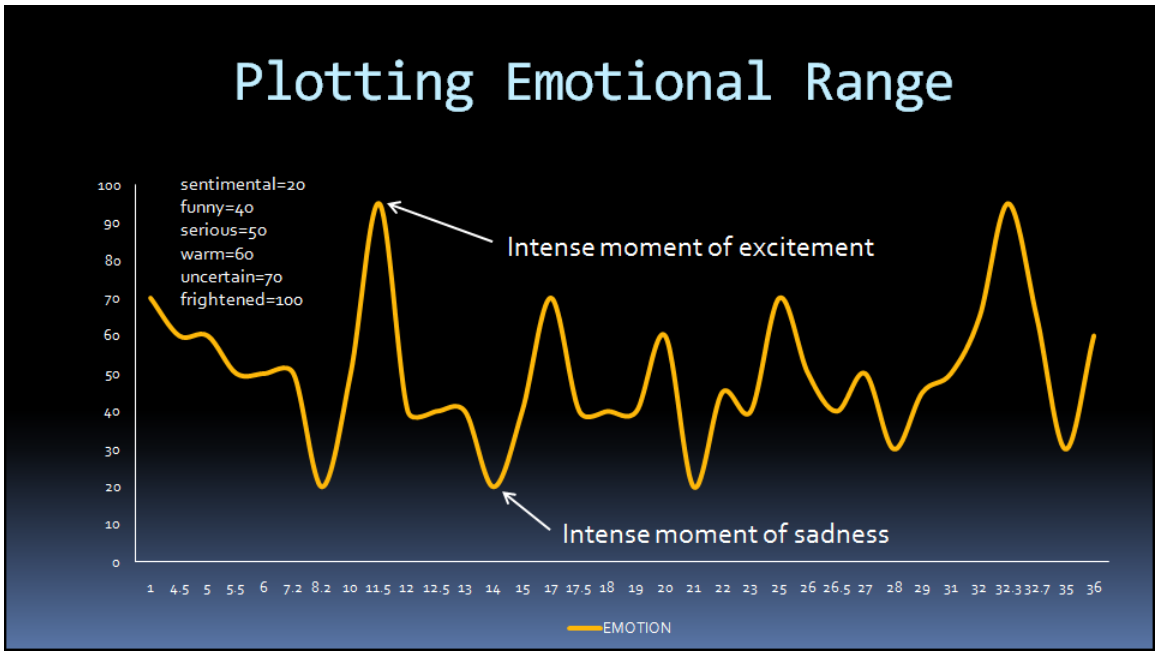

Figure 6: *Plotting emotional pace may yield a different picture that could also inform how stereo depth needs to flow across a film. In this example, excitement is mapped to aggressive depth choices, and sentimental moments are mapped to shallow depth choices.* © Disney Enterprises, Inc.

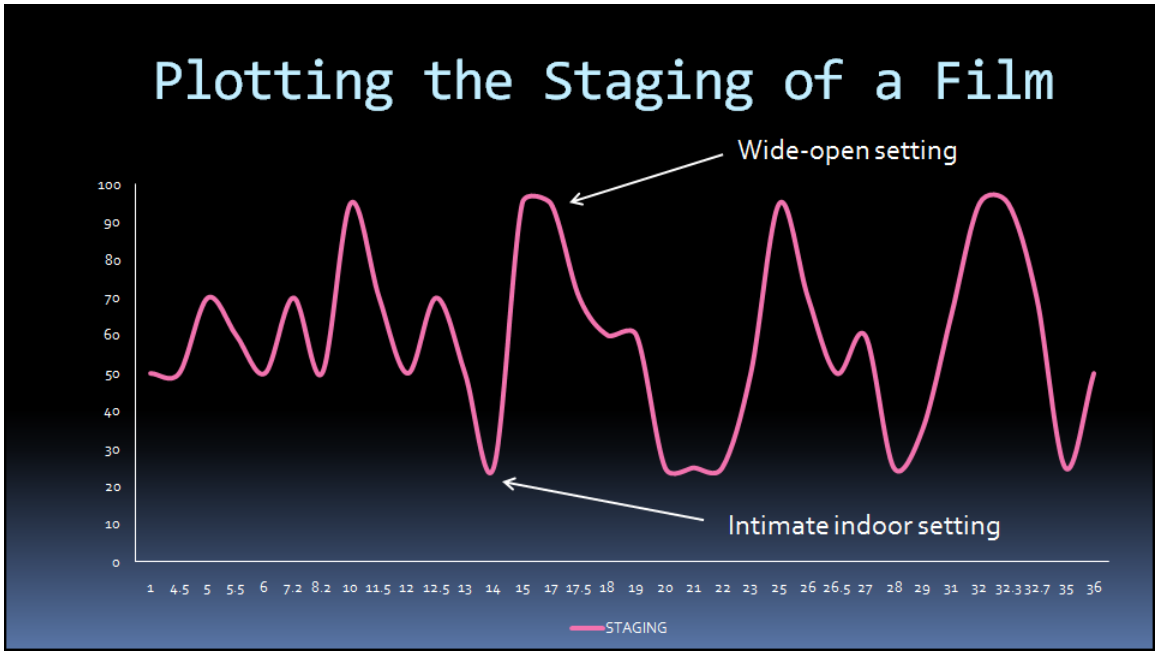

Figure 7: *Where the action takes place in a film could also factor strongly in determining depth choices. Here, cinematic moments get a generous depth treatment, whereas indoor, intimate settings get more conservative values.* © Disney Enterprises, Inc.

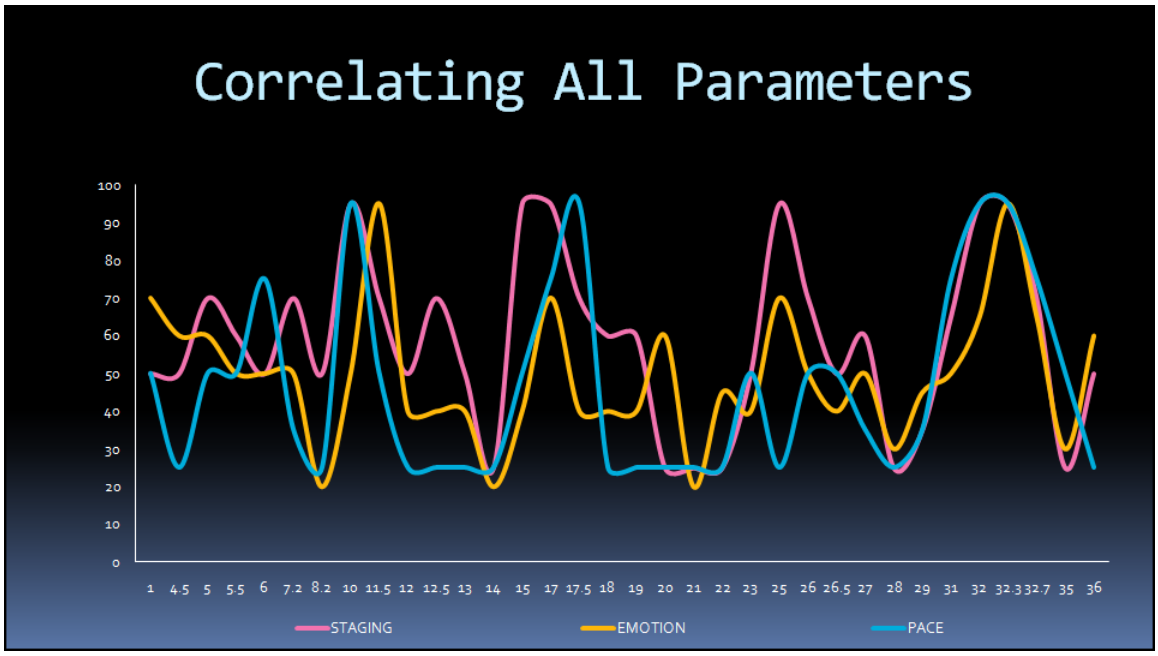

Figure 8: *Looking at all relevant parameters might help to solidify creative instincts, but in areas where there is no obvious correlation, it will necessitate further analysis and problem-solving.* © Disney Enterprises, Inc.

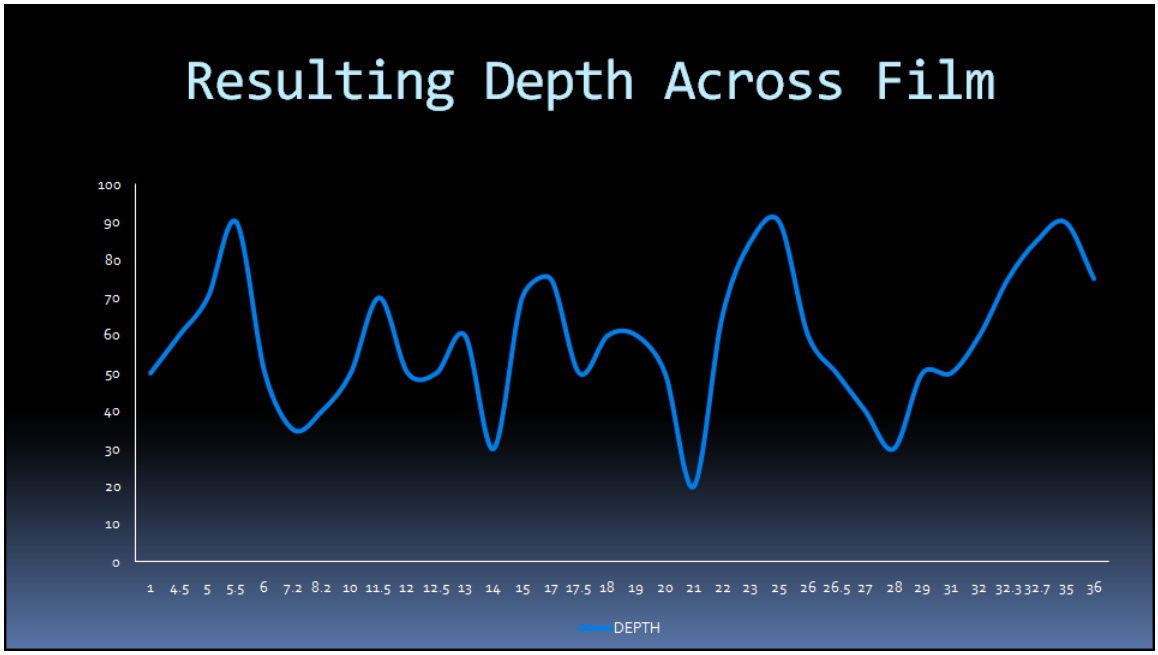

Figure 9: *In the end, the resulting depth script will come together from a set of rules the stereographer has established, and from creative deviations that justify overriding these rules.* © Disney Enterprises, Inc.

For a recent project, Sierra experimented with visualizing the flow of various film parameters such as pacing, staging and emotional content as an animated landscape. The resulting visualization provides an interesting representation of a film across time, and in future projects may provide a valuable mechanism for designing a more organic marriage between stereo and story.

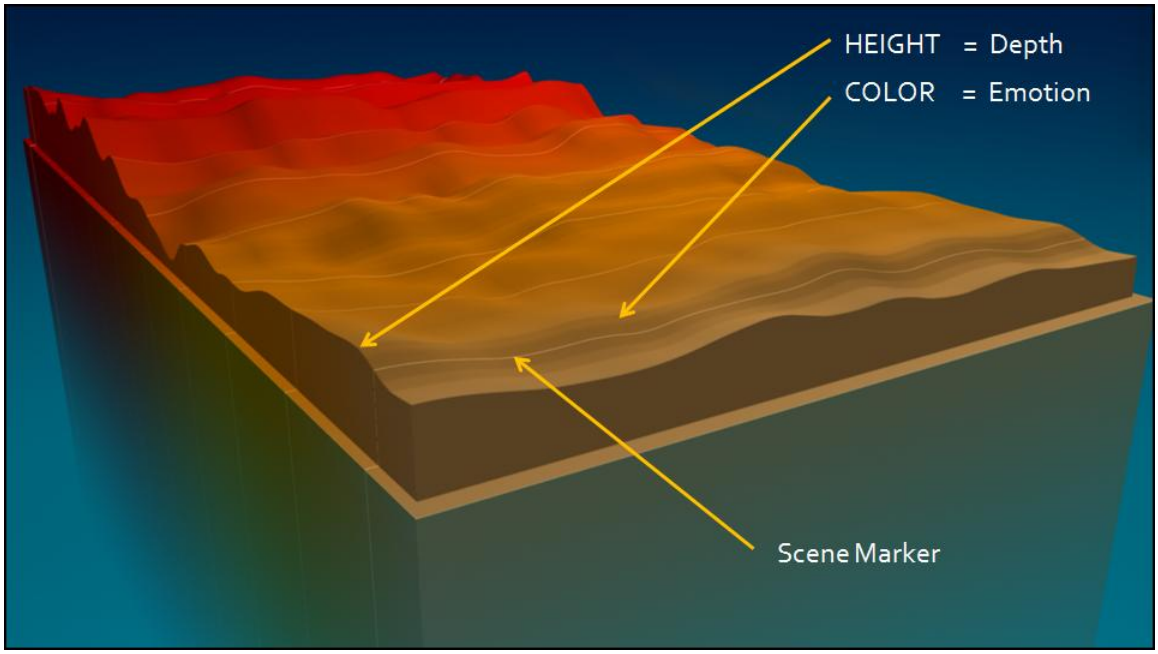

Figure 10: *A still from an animated depth script concept by Vladimir Sierra.*

Finally, regardless of how a depth script is carried out, it is important to use a visual device that will clearly convey not only an overall vision for depth in a film, but also the motivation behind that vision. In the end, a clear, concise exposition of a stereographer's understanding of the film's soul is the first step toward a dialogue that will continue to evolve throughout the length of a project.

Working with Pixel Offsets

Thus far, we've presented the idea that depth is a quantifiable visual instrument that can be dialed to strengthen the emotional arc of a film. Given that the images we project on-screen have standard pixel resolutions, it turns out that we can also express depth in terms of pixels, as opposed to depth percentages. The reason for this is very simple: Since the stereo effect in an image pair results from the fact that each image is shot from a slightly different perspective, the horizontal offset that separates each image correlates directly to the amount of depth in a scene.

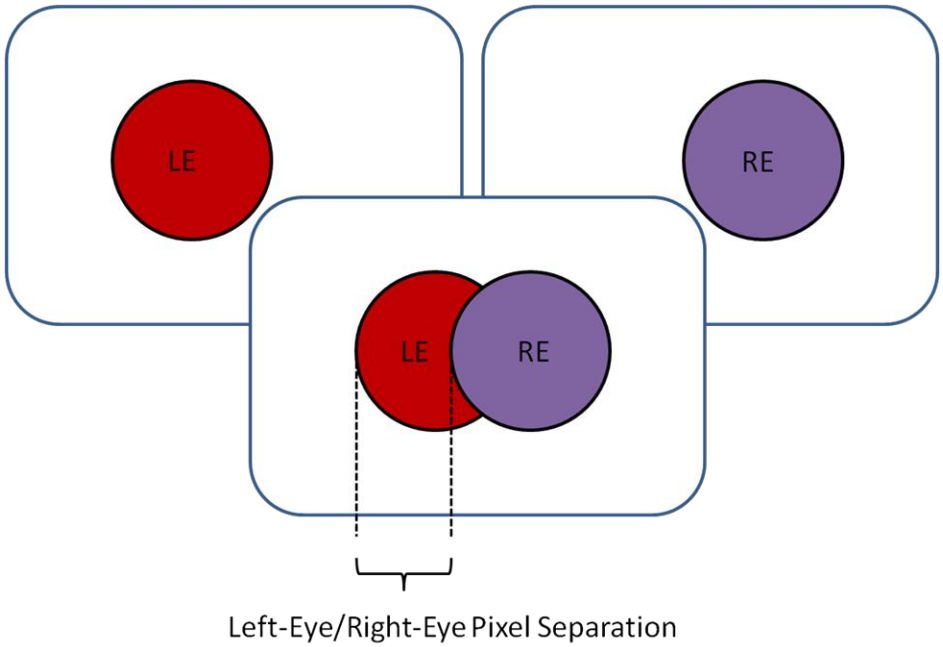

Figure 11: *When a left-eye/right-eye stereo pair is seen simultaneously, elements within the image will be offset from each other by some pixel distance.* © Disney Enterprises, Inc.

In fact, the manner in which this pixel offset manifests at any point in each image pair determines not only depth, but also where it will be manifested relative to a viewing screen. Image elements that line up perfectly (i.e. no perceptible offset) will project exactly **at screen depth**. Image elements in which the right-eye image is offset to the right of the left-eye will appear **behind the screen**; lastly, image elements in which this offset is reversed will appear **in front of the screen**. Thus, when put together, a stereo pair that features all three depth types manifests a rich depth experience. Not only is this the case, but the degree of richness in depth is directly proportional to the numeric value of these pixel offsets.

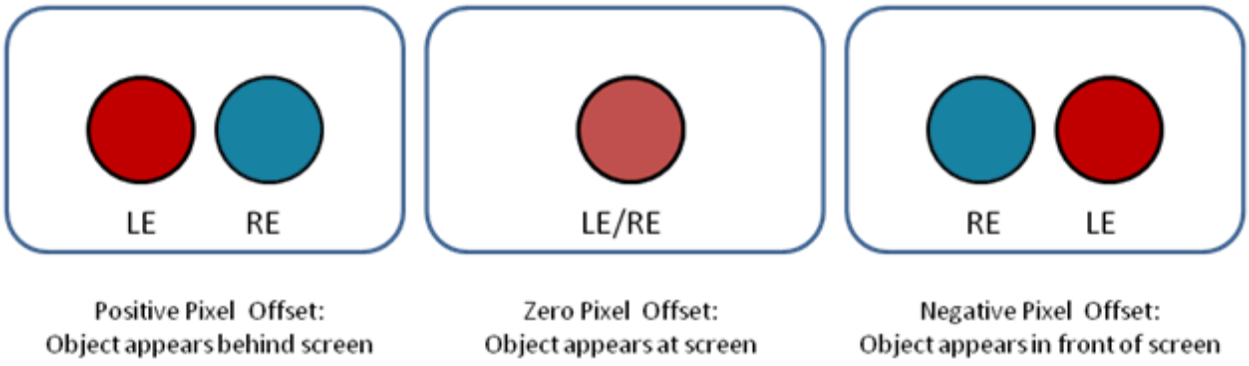

Figure 12: *The degree and direction of the pixel offset in a stereo pair determines how the depth is perceived by the human eye.* © Disney Enterprises, Inc.

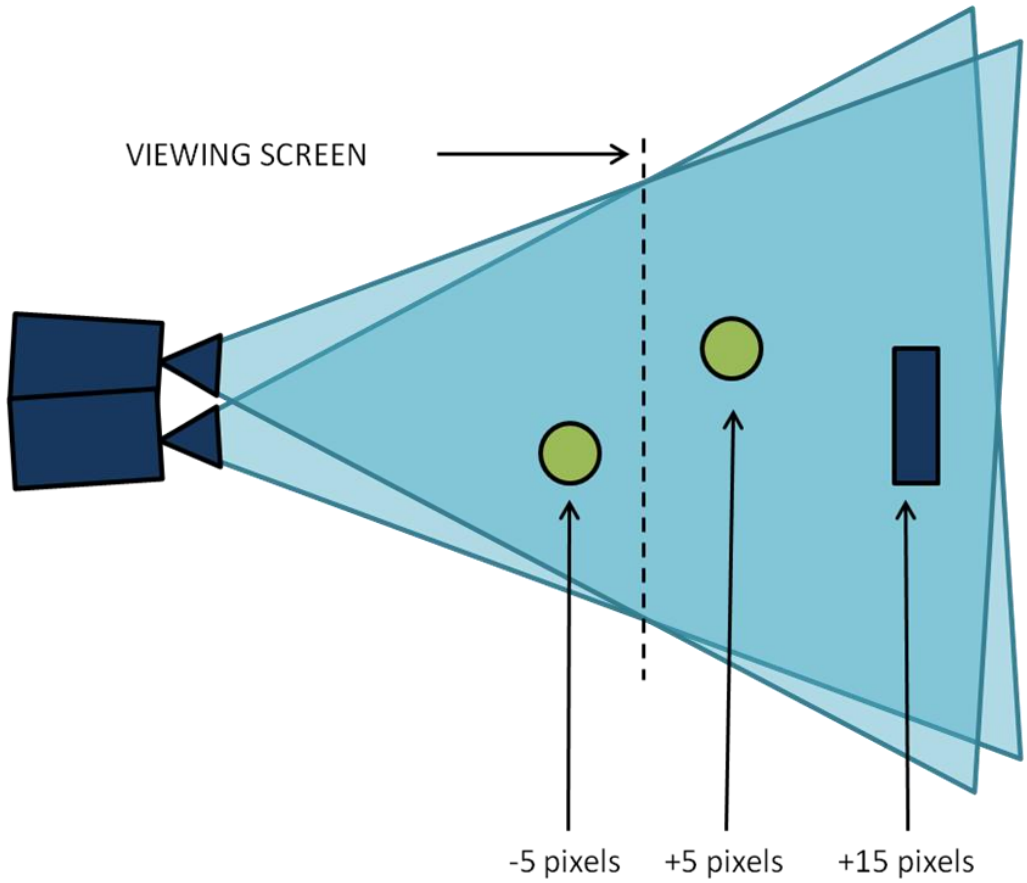

Figure 13: *Objects with positive pixel offsets appear behind the screen. By contrast, objects with negative pixel offsets get pushed in front of the screen.* © Disney Enterprises, Inc.

In a manner of speaking, from a technical standpoint stereo design can be boiled down to an intelligent application of these pixel offsets to elements within a stereo image pair. This very concept is essentially the backbone that drives the stereo conversion of 2D titles into 3D. Using a combination of math algorithms, painstaking rotoscoping of elements, and lots and lots of human craftsmanship, a whole industry has developed that specializes in creating plausibly stereoscopic versions of monoscopic films. Outside 2D to 3D conversion applications, however, (and most relevant to this course,) CG-animated films, and even live-action 3D films arrive at these offsets by rendering images out of two cameras that are offset from each other by an *interaxial distance*. The resulting left-eye/right-eye stereo pairs have within them the rich range of depth that we have become accustomed to in modern 3D cinema.

Regardless of the context, whether conversion projects, animated 3D projects or live-action 3D, a proper understanding of pixel offsets is extremely important in deciding how aggressively or conservatively we want to push stereo design. In fact, it can be said that every studio that engages in the creation of stereo content has, either through experience, research, or a mixture of the two, arrived at their own pixel charts that determine the acceptable ranges for their stereoscopic releases. In this regard, DisneyToon Studios is no exception. Below is a summary of general stereoscopic guidelines used within the Studio:

Depth Strength (pixels)	Negative Space (in front of screen)	Positive Space (behind screen)
Weak	0-8	5-15
Normal	8-15	15-30
Strong	15-25	30-40

Figure 14: *DisneyToon StudiospPixel offset targets for HD stereo content.* © Disney Enterprises, Inc.

It should be noted that the amount of perceived depth is not only a function of the pixel offsets manifested in an image pair, but also of the size of the viewing screen where the images will be projected. As a general rule, a film that modulates its depth within the ranges presented above will result in a comfortable viewing experience on screens ranging from 46” to medium-sized movie screens. These values take into account a wide range of ages in a viewing audience such that, from a comfort viewing perspective, they can be enjoyed by children and parents. Michel Benoit’s excellent paper on 3D Content Targeting provides excellent insight into stereo delivery methods. (Benoit, pp1-4)

Stereo Workflow Part I: *Pipeline Design at DTS*

Pre-Stereo Digital Pipeline at DisneyToon Studios

We would now like to present an overview of the digital pipeline at DisneyToon Studios from the perspective of how shots move through various stages to end up as beautiful finished frames. For the purposes of this talk, we begin this discussion at the point in which story is locked, the movie is entirely boarded, and all systems are go. From here, once a sequence of shots is approved by the director in editorial and is ready to move on to digital production, it gets distributed to our Previs (or Layout) department. At this point, previs takes the work of editorial and fleshes it out with low-rez digital sets and characters that are positioned and posed for key moments. Each scene then becomes an individual Maya file and gets iterated through a review process that ends when every scene in a sequence is approved by the director. One unique aspect of our pipeline is the fact that the previs stage is both the first and last time that a director will get to see entire sequences as contiguous pieces of his/her film. In other words, after previs, each shot takes on a life of its own, and the level of completeness of the film happens more or less randomly depending on how the evolution of each shot gets prioritized by each subsequent department.

From previs, (PRE) shots move onto a layout stage (LAY) where low-rez sets and rigs get swapped out for hi-rez geometry. Once approved, shots move onto a keypose stage (KEY) whose sole purpose is to block out performances prior to moving into an animation stage (ANI) where all performances are fully fleshed out to the director's satisfaction. Once each shot is approved in animation, it will go through various intermediary reviews that will eventually culminate in a color stage (COL) where shots get lit and rendered for final approval.

Figure 15: *Digital workflow at DTS prior to stereo.* © Disney Enterprises, Inc.

Stereo Pipeline for *The Secret of the Wings*

Obviously, the description presented above of the digital pipeline in place at DisneyToon Studios is a broad simplification of the tremendous amount of work that making a fully animated film requires, but, as far as thinking about creating a new placeholder for stereo, it does present an accurate picture of the major landmarks that we had to take into account in coming up with a stereo workflow. Given this relatively straightforward process, the first

question we had to answer when it came to stereo was, where would it make most sense to insert a new stage? Given that efficiency was a key factor, we found that the most obvious place to insert stereo would be right after the animation (ANI) stage since at that point each scene would, for all practical purposes, be locked.

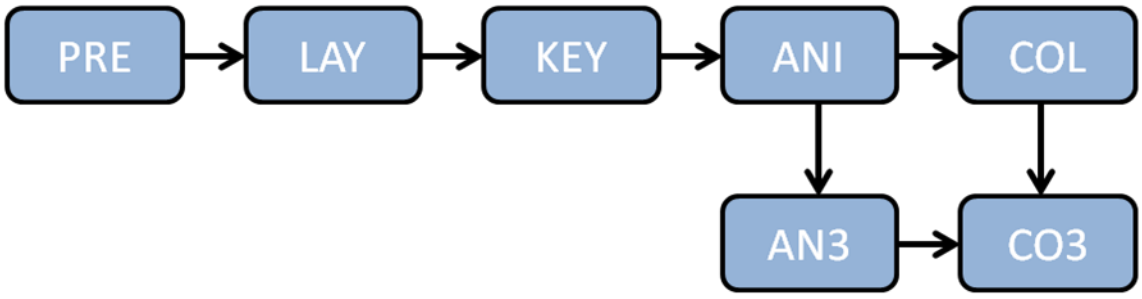

Figure 16: *Simple workflow that triggers a stereo animation stage (AN3) after animation (ANI).*

Working with animation files would ensure that we would be able to do stereo design with the exact assets, staging and camera placement that would eventually be rendered. In practice, however, using the animation stage as our stereo launching point had two significant drawbacks. Firstly, because the manner in which shots flow through our pipeline, we could never predict when an entire sequence would reach 100% approval. This meant that at no point would we have an opportunity to design stereo with *sequence continuity* in mind. From a creative perspective, this was not an ideal starting point. The second drawback became evident once we started doing tests with this stage and realized that given the file sizes in question, it would not be feasible for one artist to work through the entire film in a timely fashion.

The next obvious place to look for an insertion point for stereo was right after previs (PRE) Introducing stereo so much earlier in the pipe would allow for entire sequences to go through, and given the smaller file-size footprints, would in theory allow for much quicker turn-around times. However, designing stereo at this stage would not accommodate for the creative changes that the film would undergo as a natural part of the workflow.

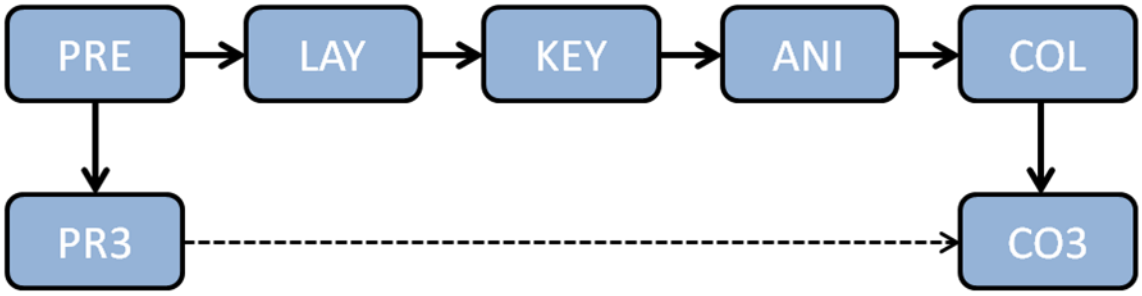

Figure 17: *Simple workflow that triggers a stereo previs stage (PR3) after previs (PRE).*

In the end, the process we put in place was a hybrid approach whereby we would design stereo right after PRE, in a stage we called PR3, out of which the only thing that we published was a stereo camera asset. Later on in the pipeline, once a shot would be approved from the animation stage (ANI) we would create a new stage AN3 into which we would import the previously saved camera from PR3. We would then create AN3 hardware stereo renders which we would compare with those created at the PR3 stage. If stereo needed to be updated, the stereographer would publish a new stereo rig. Otherwise, the original PR3 rig would move through the rest of the pipeline all the way through the stereo color stage or CO3.

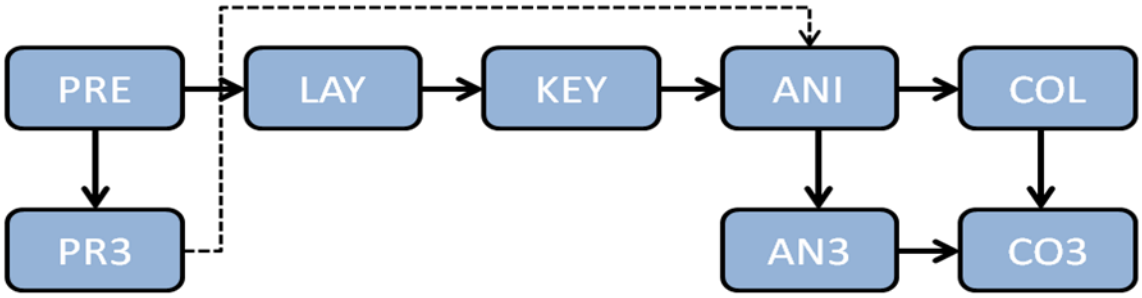

Figure 18: *Stereo workflow that triggers a stereo previs stage (PR3) and a stereo animation stage (AN3) that, if possible, reuses the camera rig created at PR3.*

This mechanism of creating a camera asset that could have its own life trajectory independent of the rest of production was instrumental in allowing us to create a pipeline that was stereo aware only at the most crucial stages. Its flexibility also made it possible for the entire film to go through a single individual during the design and final review phases.

We should note at this point that the complete stereo workflow we put in place for *The Secret of the Wings* encompasses one more stage past the CO3 color stage which we call Stereo Balancing, or BA3. This final step is a stereo grading process that we've put in place to ensure stereo continuity for the entire film. We will have more to say about this stage later on in the talk.

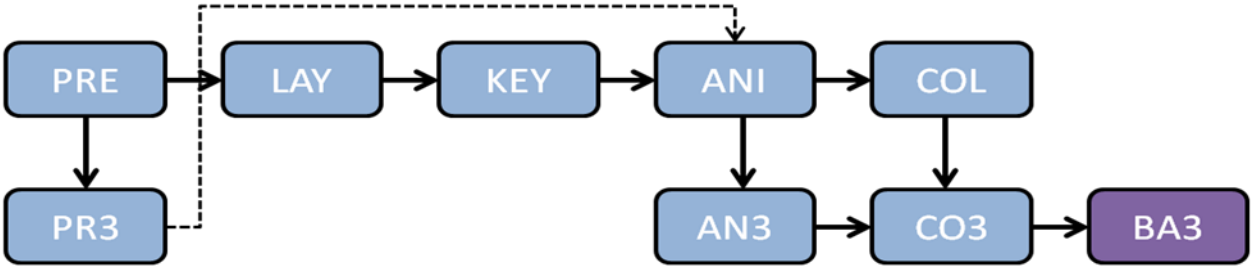

Figure 19: *Stereo workflow that incorporates a stereo balancing stage (BA3) which ensures stereo continuity throughout the film.* © Disney Enterprises, Inc

So, now that we have a very neat picture of the basic workflow we created for *The Secret of the Wings*, we'd like to say a few words about the evolution of stereo within the studio. As any producer can attest, what works brilliantly for a project, may not necessarily fit the creative needs of the next film. In fact, stereo workflow was to be one of the first areas we had to revisit when we began to plan for another film outside the Fairies world. The question we had to ask ourselves was how we could use the best of what we had built *for The Secret of the Wings* while allowing for the creative needs of the new franchise. Pretty soon, the notion of stereo workflow came down to two separate but related ideas: technology and process.

We will cover the technology aspect of stereo later on in the course, but one idea we'd like to highlight was the fact that, from a technology perspective, the studio was one hundred percent behind the toolset that we had built for *The Secret of the Wings*. In particular, one aspect of this toolset that we wanted to preserve was having a camera rig that, regardless of the underlying processes that would move this camera from stage to stage, would maintain its *structure*. In other words, when it came to stereo, at any given stage, the only currency that mattered was the camera. It was this very simple idea, in fact, that gave shape to the tools we developed to create and track this asset, and allowed us to easily build variations of the stereo pipeline we had put in place for *The Secret of the Wings*.

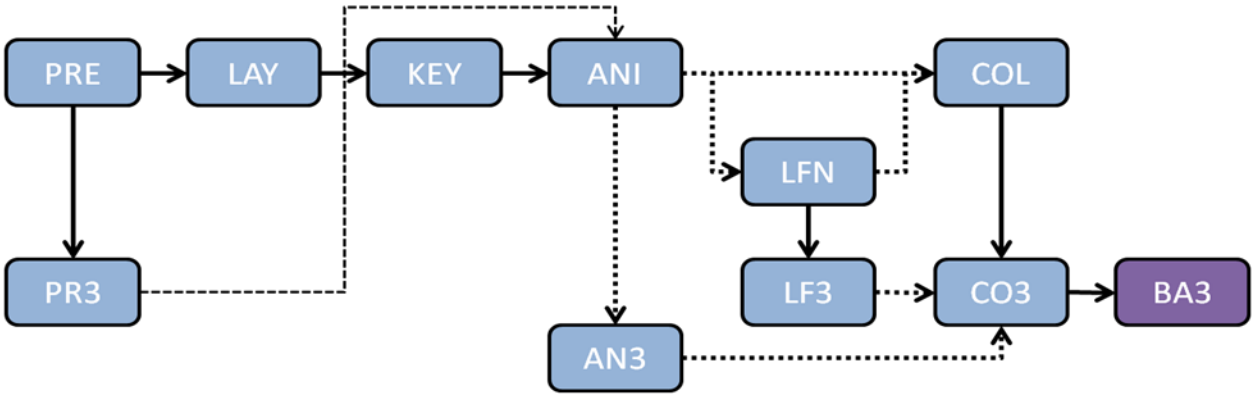

Figure 20: *A stereo workflow that may trigger a layout finaling stage (LFN) for the purposes of camera polishing, and a stereo layout finaling stage (LF3) that incorporates the camera data from LFN.*

Stereo Workflow Part II: *Workflow Principles*

Now that we have presented the stereo workflow we established for *The Secret of the Wings*, let's take a step back and consider a more global view of stereo design. Arguably, the workflow we have outlined above allows for a very generous treatment of stereo which shepherds shots across the entire production in five discrete stages:

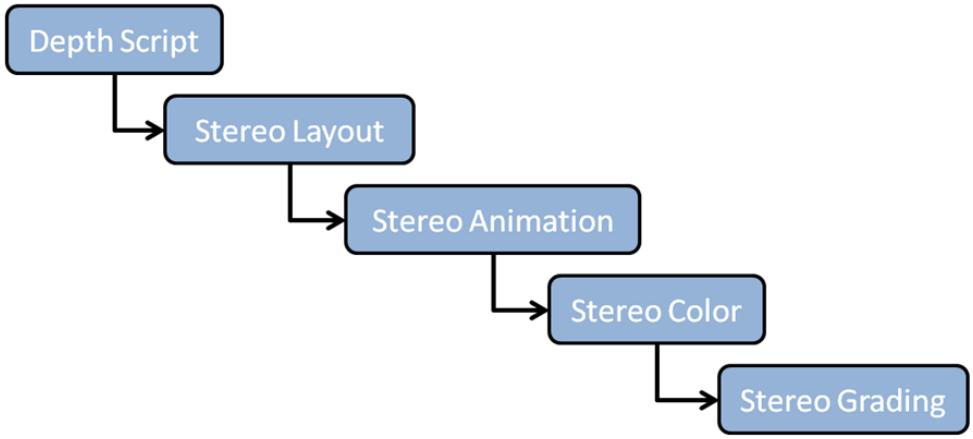

Figure 21: *Proposed stereo workflow for an animated CG film.* © Disney Enterprises, Inc.

Let's take a look at what each of these stages is designed to accomplish:

Depth Script: This is the place to look at the entirety of a film to determine how the stereo experience can enhance a film from a story perspective. Because this stage can (and should) happen very early in the film-planning stages, it offers the stereographer the possibility of influencing the style of the film to maximize the impact of stereo. More importantly, this stage is instrumental in establishing a common language with the director that the stereographer will use to communicate his/her ideas for the remainder of the film.

Stereo Layout: This is the place to execute all the ideas set forth in the depth script. It happens early enough in the film to allow room for experimentation, and as we found with *The Secret of the Wings*, it gave us the ability to design stereo for entire sequences at a time. Given that audio design was fairly mature at this stage, we were able to have a very thorough sense of how the entire film was shaping up in 3D. Generally speaking, having access to audio at this

stage opens up the possibility of using audio queues to motivate stereo choices. As an example of how this can be done, sad moments in a film could gradually sync the flattening of the spatial depth to the film's score.

Stereo Animation: Once filmic choices have been locked down in terms of shot length, staging, camera design and cinematography, stereo design should be checked against earlier choices on a shot-by-shot-basis—ideally, if the film has not changed significantly from its original concept, the percentage of shots that need to be redone could be quite small. Regardless of the number of shots that could need a revisit, the overall mission for stereo should be to maintain a flow that adheres to a common vision. When we step back and look at the role of stereo, we want to emphasize that, while artistic freedom is important, for stereo to work as an integral contributor to a film, the end result must deliver a consistent experience that an audience will enjoy. Besides serving the story, the goal of a stereographer should at minimum be to make stereo as unobtrusive and painless a presence as possible. The less eye strain we can induce in an audience the better. Incidentally, we find that the films that audiences find most difficult to experience have two characteristics: One, they make use of aggressive depth settings for extended periods of time, and two, they use depth settings that, across multiple shots, require the viewer to change focus either too drastically, too frequently, or both. For these reasons, any stereo workflow should include check points that ensure that stereo flows pleasingly from beginning to end.

Stereo Color: While all the previous stages ensure the consistency of stereo design across a show, the primary purpose of this stage is to ensure that the execution of stereo for each shot is as free of flaws as possible. Finding and fixing render inconsistencies is one of the most important roles of a stereo team, and this stage is one of the stereographer's primary weapons in terms of ensuring quality control.

Stereo inconsistencies come in many flavors, and it takes a trained eye and multiple viewings to catch them all, but generally, they boil down to four categories:

Render Inconsistencies: The kind where the same element or character will appear differently in one eye from the other. The resulting image pair may, upon first viewing, register as correct, but closer inspection will reveal a kind of shimmer that makes the source of the problem difficult to look at. The solution often entails a re-rendering of the object/character in question.

Compositing Inconsistencies: The kind where all or a portion of the image might appear at the wrong depth, mis-aligned, or in some cases, reversed. Depending on the severity of the inconsistency, the problem might be very easy or very difficult to diagnose. The problem is

almost always a manifestation of human error of some kind where different treatment was inadvertently given to one eye from the other. The solution entails revisiting the composite network to rectify the problem.

Edge-of-frame Inconsistencies: Due to the fact that the stereo effect is achieved by rendering a scene from two cameras that are offset by a small distance, each camera will see a small portion of a scene at the edges of frame that the other cannot. This results in a stereo phenomenon that is most commonly referred as a retinal rivalry. Under certain circumstances retinal rivalries can be difficult for the human brain to process. The solution to this problem is to cover these edge-of-frame violations with thin, black masks that we have come to term "stereo windows." Later on in the course, we will dedicate a section to stereo window design.

Reflective Inconsistencies: Shiny surfaces by definition reflect their environment, and the light sources that illuminate them. One problem that arises in stereo is that for a given object, depending on its distance from camera, these reflections and specular highlights will be slightly offset in each eye, thus changing the appearance of reflective objects in stereo pairs. Most of the times these inconsistencies are nearly imperceptible and can be ignored, but in more extreme situations, these inconsistencies can be distracting and must be addressed. The solution can be as simple as decreasing the amount of reflectivity to hide the inconsistency as much as possible, or as complex as implementing shaders that cheat the positioning of reflections so as to avoid the problem entirely. For *The Secret of the Wings*, we opted for the simple approach of toning down specular highlights for stereo to satisfactory effect.

While the Stereo Color Stage is certainly a good place to quality check the integrity of the stereo choices we've put in place, we should point out that it is also a good place to make special design decisions that may not necessarily make sense for the mono version of the film. For instance, in *The Secret of the Wings*, we found that the depth of field choices (DOF) that had been implemented for the mono version of the film were generally too aggressive for stereo. As a result, early on in the show we implemented a rule whereby we would, as a first order approach, cut down the amount of DOF in stereo frames by half of the amount used in mono. We would then compare the resulting stereo shot with its mono counterpart, and decide how we would proceed with each shot. Typically, reduced DOF settings worked best for stereo. Another instance in which we allowed flexibility in modifying the appearance of the stereo versus the mono version of the film had to do with the removal of framing devices that appeared distracting in stereo. Thus, items not prominently featured in shots such as edge-of-frame books, rocks, or the occasional flower, were struck from stereo shots altogether.

Ultimately, how a given production chooses to treat the stereo version of a film compared to its monoscopic counterpart has an impact on the creative choices available to the stereographer, especially at the Color stage. On *The Secret of the Wings*, one of the original design parameters

for stereo was that, in order to increase efficiency, the left eye would be used as the mono version of the film. Due to the stereo-specific needs mentioned above, what in fact took place was that the mono version of the film became a starting point for creating the left eye render. From that point on, the mono and stereo versions of the film took a life of their own. On a different production with a more streamlined workflow, it might make more sense to make all the design choices early on to better cater for stereo, and use the resulting left eye (or right eye for that matter) as the mono version of the film.

Depth Grading: One question that often comes up with regard to stereo execution goes something like this: Once the stereo version of my movie is done, is there anything I can do to change/modify stereo in a way that does not require me to revisit compositing, or worse, rendering? The answer is yes, and the process of doing so is called depth grading. It turns out that by changing the relative position of a final left-eye/right-eye pair, one can change the appearance of stereo fairly dramatically.

Below we illustrate this concept in terms of horizontal pixel offsets introduced on the right-eye of a stereo pair.

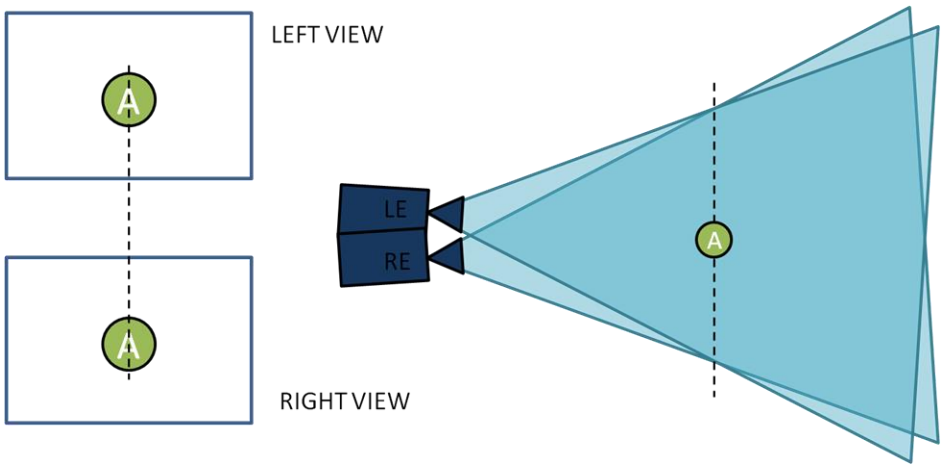

Figure 22: *Stereo pair with no pixel offsets.* © Disney Enterprises, Inc.

The figure above shows a stereo pair in which the subject is exactly at screen-space. Notice however what happens when the right-eye is shifted by some pixel distance to the right:

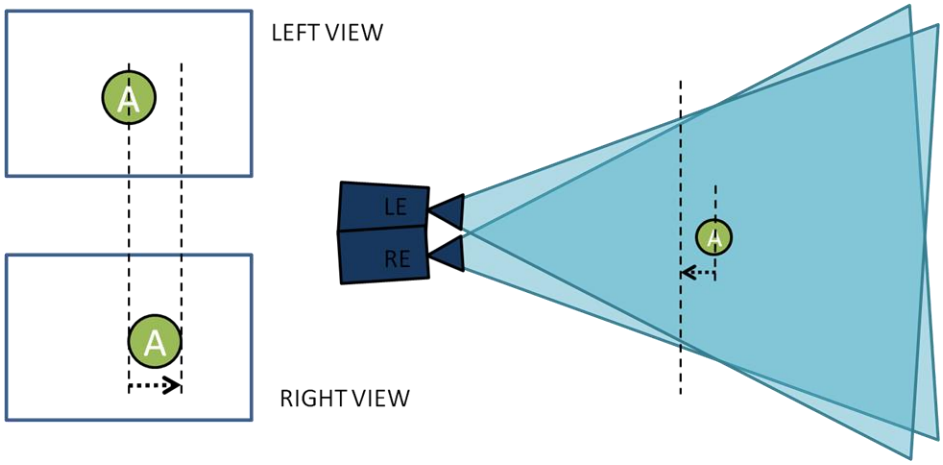

Figure 23: *Introducing a positive offset to the right eye pushes A behind the screen.*

It turns out that, introducing a positive pixel offset changes the convergence point of the stereo pair in such a way that, because of the shift, the subject appears to be pushed **back** behind screen-space. (See above)

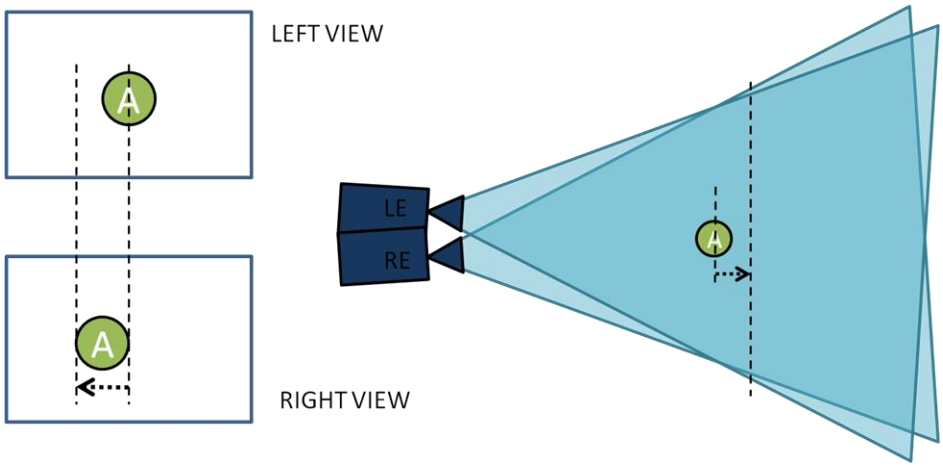

Figure 24: *Introducing a negative offset to the right eye pushes A in front of the screen.*

If, on the other hand, we introduce the same pixel offset in the opposite direction, convergence gets pushed so that the subject now appears to move **in front of** screen-space. (See above)

In terms of what manipulation of pixel offsets means for adjusting convergence in this manner, it is interesting to note that pixel offsets in the range of 5-10 pixels have quite drastic effects on stereo with minimal image quality impact. While this process can't change stereo qualities of individual elements within a shot, it does allow us to modify how stereo appears relative to the viewing screen. Next to Stereo Layout, an implementation of depth grading in a stereo pipeline is of paramount importance as a means to tune stereo across an entire film. While it is conceptually an easy process to implement on a production-grade compositing package, for it to be useful, it should ideally be carried out once the film is 100% complete in stereo. Depending on the implementation, an entire film can be depth-graded by a single individual over the course of two to four weeks. However, because this stage happens so late in the production schedule, allowing sufficient time to execute it can be challenging, but its benefits are well worth the incurred cost. On *The Secret of the Wings*, we allowed one week to depth-grade the entire film, a time frame which was overly aggressive but which ultimately we were able to meet thanks to an inspired collaborative effort that pulled in teammates from other projects.

Overview of Stereo Concepts

The way we perceive depth involves much more than our ability to fuse two images together into one volumetric view of the world. The rich history of art has, to a large extent, been a quest to create culturally relevant interpretations of these cues. With regard to photography and film, in fact it can be said that these disciplines communicate depth in the sense that photographs and motion pictures not only provide a faithful record of the world we perceive, but also render artful manipulations of cues we use daily to interpret how we view the world. (For a thorough discussion of this topic, see Lipton, pp. 54-61.) An understanding of depth then necessitates familiarity with the full range of visual elements we use to perceive it. Below is a summary of these cues:

Retinal Image Size: When presented with two similar objects, our brains interpret larger retinal images as being physically closer to us. Retinal image sizing is not only a significant depth cue, but it also forms the basis of composition.

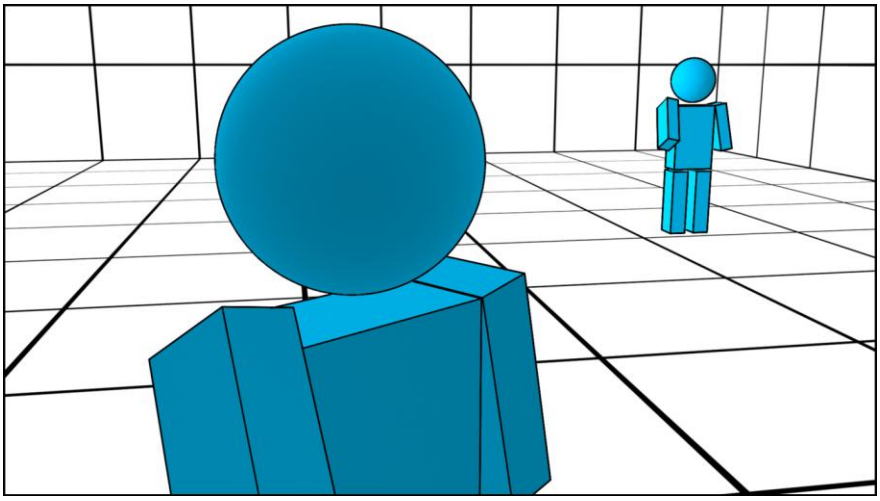

Figure 25: *Foreground character appears closer to the viewer due to his size relative to the background character.* © Disney Enterprises, Inc.

Perspective: This cue is based on the notion that objects diminish in size as they recede from the observer. In the arts, the development of this concept took many centuries to develop, but a true mastery of it became widespread during the 15th Century amongst Renaissance artists.

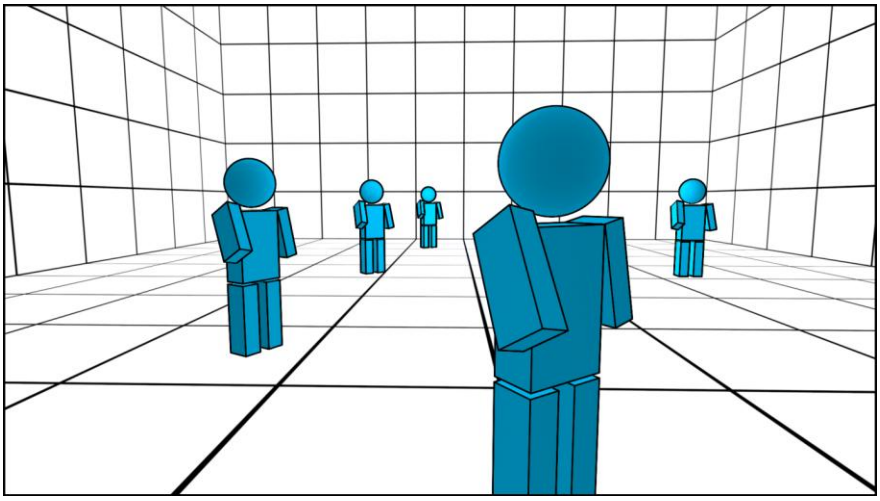

Figure 26: *Objects diminish in size as they recede from the observer.* © Disney Enterprises, Inc.

Interposition or Overlapping: When presented with a complex scene, our brain uses this cue to help us discern which objects are closer to us by virtue of the fact that objects that are near us may partially or fully hide objects that are behind.

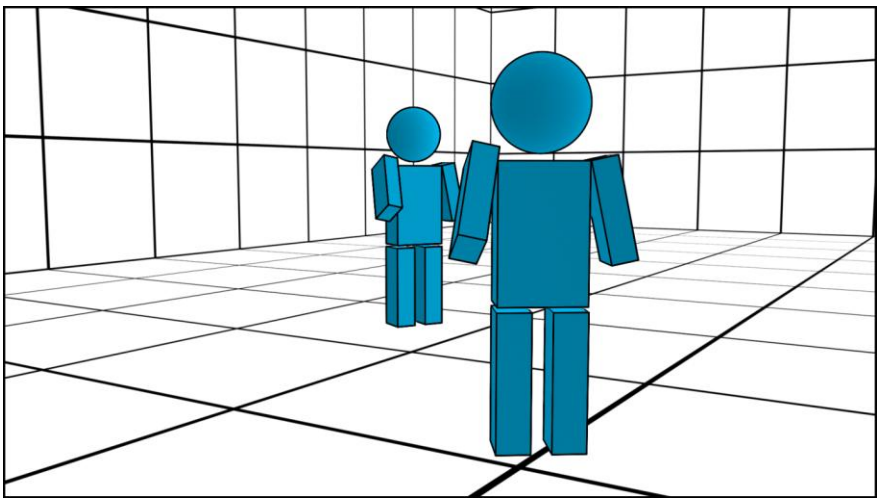

Figure 27: *Overlapping as a visual cue.* © Disney Enterprises, Inc.

Aerial Perspective: This cue tells our brain that, by virtue of the fact that, with distance, we can actually perceive the air we breathe as haze, the farther an object is from us, the more it is hidden by the surrounding atmosphere. Thus, as in the example below, mountains that are farther away have a qualitatively more subdued color quality than mountains that are near to us.

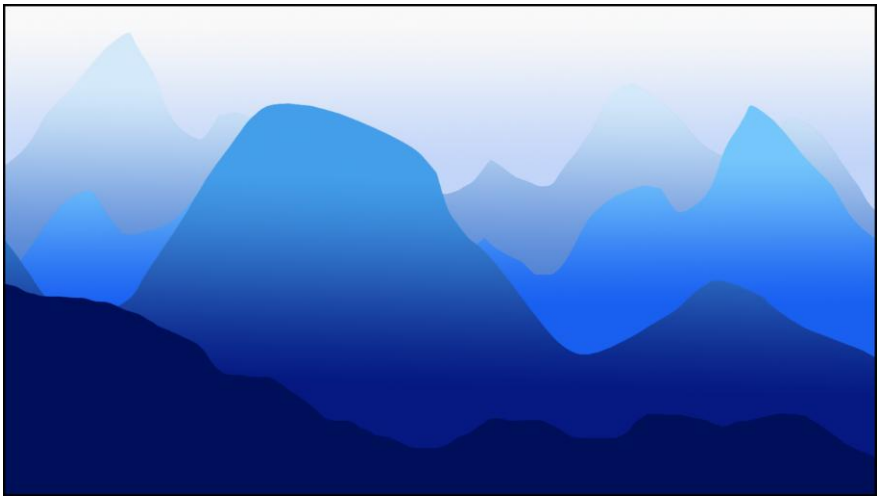

Figure 28: *Atmospheric haze de-saturates the color of distant mountains.* © Disney Enterprises, Inc.

Light and Shade: These depth cues provide information about the volume of space filled by objects, and how far apart one object is from another. Most importantly, how objects are lit gives our brains essential clues as to their volumetric qualities.

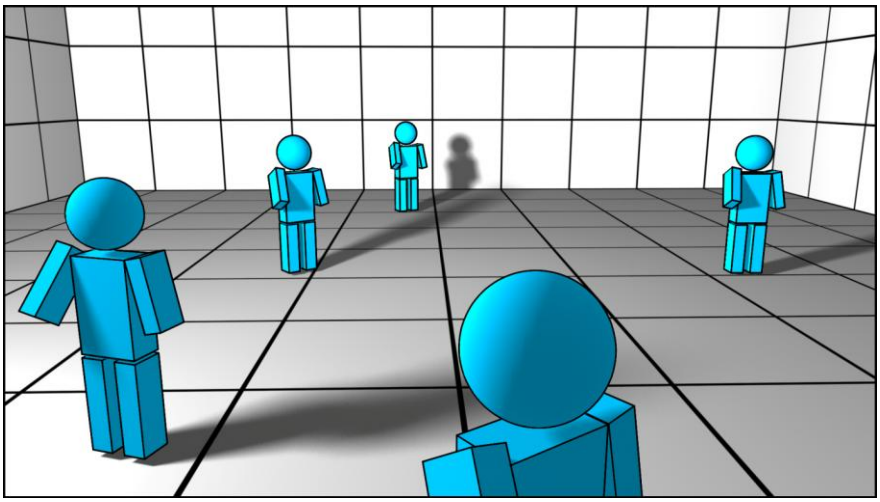

Figure 29: *Shading and shadowing give valuable depth cues.* © Disney Enterprises, Inc.

Textural Gradient: This depth cue tells us a great deal of information about how objects fill space. For instance, in the picture below, we can see in detail the qualitative aspects of the grid where the foreground character stands. However, as we recede back into space, this information becomes simplified, and, if we move far enough in distance, lost entirely.

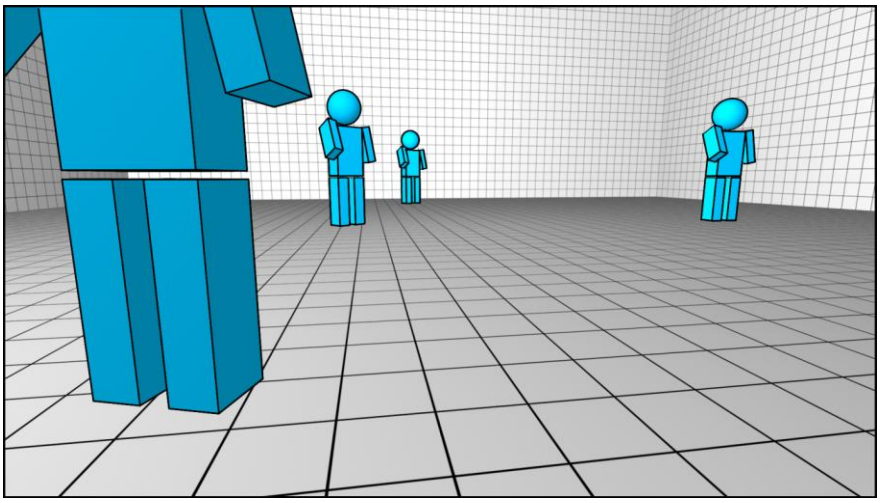

Figure 30: *Texture detail is more prominent in the foreground than in the background.*

Motion parallax: This cue is a post-Renaissance discovery, and similar to retinal size, it tells the brain that the movement of objects that are closer to us is more dramatic than that of objects that are far away. In the picture below, the right-to-left camera motion makes the left-right translation of the foreground character appear significantly larger than that of the mid or background characters.

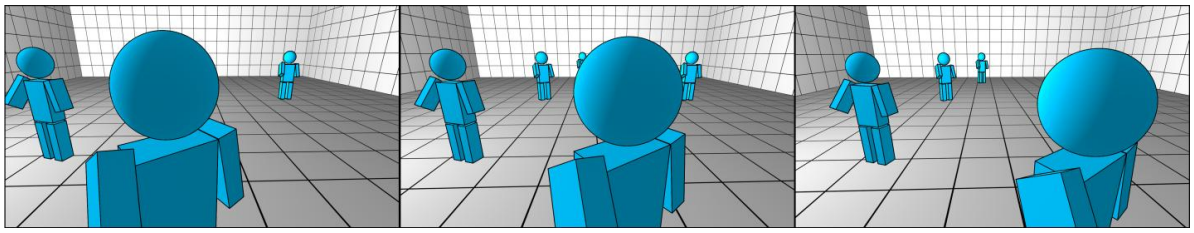

Figure 31: *The closer an object moves relative to the camera, the more motion parallax it will exhibit.*

Accommodation: In the same manner that a camera lens focuses by moving closer or farther away from the film, our eyes focus images on the retina by changing the shape of our lenses. This muscular effort involved in focusing also helps our brain in gauging depth.

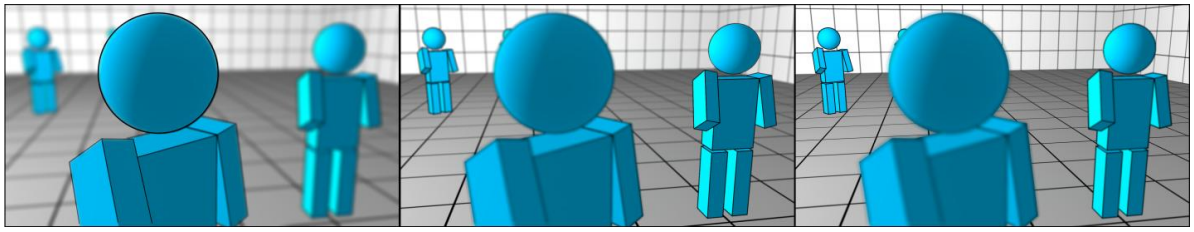

Figure 32: *Depth and focusing. Focus occurs at 1) foreground on the left view, 2) the midground on the center view and 3) on the background at far right.* © Disney Enterprises, Inc.

Convergence and Disparity: These set of cues depend on the brain's ability to fuse the two separate views of the world it receives from both eyes. The point at which we direct our eyes to focus on a scene becomes a point at which both left and right views see exactly the same thing. This phenomenon is called *convergence*. As a result, when the eyes converge on a particular object in space, all other objects in front of or behind the point of convergence become double images. This phenomenon is called *disparity*. As illustrated in the anaglyph triplet below, convergence moves from the foreground character in the first image, to the midground character on the second, and finally, the left background character in the last image.

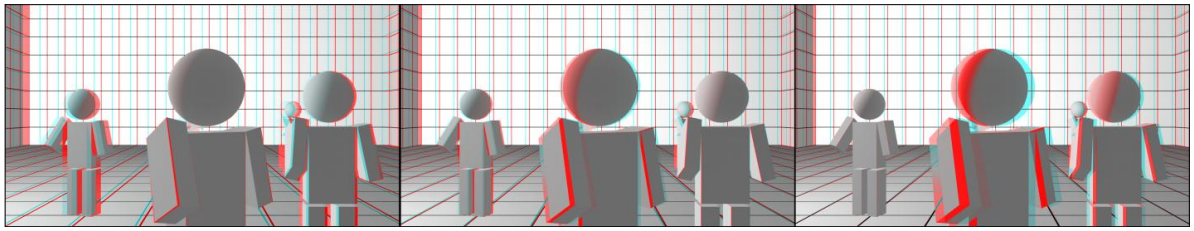

Figure 33: *Convergence occurs at 1) foreground on the left view, 2) on the midground on the center view and 3) on the background at far right.* © Disney Enterprises, Inc.

Stereopsis: The ability for the brain to fuse left-eye/right-eye image pairs into one volumetric view of the world is called stereopsis. Of all the ways in which the brains processes visual input to gauge depth, only convergence, disparity and stereopsis depends on our having two eyes. Interestingly, although all species that have eyes possess two eyes, there is a very small percentage of animals that can actually see in stereo. For stereopsis to be possible both eyes must be able to converge on an object so that the image can be fused into a single three-dimensional view. For humans, the development of stereoscopic vision appears to have been one of the essential factors in the evolution of the human intellect.

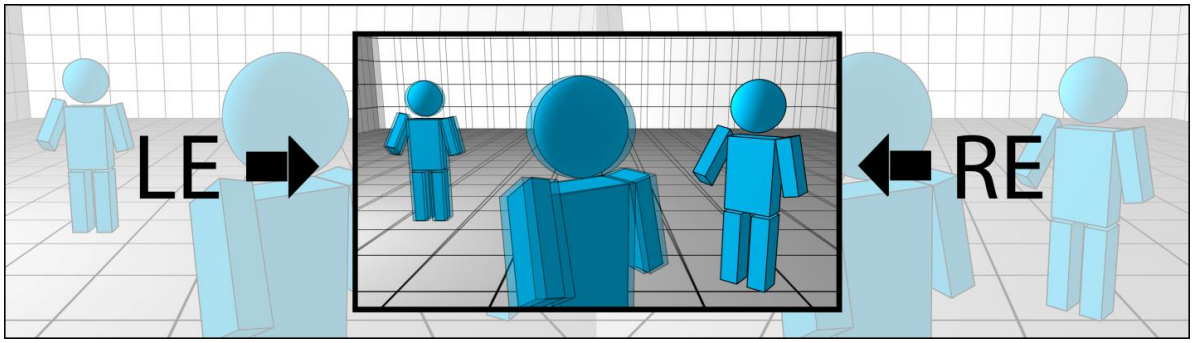

Figure 34: *A left-eye/right eye image pair fused into one volumetric image.* © Disney Enterprises, Inc.

Stereo Techniques: Creating Stereo Pairs

We have thus far established that what is needed for an enjoyable stereo experience involves the creation of image pairs that we can comfortably fuse into one volumetric image. During this portion of the course, we will provide an overview of the different methods available for the creation of stereo pairs.

- *Two-Camera rendering:* For 3D productions in which all of the content exists digitally as a part of a scene file, this is by far the easiest, most straightforward method of creating stereo. In addition to a 3D set, all that is required is a left-eye/right-eye camera rig, as well as some method of controlling the distance between these cameras (interaxial separation) and the distance at which these two cameras converge (convergence). Below we illustrate the effects of changing interaxial distance and convergence on stereo:

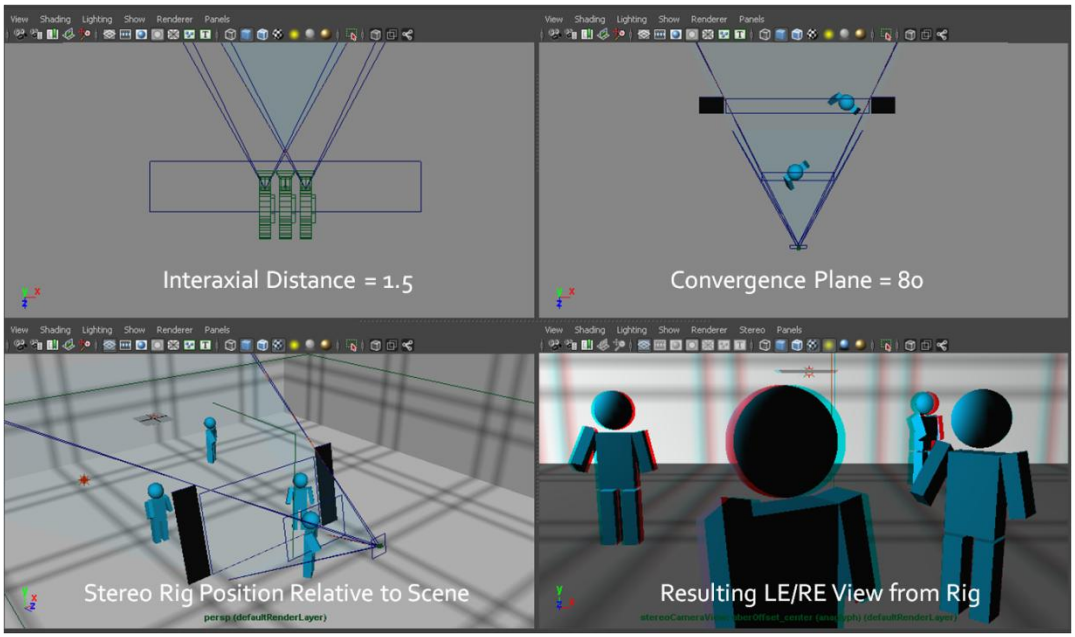

Figure 35: *Setting interaxial distance and convergence to generate a stereo pair.*

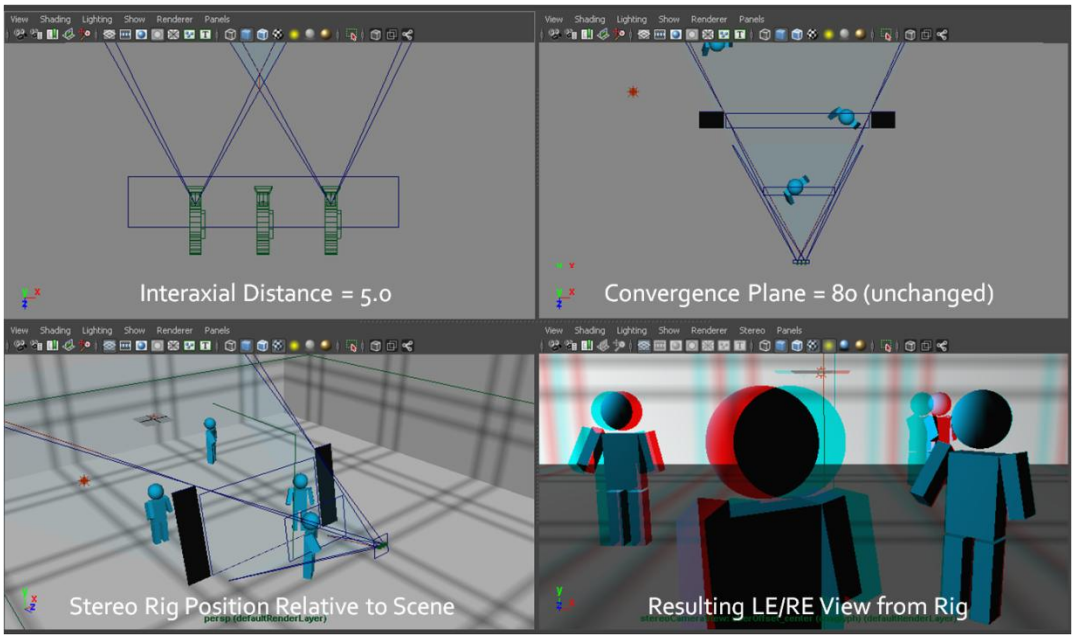

Figure 36: *Adjusting interaxial distance while keeping convergence fixed on the midground character. Notice the large amount of resulting disparity on the foreground and background character.* © Disney Enterprises, Inc.

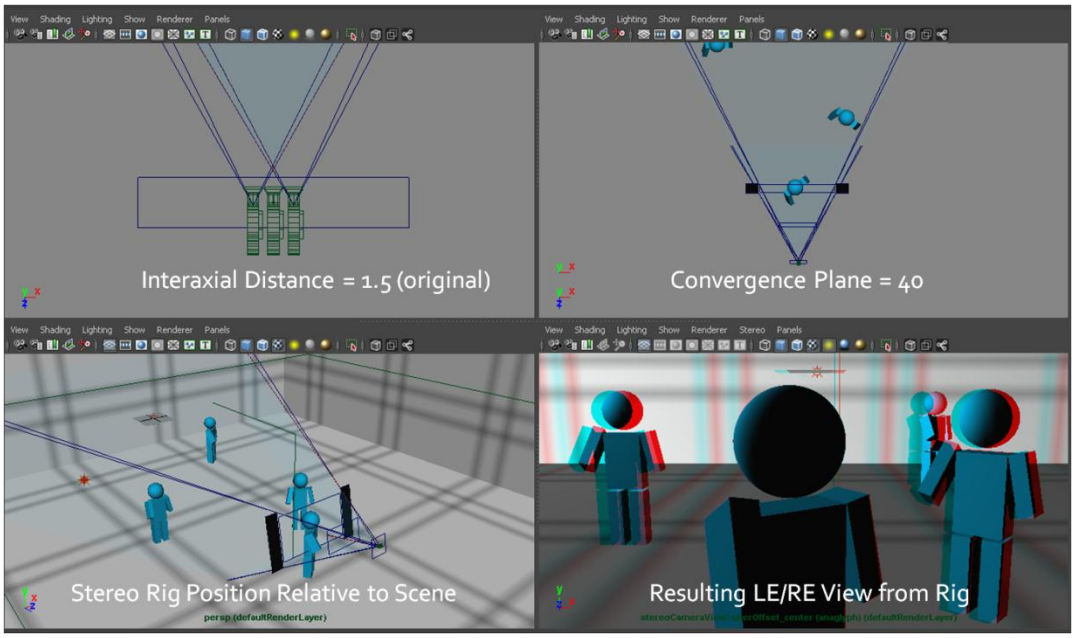

Figure 37: *Adjusting convergence while keeping interaxial distance fixed. Notice the large amount of resulting disparity on the background characters.* © Disney Enterprises, Inc.

The advantages to using this method are numerous and, when set up as an integral part of a digital pipeline, provides the most straightforward mechanism for creating a stereoscopic movie. However, without adequate resources and a realistic

timeline, this technique can quickly become difficult to manage, and the sheer amount of data generated (think every layer of every shot must now be created twice) can bring a digital production to its knees. In addition, from a compositing point of view, the management of left-eye/right-eye resources must be kept under tight supervision.

- *Z-Depth compositing:* This method can be used in 3D productions, mostly as a resources-saving alternative. In concept, this method re-purposes all the elements that were generated for the creation of one eye and, through the use of individually rendered depth-maps, a second perspective view can be generated. The advantages of this method are measured strictly in terms of time and budget savings from rendering just one eye. The disadvantages however quickly add up, as there is in fact a very small segment of shots for which this technique works convincingly without too many visual artifacts that degrade stereo pairs.

- *Stereo conversion:* The process of making stereo images out of traditional 2D content (live action or animated) is called stereo conversion, or dimensionalization. With the growing popularity of 3D titles, this process has become a very popular way to create stereo. (For an excellent article on the different methods of post-conversion, please refer to Mike Seymour's excellent article on www.fxguide.com included in the references.) Essentially, this technique achieves the stereo effect by processing a 2D image, and using the information within it to create a secondary view that will complete a stereo pair. Starting in 1998, variations of this technique have been put to use to convert many popular films such as *Alice in Wonderland* (2010), *The Green Lantern* (2011) and *Thor* (2011).

  Stereo conversion is a very labor intensive process because all the depth subtleties found in two-camera renders must now be created *by hand* on a frame-by-frame basis. Doing this convincingly and consistently is, needless to say, very difficult to achieve. As of the writing of this paper, there are several companies that specialize in stereo conversion, each of which has in place proprietary tools and teams of artists to help them achieve nuanced results. Among the most well-known conversion companies are Stereo D, Legend3D and Prime Focus.

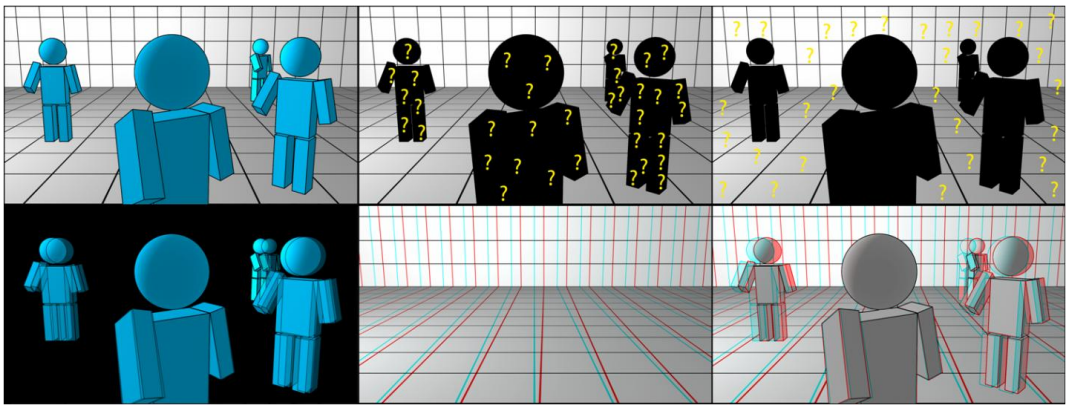

Figure 38: *Stereo conversion in a nutshell. Start out with a 2D image and, using any number of techniques, as well as painstaking artistry, break up image into separate elements and/or layers, all the while ensuring that wherever objects overlap, new image data can be generated in the occluded regions for a second eye projection. The reconstructed image is then shot through an offset camera to complete a stereo pair.* © Disney Enterprises, Inc.

- *Compositing via pixel offsets:* This technique is only possible when working with existing stereo composites. It is primarily useful for modifying the stereoscopic qualities of *individual elements* within a composite. For instance, if a foreground element comes out into audience a little too aggressively, dialing the horizontal position of one of the eye renders can push that element into a more appropriate depth space. While this technique can be used to drastically change the stereoscopic quality of a shot, its primary disadvantage is that, without proper supervision, it can inadvertently introduce conflicting stereo cues that will break a shot.

- *Depth Grading:* This technique, as stated earlier in this course, can be used to change the convergence point of a stereo-pair simply by introducing a pixel offset to one of the eyes. This is a very powerful technique that is most useful to balance the stereo space across an entire movie quickly and with relatively few resources.

Stereo Techniques: Design Principles

If there is one overarching idea we would like to communicate throughout this course, it is the fact that, while creating stereo is, thanks to the wide selection of tools the modern film-maker has at his/her disposal, easy to achieve, creating a great stereo experience requires a knowledgeable eye, and the means to execute a vision that spans the entirety of a film. We have thus broken down the different areas in a given production where a film-maker can make stereo choices starting from the depth script, then stereo layout, and progressing all the way through stereo color and finishing with depth grading. However, we have made little mention of the nature of those choices or how to achieve them. Therefore, at this point in the course, we will turn our attention to stereo design, and how adhering to some simple rules early on in the film-planning stages can open up creative choices that will be useful for stereo further down the pipe.

Depth Budget

One of the first questions that must be addressed early on in the planning stages for stereo is what is the depth budget (in terms of pixels) for the film? In other words, what pixel ranges will map to *strong*, *normal* and *weak* stereo moments. Earlier in the course, we introduced the parameters that DisneyToon Studios generally adheres to:

Depth Strength (pixels)	Negative Space (in front of screen)	Positive Space (behind screen)
Weak	0-8	5-15
Normal	8-15	15-30
Strong	15-25	30-40

Figure 39: *Depth strength for DisneyToon Studios.* © Disney Enterprises, Inc.

For *The Secret of the Wings*, these parameters provided a starting point that stereographer Sierra used to begin shaping his vision of the film. Then, as director Holmes' vision became clearer, these parameters were modified to fit the final stereo treatment. The resulting depth budget leaned heavily on the usage of negative space, while positive space (i.e. in front of the screen) moments were kept to a minimum.

Depth Strength (pixels)	Negative Space (in front of screen)	Positive Space (behind screen)
Weak	0	0-8
Normal	0-5	8-20
Strong	8-10	20-30

Figure 40**:** *Depth strength used for* ***The Secret of the Wings.*** © Disney Enterprises, Inc.

Creating Layered Depth

Stereo is most effective when a filmmaker can immerse an audience in a rich environment that offers visual cues in the foreground, middle ground, and background. In the illustrations below, it is compositionally obvious that the picture on the right offers more visual depth cues than in the picture on the left. It should be of no surprise then that the composition on the right would lead to a richer stereo experience.

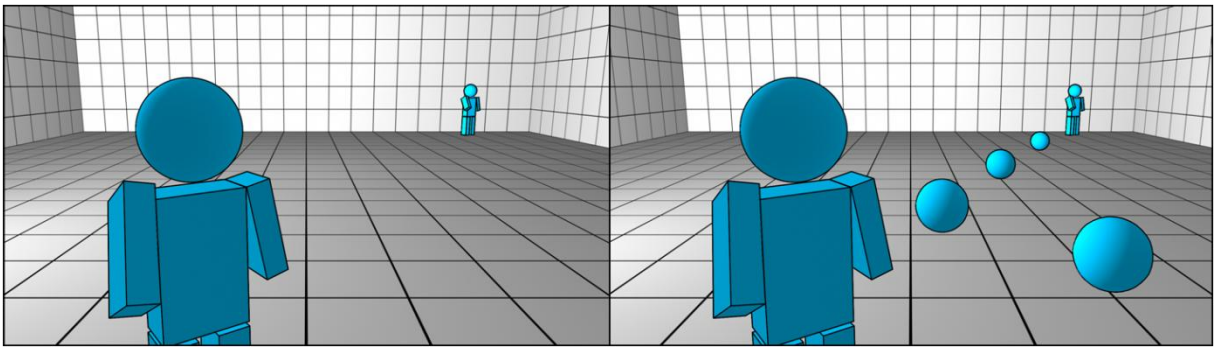

Figure 41: *The more visual cues a composition has, the richer the resulting stereo experience.* © Disney Enterprises, Inc.

As a general rule, then, it is a good idea from a stereo perspective for shot compositions to have well-defined foreground, middle ground and background elements. Compositions that have one or more of these missing may be more difficult to work with. For instance, the illustration below shows a composition which only has foreground and background components.

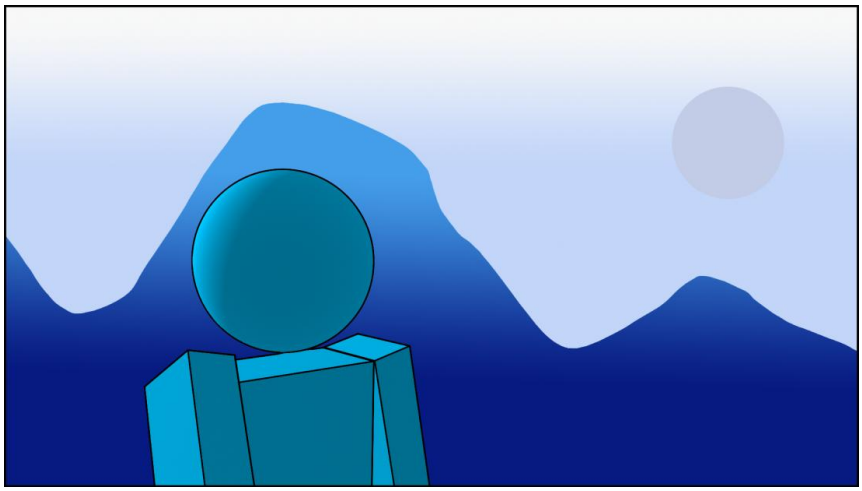

Figure 42: *A character standing in front of a distant background is not ideal for stereo.* © Disney Enterprises, Inc.

Given the scale of the scene, just about the only option that a stereographer would have in terms of design is to dedicate most of the stereo budget on the foreground character. The result would end up looking like a fully-rounded character standing in front of a flat background. Of course, in every film, there will invariably be shots for which the number of options available to the stereographer is just as limited, and that's okay. However, managing the number of such shots is ultimately a team effort in which everyone must be made aware of how composition in general affects stereo.

Layered Depth and Multi-Rigging

One of the most powerful tools available to the stereographer in terms of creating layered depth is the multi-rig. Multi-rigging allows a stereographer to create stereo settings that can be tied to specific geometry. The power of this tool is best illustrated by an example:

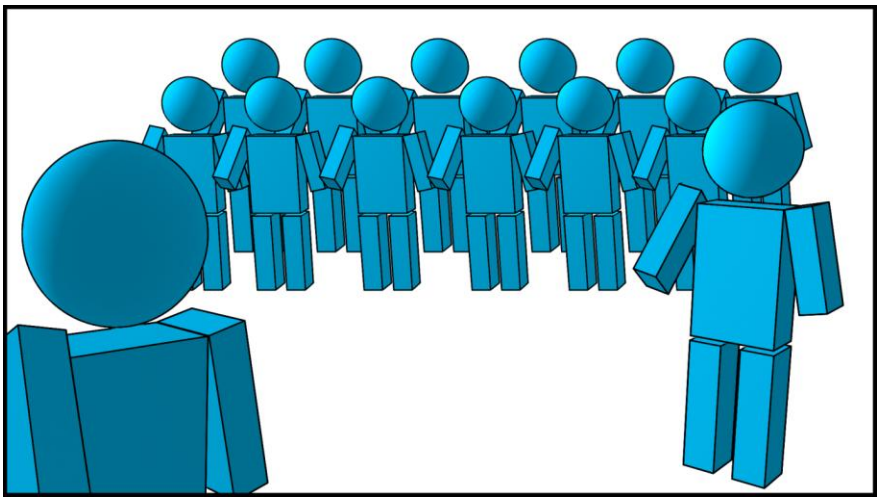

Figure 43: *A simple scene with well-defined foreground, middle ground and background.*

The composition above looks ideal for stereo, as it is clearly split up into foreground, middle ground and background elements. However, if the stereographer desires to set depth to put emphasis on the middle character, the resulting depth in certain situations could yield a stereo pair that puts the foreground element too far into audience-space:

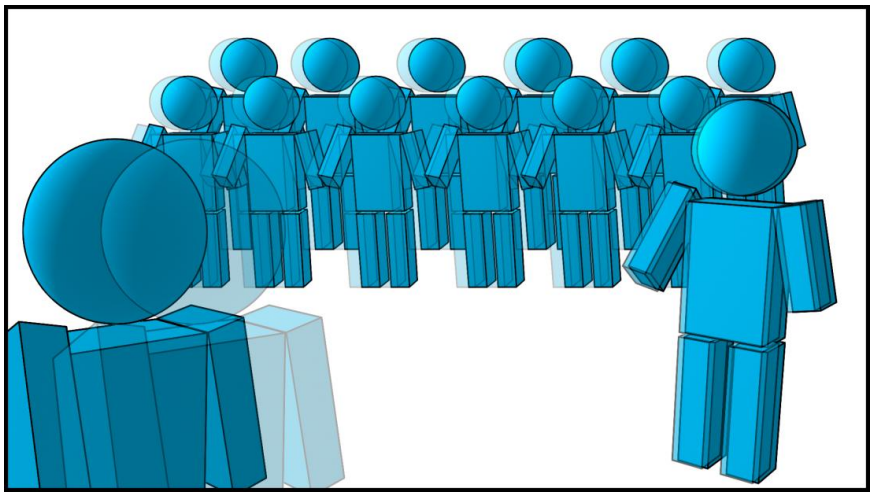

Figure 44: *A single stereo rig may result in unsatisfactory depth. In this case, the foreground character may be too far into audience space.* © Disney Enterprises, Inc.

While the resulting foreground and background may look great, the foreground element could cause extreme discomfort. Changing the stereo settings to fix this problem may yield a more satisfactory treatment for the background, but may flatten the rest of the composition too much, as shown below:

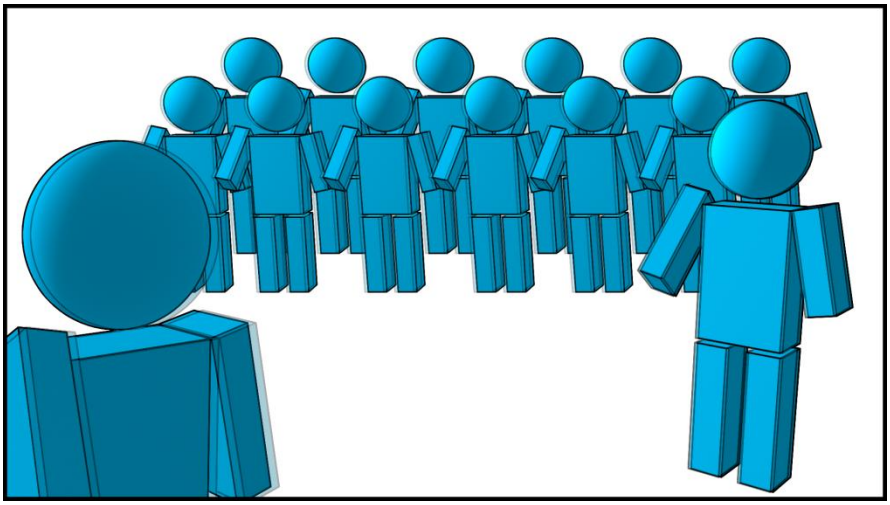

Figure 45: *Adjusting depth to accommodate foreground, flattens scene unacceptably.*

© Disney Enterprises, Inc.

Essentially, what the multi-rig concept allows the stereographer to do is to split up a scene into different regions that will be seen through separate stereo cameras. So, in the case of this scene, the foreground (red) could be shot with very conservative stereo settings, the

midground (yellow) with stronger settings, and the background (green) with more aggressive settings.

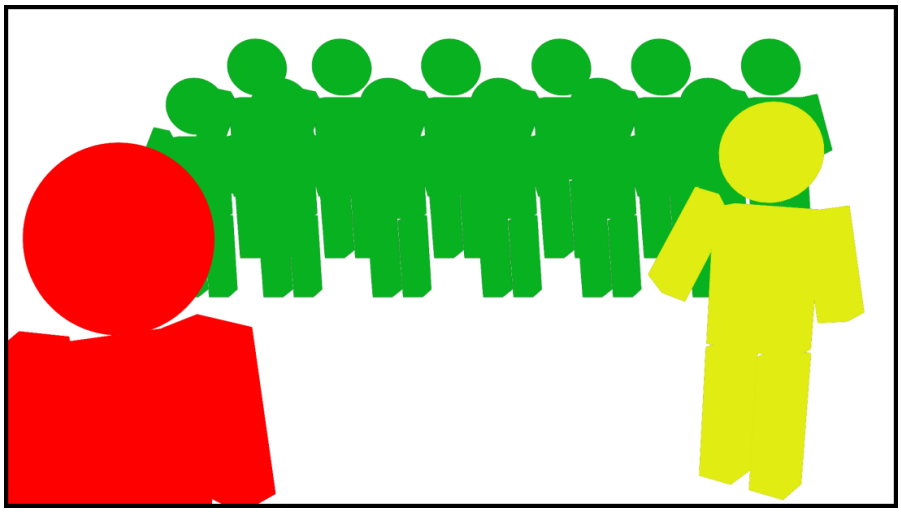

Figure 46: *Breaking up a scene into regions to be rendered by multiple stereo rigs may open up creative stereo design options.* © Disney Enterprises, Inc.

The end result would, if coordinated carefully, balance out the depth for all the layers, thus resulting in a richer stereo experience.

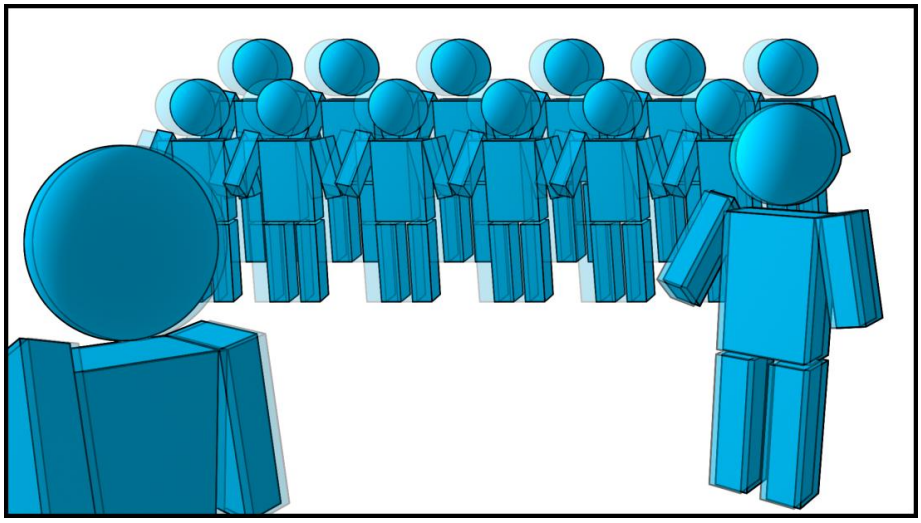

Figure 47: *A balanced stereo treatment rendered with multiple stereo rigs.* © Disney Enterprises, Inc.

Stereo and World Scale

As we have established throughout this course, the two means of changing depth in a two-camera set-up are by means of setting **convergence** and **interaxial distance**. Convergence allows us to define where left and right-eyes will appear as one, and interaxial distance defines how far apart our two views are from one another. These concepts are not just essential for stereo design, they're essential for stereo vision. The distance between our eyes is our built-in interaxial distance, and our eyes' ability to focus at different distances constitutes our convergence. These two facts play a pivotal role in how we perceive the world in terms of scale. A model of a car that we can hold in our hands appears qualitatively different to us than a life-size precisely because of their scale relative to us. If our eyes were further apart, the way we perceive our environment would appear proportionally smaller, and vice-versa.

This question of interaxial distance and scale poses an interesting problem for the stereographer in terms of how he/she chooses to interpret stereo for a given scene. While conceptually speaking, it might make sense to set the interaxial distance for a shot to be the same as the interaxial average for a human being, one would find very quickly that doing this in most cases would lead to a very flat and uninteresting stereo experience. On the other hand, modifying stereo parameters too aggressively might lead to a problem that's come to be known as miniturization, in which sets and characters take on diminutive dimensions that may rob it of its intended cinematic impact.

The solution, as is often the case with any other decisions about stereo, comes down to being aware of the potential pitfalls at both extremes, and after doing some experimentation, coming up with a set of rules that will make sense for the film. For instance, on *The Secret of the Wings*, the basic rule that drove all other decisions was to give volume to the characters in such a way as they would consistently appear at screen-space. Regardless of world scale, this was a very clear directive that had a huge influence in the stereoscopic design of the finished film. Given that the world of fairies leans toward the fantastical, our latitude to play with scale gave us more creative license than we might have had otherwise.

Camera Staging

CG animated films have the advantage that, given that by design they are created in 3D space, they are naturally suited for stereo. As a general rule, therefore, one can say that, if a shot

looks good in 2D, it will likely also look good, if not better, in 3D. However, in terms of camera placement and composition, there are a few things to be aware of when creating a shot for stereo:

- *Camera Height:* The closer a camera is to the ground, the more limited the options become available for stereo design. If a camera clips the ground too close to the lens, a shot can end up with too much ground detail coming out into audience space, or stereo settings that are too flat compared to surrounding shots. Usually raising the camera by a few units can maintain the intended composition while also opening up room for stereo

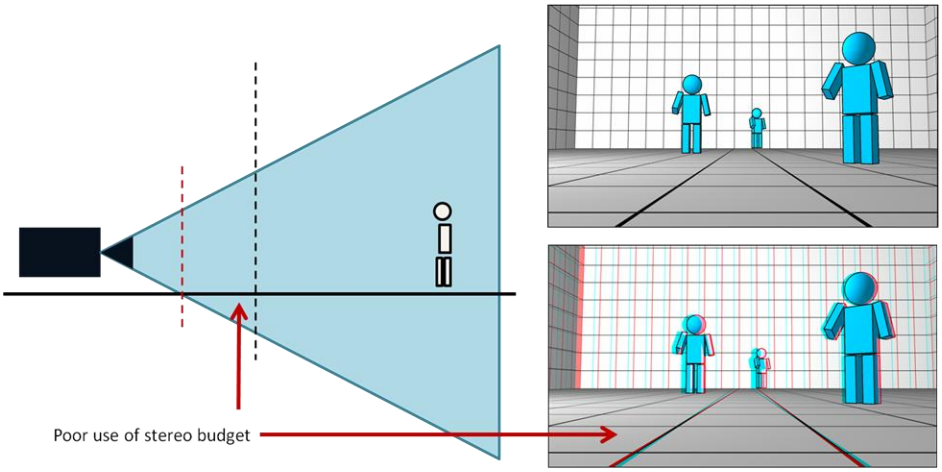

Figure 48: *Placing the camera too close to the ground limits stereo choices.* © Disney Enterprises, Inc.

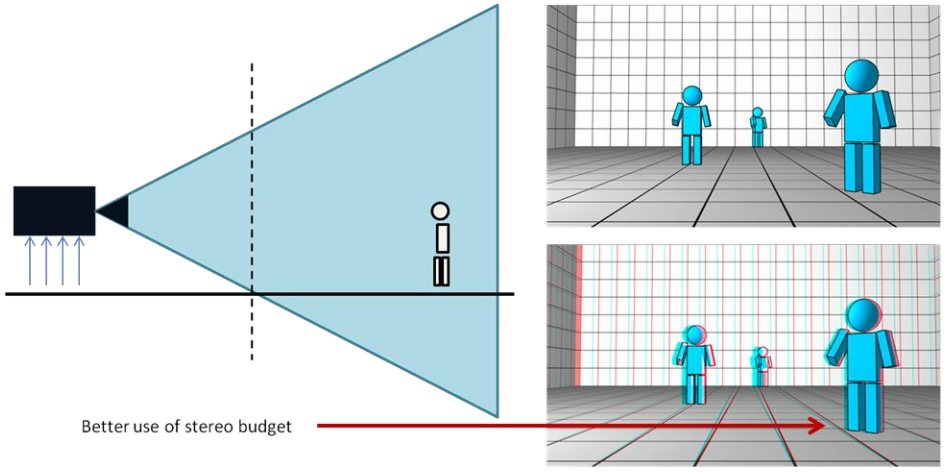

Figure 49: *Raising the camera slightly opens up options for stereo.* © Disney Enterprises, Inc.

- *Edge Composition:* Generally speaking, it is a good idea to keep prominent elements in a shot away from left/right edges of frame. This is especially true for locked shots.

Besides minimizing the potential for edge-of-frame violations, following this rule results in compositions that are more pleasing to look at in stereo

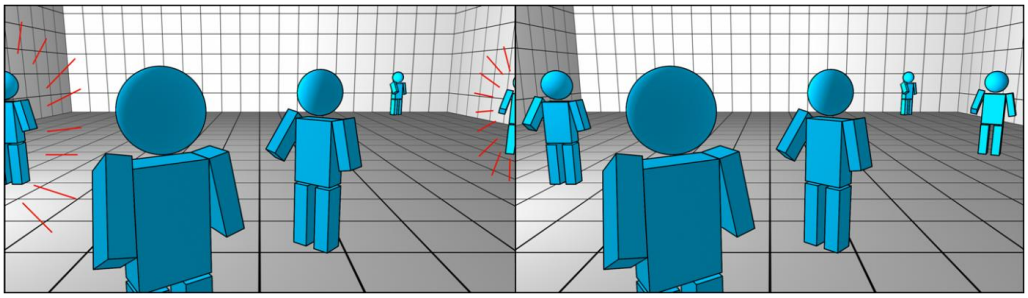

Figure 50: *Keeping edge-of frame clean helps create a more pleasing stereo shot.*

- *Stereo and Focal Length:* The type of lens through which a scene is shot not only affects cinematography; it also affects the quality of the resulting stereo. Broadly speaking, short lenses produce round, deep stereo, whereas long lenses yield stereo that is card-like and layered. The images below illustrate the results of the same scene shot with a 35mm lens vs. a 100 mm lens.

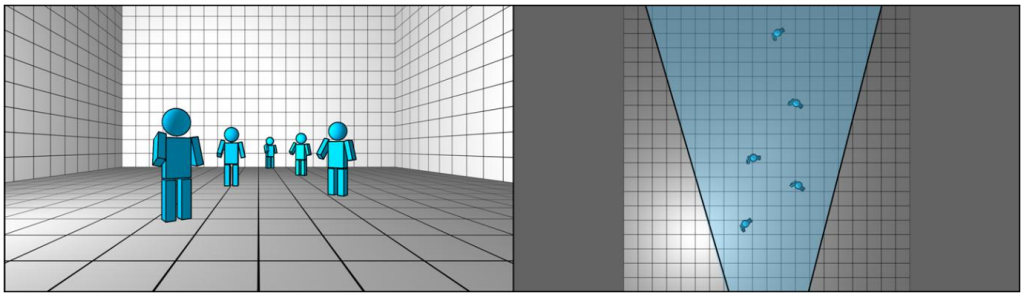

Figure 51: *Using a 35mm lens results in a wide and deep stereo space.*

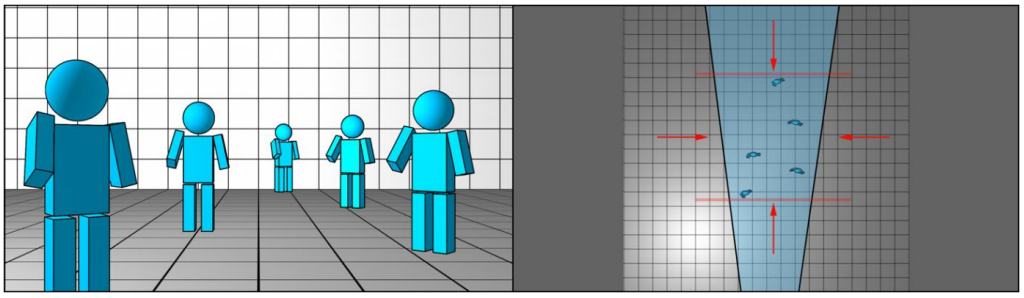

Figure 52: *Using a 100mm lens results in a more compressed, flattened stereo space.*

It is of course desirable to have as rich a stereo experience as possible, but oftentimes, the needs of a film are not always in line for what is best-suited for stereo. It is important, however, to know the limitations of the medium, and becoming intimately familiar with the effects of lensing for stereo is essential. Below are some rough lensing guidelines to keep in mind (Mendinburu, p.110):

Focal Length	Lens Type	Subject Type
21-35	Wide	Landscape
35-70	Normal	Street
70-135	Medium Telephoto	Portrait
135+	Telephoto	Sports

Figure 53: *Focal length ranges and some suggested subjects.* © Disney Enterprises, Inc.

It should be pointed out at this point that the use of animated zooms in a shot has the effect of taking a "round" world and rendering it flatter for stereo. While this may have a desired effect for a shot, countering this flattening effect could be achieved by means of a multi-rig set up.

- *Camera Motion:* Almost every kind of camera movement can be well-suited for stereo so long as the resulting motion is not too fast or too jarring. A camera that moves too quickly through a scene can be difficult for the human brain to be able to fuse into a cohesive image. Similarly, objects or characters that move too quickly across a frame may produce a strobing quality that adversely affects stereo. Motion blur can soften the degree of strobing, but ultimately, the ability for an audience to discern stereo may still be impacted adversely. For this reason, as a general rule, the pace of a film should weigh heavily in the overall stereo design planned from the outset. Below are different types of motion, and some general thoughts about their effect on stereo.

    - Rotation vs. Translation: When given the option between capturing action by rotating a camera vs. translating it in pace, translation results in parallax cues that make the shot easier to read in stereo. By contrast, camera rotations alone don't yield visual cues that communicate depth as clearly.

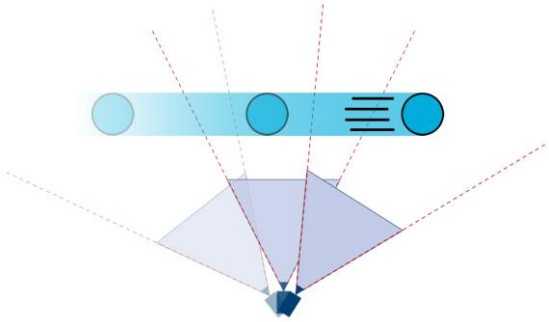

Figure 54: *Camera rotation does not generate parallax cues.*

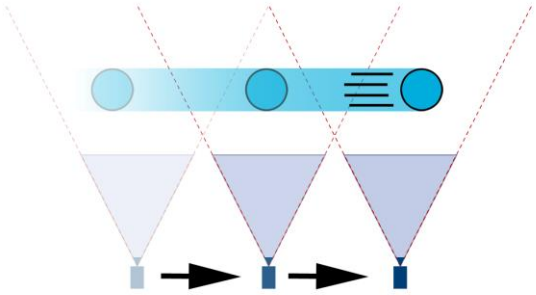

Figure 55: *Camera translation generates parallax cues useful for stereo.*

- *Object Motion Across Frame:* In the same manner that camera motion affects the resulting stereo experience, the way an object moves across a frame has a similar effect on stereo. To start with, an object that moves perpendicular to the camera will result in a left-to-right (or right-to-left) motion across the frame. The closer this object is to the camera, the faster it will appear to move across, and hence, the more difficult it will be for a viewer to fuse the resulting stereo pair into one. Angling the motion of the same object such that its entrance or exit from frame is closer to camera results in motion that is more interesting and arguably generates stereo cues that are easier to read for stereo.

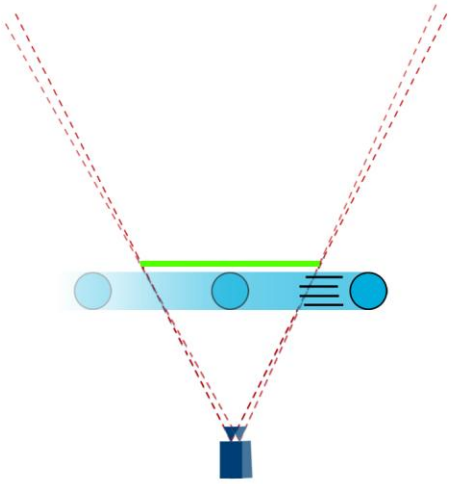

Figure 56: *Object translation perpendicular to camera is somewhat uninteresting.*

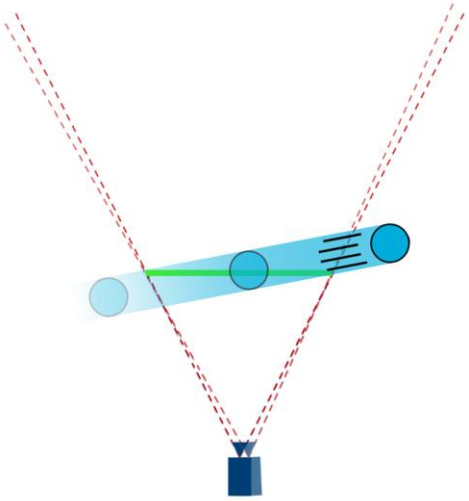

Figure 57: *Object translation slightly angled to camera results in easier to read motion.*

- *Object Motion Through Frame:* The movement of objects in Z, whether toward camera or away from camera, can yield optimal conditions for stereo. However, in designing this kind of shot it is important to keep in mind that how far an object moves into negative (or audience) space, must stay within the parameters that the stereographer has defined for a show. In other words, if the maximum amount of negative parallax that a particular sequence allows is -10 pixels, by the time the object in question moves out of frame, it should at most have a 10-pixel offset. In

order to achieve a good sense of depth throughout the shot, as well as hitting the required pixel offset parameters, it may be necessary to animate one or both convergence and interaxial separation.

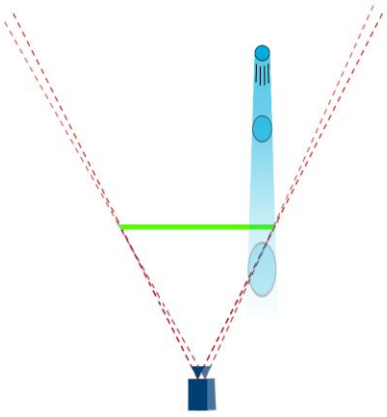

Figure 58: *Object translation in Z from off-screen toward infinity.*

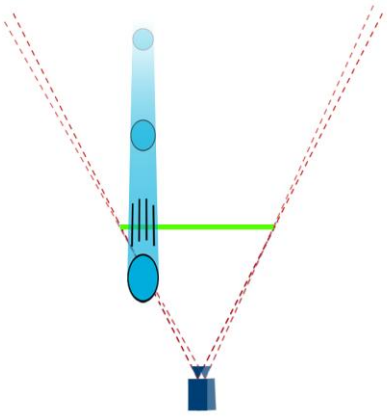

Figure 59: *Object translation in Z from far away toward camera.*

Finally, with regard to motion in Z, now that stereo has grown up beyond the gimmick domain, it is almost never a good idea to have objects/characters come directly toward camera beyond the depth parameters previously established.

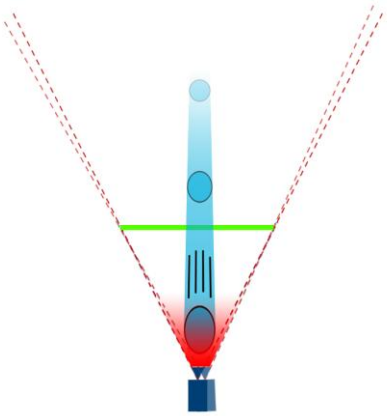

Figure 60: *As a general rule, avoid object translation in Z directly toward audience.*

Direct/Indirect Stereo: As we have established throughout this course, creating stereo is very much a matter of understanding pixel offsets. Specifically, setting a value for interaxial distance and convergence results in a stereo profile in which left-eye/right-eye objects will be offset by some pixel amount depending on where they are relative to the convergence plane. A stereo camera rig, therefore, should not only provide easy access for the setting of interaxial distance and convergence, it should also provide the means to evaluate what setting those two values will result in terms of left-eye/right-eye pixel offsets. In other words, a stereo rig should provide a means to pick a point in space, and output the corresponding pixel offset.

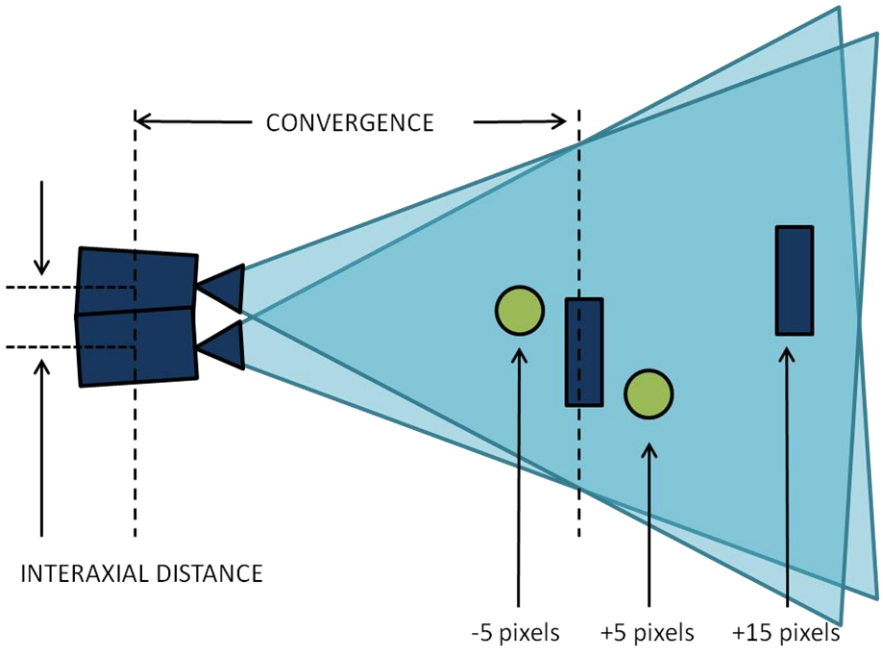

Figure 61: *Setting interaxial distance and convergence to create pixel offsets at arbitrary points within the camera's visual frustum.* © Disney Enterprises, Inc.

This mechanism of obtaining pixel offset values from the setting of stereo parameters (interaxial distance and convergence) is what in this paper we define as Direct Stereo, and a camera rig that works through direct control of these parameters is called a Direct Rig. Most stereographers become very comfortable with this method of working, and find the notion of thinking in terms of interaxial distance and convergence very intuitive. There is, however another way of thinking about stereo, and it is what we in this paper define as an Indirect Stereo.

The concept behind an Indirect Rig is very simple. Since it is possible to derive pixel offsets from the geometry resulting from setting interaxial and convergence, then it is also possible to work backward from pixel offsets and create the geometric conditions that would physically make those offsets possible. In this manner, a stereographer can look at a scene, define the pixel profile that he/she desires to obtain at both the positive and negatives ends of the spectrum, and let math do the rest.

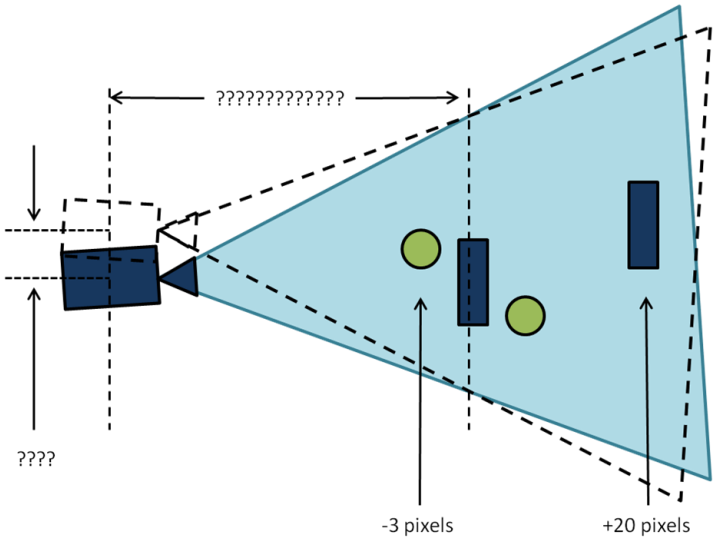

Figure 62: *Setting negative and pixel offsets to drive convergence and interaxial separation.*

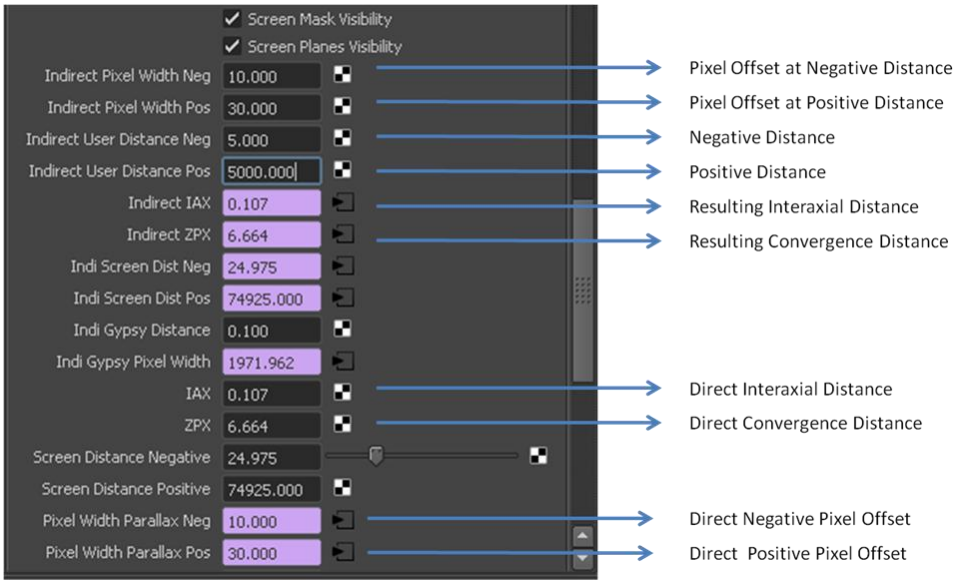

Figure 63: *Sample direct/indirect rig controls.*

It should be noted that the derivation of pixel offsets from interaxial and convergence settings is both dependent on the pixel-size of the output frame, and the stereo formulas that drive the stereo rigs. For a description of the Maya camera model, please see the Maya Stereoscopic Whitepaper listed in the References section.

Stereo Windows: The motivation behind the use of stereo windows is one of consistently creating stereo pairs that are as free of conflicts as is possible. As we have established throughout the course, a stereo pair consists of images resulting from cameras that are slightly offset by some distance, as shown below.

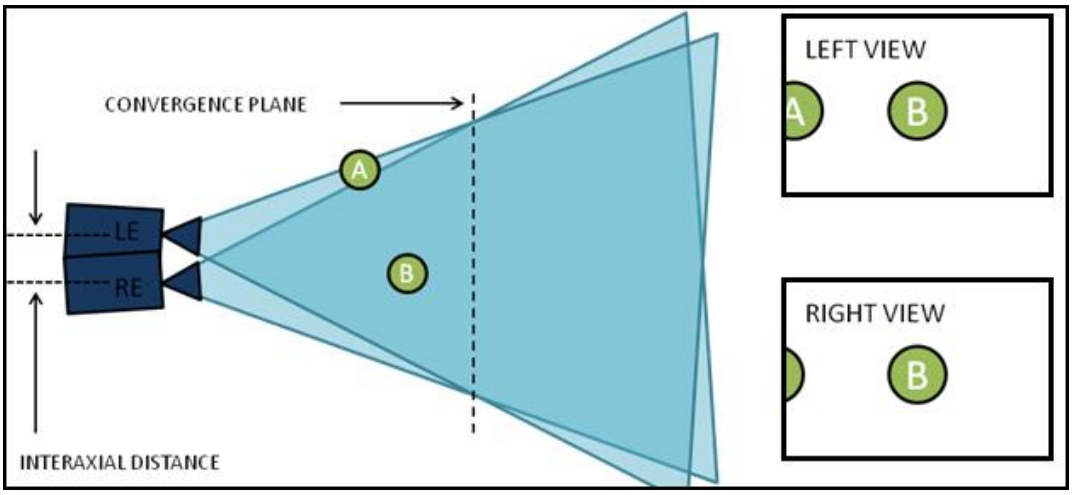

Figure 64: *Object A is more visible in the Left-Eye View compared to the Right-Eye View.*

You will notice from the figure above that, while object B is equally visible in both cameras, object A is more visible in the left view than in the right view. In fact, the wider the interaxial separation, the more object A will be hidden from one or both of the cameras (depending on how the stereo rig is set up). In this case, the right view sees less of object A than the left. This gap produces a phenomenon known *as retinal rivalry*. Our brains process retinal rivalries every day without any difficulties whenever one object partially occludes another. This occlusion provides a visual cue (i.e. the retinal rivalry) that helps us judge depth. However, retinal rivalries that occur when objects in front of the convergence plane are partially hidden from one of the cameras produce a kind of retinal rivalry that is difficult for the brain to process. This is exactly what is taking place in the image above. These kinds of retinal rivalries can be easily solved by introducing a masking element that blocks the extraneous detail that is being seen by

one eye. This masking element has come to be known as a stereo window or a floating window, and is illustrated below:

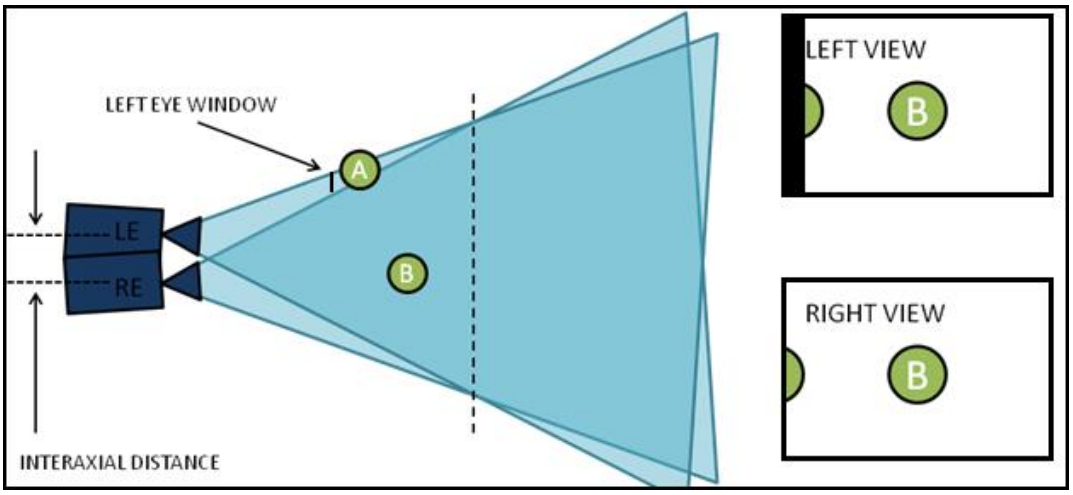

Figure 65: *The insertion of a mask that bridges the problem gap gets rid of the retinal rivalry.*

Lessons Learned on *The Secret of the Wings*

As outlined earlier in the course, the stereo process that we set up for *The Secret of the Wings* consisted of four stages:

- **PR3** Previs 3D
- **AN3** Animation 3D
- **CO3** HD Color 3D
- **BA3** HD Balance 3D

Over the course of nearly eighteen months allocated for stereo production, each of the 1163 shots that comprised the finished film had to go through each of these stages at least once, and undergo a strict review process that, in many instances required shots to through one or more stages multiple times.

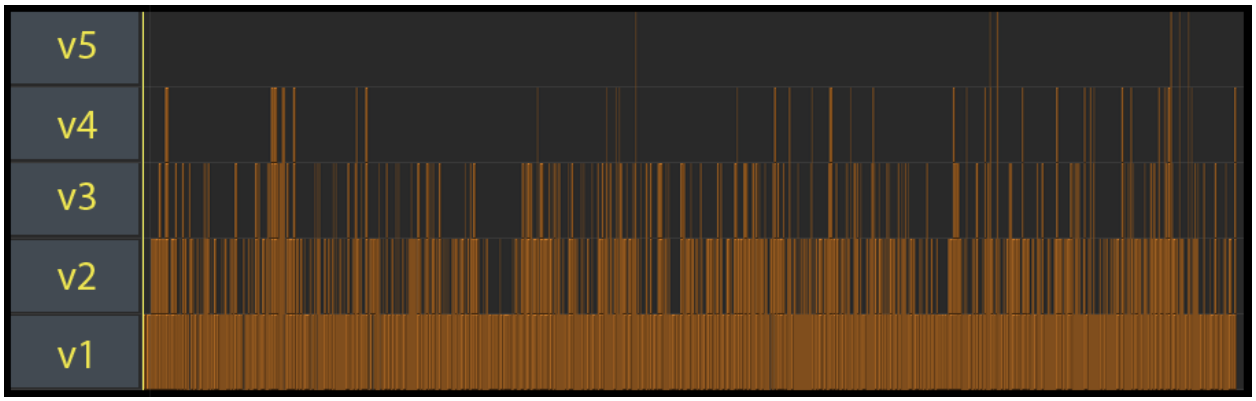

Figure 66: *Final Stereo Color Approvals for* ***The Secret of the Wings****. 47% of the film was finaled in one take, 36% in two takes, and the remaining 17% was approved in three or more takes.* © Disney Enterprises, Inc.

Given the workflow that we established for *The Secret of the Wings*, we solved the majority of the production challenges during the course of a shot's progression through the monoscopic pipeline. In this sense, the stereoscopic version of the film was spared from a lot of the approval cycle that resulted in the finished monoscopic film. However, the aggressive schedule that was set aside for stereo created its own set of challenges for each stage. Below is a summary of the production experience for each of the stages described above, as well as some general impressions about how each stage could benefit from improvements.

PR3: After setting the foundation for the stereo vision through our depth script, the Stereo Previs stage was instrumental in fleshing out how that vision would translate across entire sequences. The primary challenge we encountered at this stage was the fact that every shot needed to go through one person (namely the stereoscopic lead Vladimir Sierra) for a first camera pass. Despite the fact that during the previs stage, the scene files were at their lightest in complexity they would ever be, the sheer volume of shots that needed to be processed weekly resulted in a bottleneck that may have been otherwise avoided. In the future, implementing a performance-enhancing mechanism like geometry-caching at the PR3 level would allow us to work through footage more quickly.

AN3: Stereo Animation (AN3) approvals went fairly smoothly on *The Secret of the Wings*, given that, for a large portion of the film, scene layout did not change too drastically between the previs stage (PRE) and the animation stage (ANI). The design of the ANI stage was such that, when animation was approved, our animators would import the stereo rigs created at the PR3 stage, and create QuickTime movies that the stereo lead would then review. If the resulting stereo design adhered closely to its first incarnation during PR3, then a shot could simply proceed onto the Stereo Color stage. If, however, stereo had changed too drastically, then a new camera rig would need to be created using the latest animation files. Similar to PR3, the challenge at this stage came down to throughput of content by a stereo department of one. Luckily, unlike the PR3 stage, for AN3, the percentage of geometry files that needed to be processed was significantly smaller. Once again, an optimization in the way we work with dense geometry sets will make this stage go a lot more smoothly in future shows.

CO3: Because this stage was designed to be triggered after Color approval, as well as Animation approval, the primary function of this stage was to be a checkpoint for quality control. With every shot, the primary questions we needed to ask were as follows:

- Does stereo match the previously approved AN3 design?
- Does each stereo shot match the previously approved COL design for mono?
- Are there technical glitches that break stereo?

Once again, although not strictly a CO3 problem, we encountered that the gap between the approval of AN3 and the first version of CO3 did not give the stereo department sufficient opportunity to provide input that could prevent potentially unrecoverable problems. This was particularly the case with regards to the appearance of particle effects in stereo versus in mono. On *The Secret of the Wings*, for instance, every snow shot was essentially an FX shot. In the future, for an effects-heavy production, it would be very advantageous to have a Stereo Effects stage or something of the like where potential simulation issues could be addressed earlier rather than later.

BA3: This was our depth grading stage, which internally at DisneyToon Studios we came to call stereo balancing. While conceptually, balancing shots by means of pixel offsets is easy to implement, the biggest challenge for us in balancing *The Secret of the Wings* was setting up a mechanism that would allow us to do so across entire sequences. Ultimately, we implemented a compositing workflow in Nuke that would ingest entire sequences, and, through a simple interface, we created controls that would affect the positioning of both the right eye, and the right stereo window on a shot by shot basis. In this fashion, we could scrub across multiple shots, and by animating small pixel offset adjustments, we would be able to preview the flow of stereo continuity across entire sequences. Once we arrived at values that yielded pleasing results, we would send out entire sequences to the farm where they would be rendered and prepared for final output. Altogether, balancing the entire film took a little under a week with a total of four artists doing the bulk of the work.

Closing Thoughts

Throughout this course, we have highlighted the breadth of implications that setting up an efficient stereo workflow has on an existing pipeline. It is our hope that, in sharing the lessons we learned during the production of *The Secret of the Wings*, we have made a small contribution toward the furthering of stereography as an integral part of modern filmmaking. Also, it is our hope that the insights we gained in understanding this art form will be of use to the growing community of stereographers worldwide that will shape the future of stereo.

Given the relative infancy of stereo production at DisneyToon Studios, the stereo pipeline we've put in place is robust and is a natural extension of our existing filmmaking process. Most importantly, our stereo process leverages the resources we have in-house while maximizing the amount of creative control of each film created at DisneyToon Studios. While we continue to improve upon the suite of stereo tools to enhance our creative range, we are constantly evaluating and improving the production tools we have in place to track our films. The smoother we are able to integrate the stereo workflow with our constantly evolving studio, the better we will be able to deliver the kind of entertainment that our kids will want to share with their kids, and the generations that will come long after we are gone.

References

Autodesk, Inc. , **Stereoscopic Whitepaper**, 2008, www.autodesk.com. [Online]. http://images.autodesk.com/adsk/files/stereoscopic_whitepaper_final08.pdf

Michel Benoit, **Production Issues With 3D Content Targeting Cinema, TV and Mobile Devices**, 2009, 3D Stereo Media

Lenny Lipton, **Foundations of Stereoscopic Cinema, A Study in Depth**, 1982, Von Nostrand Reinhold Company, ed, ISBN: 0-442-24724-9

Bernard Mendiburu, **3D Movie Making**, 2009, Focal Press, ed, ISBN: 978-0-240-81137-6

Mike Seymour, Art of Stereo Conversion, 2012, www.vfxguide.com. [Online]. http://www.fxguide.com/featured/art-of-stereo-conversion-2d-to-3d-2012/